

Topological Paths, Cycles and Spanning Trees in Infinite Graphs

Reinhard Diestel

Daniela Kühn

Abstract

We study topological versions of paths, cycles and spanning trees in infinite graphs with ends that allow more comprehensive generalizations of finite results than their standard notions. For some graphs it turns out that best results are obtained not for the standard space consisting of the graph and all its ends, but for one where only its topological ends are added as new points, while rays from other ends are made to converge to certain vertices.

1 Introduction

This paper is part of an on-going project in which we seek to explore how standard facts about paths and cycles in finite graphs can best be generalized to infinite graphs. The basic idea is that such generalizations can, and should, involve the ends of an infinite graph on a par with its other points (vertices or inner points of edges), both as endpoints of paths and as inner points of paths or cycles.

To implement this idea we define paths and cycles topologically: in the space \overline{G} consisting of a graph G together with its ends, we consider arcs (homeomorphic images of $[0, 1]$) instead of paths, and circles (homeomorphic images of the unit circle) instead of cycles. The topological version of a spanning tree, then, will be a path-connected subset of \overline{G} that contains its vertices and ends but does not contain any circles.

Let us look at an example. The *double ladder* L shown in Figure 1 has two ends, and its two sides (the top double ray R and the bottom double ray Q) will form a circle D with these ends: in the standard topology on \overline{L} (to be defined later), every left-going ray converges to ω , while every right-going ray converges to ω' . Similarly, the edge vw forms a circle with the end ω' and the two right-going subrays of R and Q starting at v and w , respectively.




Figure 1: The double ladder L

Which subsets of \overline{L} would be topological spanning trees in L ? The ‘infinite comb’ consisting of R , the ends ω and ω' , and all the vertical edges of L would be one example; the arc $uR\omega Q\omega'Rv$ obtained from D by deleting the edge uv another. The ordinary spanning tree $R \cup Q + vw$ of L , however, would not qualify, because it fails to contain the ends ω and ω' . (And we cannot simply add the ends, since that would create infinite circles.)

When G is locally finite, then those of its ordinary spanning trees whose closure in \overline{G} qualifies as a topological spanning tree are precisely its *end-faithful* spanning trees (see Section 7). In [4] we showed that these are precisely the spanning trees of G whose fundamental cycles generate its entire cycle space (including infinite cycles). Thus, topological spanning trees are not merely natural objects to study in an infinite graph but came up as the solution to a problem: the problem of how to generalize a basic fact about finite spanning trees and cycles to infinite graphs.

When G is not locally finite, however, things are more complicated. The complications which arise require either restrictions to the notion of the cycle space that are needed in some cases but seem unnecessary in others, or a different topology on \overline{G} . The first of these approaches was followed in [5], while it is the purpose of this paper to explore the other. One of our first tasks will be to motivate our new topology on \overline{G} in terms of the problems indicated above, and this will be done in Section 3. However, there is yet another way to motivate that topology, independent of those problems, which we indicate now.

The double ladder L satisfies Menger’s theorem for ω and ω' : these ends can be separated by two vertices (such as v and w), and they are joined by the two independent arcs $\omega R\omega'$ and $\omega Q\omega'$. However, when we contract R to the edge uv (Fig. 2), the resulting graph G no longer contains two independent arcs between ω and ω' , although we still need two vertices to separate them. Our way to restore the validity of Menger’s theorem here will be to identify ω with u and ω' with v . Or put another way: we shall define the space \overline{G} not by adding ω and ω' to G as extra points and then applying the standard topology (see Section 2), but by choosing a topology

on G itself in which the left and right subray of Q converge to u and v , respectively. Then u and v are joined by the two arcs uv and Q (which together form a circle), and G again satisfies Menger's theorem.




Figure 2: A Menger problem for ω and ω'

More generally, the topological space for a graph G and its ends that we propose here will be the quotient space obtained from \overline{G} with its standard topology (in which all the ends are new points) by making all vertex-end identifications in situations as above. In this space, only ends that are not ‘dominated’ by a vertex (in the way ω is dominated by u in Fig. 2: there are infinitely many paths which join u to a ray in ω and which only meet in u) will be new points. As it happens (see [6]), these are precisely the ends of G that satisfy Freudenthal’s [9] original topological definition of an end.

We shall see later that our identification topology is not just an ad-hoc device to deal with problems such as the Menger example above. Roughly speaking, it is with this topology that standard finite results such as the generation of the cycle space by fundamental cycles can be generalized to the largest class of graphs that are not necessarily locally finite. But the example of Figure 2 already indicates why this is not unexpected: the identification topology on \overline{G} merely extends to vertex-end pairs what is already the case in the standard topology for pairs of ends of rays in G , namely, that two such points are to be identified if they cannot be finitely separated.

We have organized this paper as follows. In Section 2 we define the concepts to be used, in particular our topological versions of paths, cycles, and spanning trees, and introduce the topology on \overline{G} that is standard in the literature. In Section 3 we recall some results from [4] about topological spanning trees and the cycle space of locally finite graphs, and describe the obstructions that arise when we try to extend these results to graphs with infinite degrees. The identification topology motivated by these obstructions (as well as by the considerations above) is introduced in Section 4. In Section 5 we prove that topological spanning trees exist in all graphs in which their existence is not ruled out trivially by some obvious obstructions. As a spin-off of our methods we obtain that closed connected subsets of \overline{G} are path-connected. (This was unknown even for locally finite graphs under the

standard topology and is false in general for graphs with infinite degrees; see Section 2 for an example.) In Section 6 we prove our main results on topological cycles and spanning trees. These extend our locally finite results from [4] to a larger class of infinite graphs, which will be seen to be essentially largest possible. In Section 7, finally, we relate topological spanning trees to the existing literature on end-faithful spanning trees, and briefly address the general existence problem of topological spanning trees under the standard topology.

2 Basic concepts, and the standard topology

The terminology we use is that of [1]. A 1-way infinite path will be called a *ray*, a 2-way infinite path a *double ray*. The subrays of rays or double rays are their *tails*. Two rays in a graph G are *end-equivalent* if no finite set of vertices separates them in G . This is an equivalence relation on the set of rays in G ; its equivalence classes are the *ends* of G . We denote the set of ends of G by $\Omega(G)$. A vertex $v \in G$ is said to *dominate* an end ω if for some (and hence every) ray $R \in \omega$ there are infinitely many v - R paths in G that meet pairwise only in v ; such a set of paths is a v - R *fan*.

We shall freely view a graph either as a combinatorial object or as the topological space of a 1-complex. (So every edge is homeomorphic to the real interval $[0, 1]$, and the basic open neighbourhoods of a vertex x are the unions of half-open intervals $[x, z)$, one from every edge $[x, y]$ at x ; note that we do not require local finiteness here.) When E is a set of edges we let $\overset{\circ}{E}$ denote the union of their interiors, i.e. the set of all inner points of edges in E .

A homeomorphic image in a topological space X of the closed unit interval $[0, 1]$ will be called an *arc* in X ; a homeomorphic image in X of the unit circle is a *circle* in X ; and a homeomorphic image in X of the interval $[0, 1]$ is a *topological ray* in X . A continuous (but not necessarily injective) image of $[0, 1]$ is a *topological path*. If x and y are distinct points on an arc A , we write xAy for the subarc of A between x and y . Note that an arc inherits a linear ordering of its points from $[0, 1]$ (up to reflection). Given two sets $Y, Z \subseteq X$, we say that A is a Y - Z *arc* if one endpoint of A lies in Y , the other lies in Z , and the interior of A avoids $Y \cup Z$.

We shall frequently use the following lemma from elementary topology [11, p. 208].

Lemma 2.1 *Every topological path with distinct endpoints x, y in a Hausdorff space X contains an arc in X between x and y . □*

Our objects of study will be Hausdorff spaces \overline{G} consisting of a graph G and some or all of its ends. More precisely, we will either add all ends to G and endow this set with the standard topology, or add only those ends that correspond to the topological ends of G as a 1-complex. In the first case, the topology which \overline{G} induces on G will be the original 1-complex topology of G , while in the latter some rays may converge to vertices. In both cases, however, all the rays in an end ω will converge to a common point: either to ω (if $\omega \in \overline{G} \setminus G$), or to the unique vertex dominating ω .

Any circle D in \overline{G} will have the property that it contains every edge of which it contains an inner point. The set $C(D)$ of edges contained in D will be called its *circuit*. Since we intend to study the circles in \overline{G} combinatorially in terms of their circuits, it will be important that no two circles have the same circuit. To ensure this, we shall require that the topology on \overline{G} satisfies the following condition:

For every circle $D \subseteq \overline{G}$, the union $\bigcup C(D)$ of its edges is dense in D . (1)

Thus every circle D is the closure in \overline{G} of its circuit C , and is therefore uniquely determined by C .

Let us call a family $(C_i)_{i \in I}$ of circuits *thin* if no edge lies in C_i for infinitely many i , and let the *sum* $\sum_{i \in I} C_i$ of these circuits be the set of those edges that lie in C_i for an odd number of indices i . We now define the *cycle space* $\mathcal{C}(\overline{G})$ of \overline{G} as the set of sums of circuits in \overline{G} ; this is a subspace of the edge space of G just as in the finite case. In Section 6 we show that, for the topology considered in this paper, $\mathcal{C}(\overline{G})$ is closed also under infinite sums.

Finally, a *topological spanning tree* of \overline{G} is a path-connected subset T of \overline{G} that contains all the vertices and ends of \overline{G} , contains every edge of which it contains an inner point, and does not contain a circle. Note that T is closed in \overline{G} . Its subset $T \cap G$ is a subgraph of G but need not be connected. (However, topological spanning trees for which this is the case, ie. where $T \cap G$ is an ordinary spanning tree of G , may be of particular interest.) We write $E(T)$ for the set of all edges contained in T . For every edge $e = xy$ not in $E(T)$, Lemma 2.1 ensures that T contains a (unique) arc between x and y , and so $T \cup e$ contains a unique circle D . We call every such D a *fundamental circle*, and its circuit C_e a *fundamental circuit* of T .

We will use the following standard lemma about infinite graphs; the proof is not difficult and is included in [3, Lemma 1.2].

Lemma 2.2 *Let U be an infinite set of vertices in a connected graph G . Then G contains either a ray R with infinitely many disjoint U - R paths or a subdivided star with infinitely many leaves in U . \square*

Let X be a Hausdorff space. We denote the closure of a set $Y \subseteq X$ by $\text{cl}(Y)$. Given a topological ray R in X , an infinite sequence x_1, x_2, \dots of distinct points, and for all $x_i \notin R$ disjoint x_i - R arcs Q_i such that the sequence consisting of the preimages under the homeomorphism $[0, 1) \rightarrow R$ of the endpoints on R of these paths and the preimages of all x_i on R converges to 1, we call the union of R with all the Q_i a *topological comb* in X with *back* R and *teeth* x_1, x_2, \dots (including the x_i on R). A *topological* \aleph_0 -*star* in X is any union S of \aleph_0 arcs in X meeting pairwise exactly in their first point. This point is the *centre* of S , the other endpoints of those arcs are its *leaves*. Lemma 2.2 thus states that, for every infinite set U of vertices, G contains either a topological comb with teeth in U (and back a ray) or a topological \aleph_0 -star with leaves in U .

The following lemma generalizes Lemma 2.2 to arbitrary path-connected Hausdorff spaces. We omit its straightforward proof, which is similar to that of Lemma 2.2.

Lemma 2.3 *Let U be an infinite set of points in a path-connected Hausdorff space X . Then X contains either a topological comb with all its teeth in U or a topological \aleph_0 -star with all its leaves in U . \square*

A rooted (ordinary) spanning tree T of G is *normal* if the endvertices of every edge of G are comparable in the tree order induced by T ; see [1]. Countable connected graphs are easily seen to have normal spanning trees, but not all uncountable ones do; see [8] for details. For our purposes, we shall need the following existence theorem of Halin [10]:

Lemma 2.4 *Every connected graph containing no subdivision of K_{\aleph_0} has a normal spanning tree. In particular, every connected graph in which every end is dominated by at most one vertex has a normal spanning tree.*

The usefulness of normal spanning trees is based on the following easy lemma:

Lemma 2.5 [7] *Let $x_1, x_2 \in V(G)$, and let T be a normal spanning tree of G . For $i = 1, 2$ let $[x_i]$ denote the path in T that joins x_i to the root of T . Then $[x_1] \cap [x_2]$ separates x_1 from x_2 in G .*

We now define the topology TOP on \overline{G} that is standard in the literature for $\overline{G} = G \cup \Omega(G)$. We refer to [2] and, especially, Polat [13] for more background on TOP. Consider a finite set $X \subseteq V(G) \cup E(G)$.

For every end ω of G there is exactly one component C of $G - X$ that contains a tail of every ray in ω ; we say that ω *belongs* to C . Ends or vertices belonging to different components of $G - X$ are *separated* by X . When y is either an end or a vertex in $G - X$, we write $C_G(X, y)$ for the component of $G - X$ to which y belongs, and $E_G(X, y)$ for the set of edges that either join $C_G(X, y)$ to vertices in X or else are edges in X incident with $C_G(X, y)$. The ends of C correspond naturally to the ends of G belonging to C , and we do not normally distinguish between them. Finally, we define

$$\widehat{C} := \widehat{C}_G(X, y) := C \cup \Omega(C) \cup E'(X, y) \subseteq \overline{G},$$

where $E'(X, y)$ is the union of any maximal set of internally disjoint half-edges $(z, v) \subset e$ with $e \in E(X, y)$, $z \in \mathring{e}$, and $v \in V(C)$. (Thus $E'(X, y)$ contains two half-edges for every edge $e \in X$ joining two vertices of C and one for every other edge in $E(X, y)$.) When U is a union of components of $G - X$, we similarly write \widehat{U} for any union of sets \widehat{C} , one for each component $C \subseteq U$.

Now let TOP denote the topology on \overline{G} that is generated by the open sets of the 1-complex G and all sets of the form $\widehat{C}_G(S, \omega)$ with S a finite set of vertices. Thus for each end ω , the sets $\widehat{C}_G(S, \omega)$ are the basic open neighbourhoods of ω . It is not difficult to check [4] that TOP satisfies all our earlier requirements on \overline{G} . In particular, TOP satisfies (1), so the circles in \overline{G} correspond bijectively to its circuits. When G is locally finite and connected, \overline{G} is compact under TOP.

We close this section with a general observation concerning TOP that we have found surprisingly difficult to prove:

Theorem 2.6 *When G is locally finite, every closed connected subset of \overline{G} is path-connected.*

Note that \overline{G} is locally path-connected and so every open connected subset of \overline{G} is path-connected (even if G is not locally finite). Theorem 2.6 is a special case of Theorem 5.3, to be proved below. We expect that it extends to sets X that are neither open nor closed, but our proof of Theorem 5.3 depends on the assumption that X is closed.

When G has vertices of infinite degree, \overline{G} can have closed connected subsets that are not path-connected. For example, if G is obtained from a ray R by adding a new vertex x and infinitely many x - R paths of length 2 that meet only in x , then deleting from these paths the edges incident with R results in a subspace of \overline{G} that is connected (because every neighbourhood of the unique end contains a tail of R and almost all the neighbours of x) but not path-connected.

3 Cycles and trees in the standard topology

Let G be a locally finite graph, and consider $\overline{G} := G \cup \Omega(G)$ with TOP. Here are some results concerning infinite cycles in \overline{G} that we would like to generalize sensibly to graphs that are not locally finite.

Theorem 3.1 *The fundamental circuits of any topological spanning tree of \overline{G} span its cycle space $\mathcal{C}(\overline{G})$.*

This is the locally finite case of Theorem 6.1 below. It was proved in [4] for end-faithful spanning trees of G , ie. for topological spanning trees $T \subseteq \overline{G}$ such that $T \cap G$ is connected (cf. Theorem 7.3).

Cycle-cut orthogonality in finite graphs generalizes too:

Theorem 3.2 [4] *$\mathcal{C}(\overline{G})$ consists of precisely those sets of edges that meet every finite cut in an even number of edges.*

Nash-Williams [12] proved that the edge set of any graph (not necessarily locally finite) decomposes into finite circuits if (and only if) the graph has no odd cut. If the entire edge set $E = E(G)$ is an element of $\mathcal{C}(\overline{G})$, then this implies with Theorem 3.2 that E is a sum of disjoint (finite) circuits. For arbitrary elements of $\mathcal{C}(\overline{G})$ this is no longer clear (even admitting infinite circuits in the sum), since the graph on $V(G)$ induced by an infinite circuit is just a disjoint union of rays, which has lots of odd cuts. The fact that arbitrary elements of $\mathcal{C}(\overline{G})$ have disjoint-circuit decompositions is one of the main results of [5]:

Theorem 3.3 [5] *Every element of $\mathcal{C}(\overline{G})$ is a union of disjoint circuits.*

How do the above results generalize to graphs that are not locally finite? Consider the plane graph G shown in Figure 3. There, the finite circuits




Figure 3: The edge xy is the sum of all the facial cycles.

bounding a face form a family in which the edge xy occurs in one circuit and every other edge occurs in two circuits. So these circuits sum to the single

edge xy —which would thus be an (unwelcome) element of $\mathcal{C}(\overline{G})$ according to the definition given in Section 2. (The unwelcomeness is not just a matter of taste: of the above three theorems only Theorem 3.2 generalizes to this graph.)

In [5], we dealt with this phenomenon by restricting the notion of the cycle space, disallowing sums in which infinitely many terms share a vertex. With this restriction, Theorem 3.3 generalizes to arbitrary infinite graphs, while Theorem 3.2 adapts with a trivial modification. But Theorem 3.1 no longer works for all topological spanning trees: in the graph of Figure 4, all fundamental circuits of the topological spanning tree T contain x , but no finite sum of these circuits generates the infinite circuit $E(RyQ)$.




Figure 4: The circuit $E(RyQ)$ is not a sum of fundamental circuits.

This problem is not easily overcome just by allowing more sums in the definition of $\mathcal{C}(\overline{G})$. Indeed, any sum $\sum C_e$ of fundamental circuits yielding $E(RyQ)$ would have to be over precisely the edges $e \in R$, because these are the edges of RyQ that are not in T . But clearly $\sum_{e \in R} C_e = E(xyR) \neq E(RyQ)$, no matter whether this sum is a legal element of $\mathcal{C}(\overline{G})$ or not.

As soon as we identify the end ω with the vertex x dominating it, however, the problem disappears: now T is no longer a topological spanning tree (because xyR is now a circle contained in T), but $T - xy$ is a topological spanning tree, and its fundamental circuits generate all circuits, including $E(RyQ)$.

In this paper we show that, for all graphs G not containing the trivial obstruction of Figure 3, identifying every end with the (unique) vertex dominating it yields the ‘right’ space for our desired generalization of Theorems 3.1–3.3: we prove that all three theorems continue to hold in this space \tilde{G} , even with the original (unrestricted) definition of the cycle space that includes arbitrary sums of thin families of circuits.

Because of the unavoidable problems associated with Figure 3, we will only consider graphs which satisfy the following condition:

No two vertices are joined by infinitely many independent paths. (2)

(Some additional motivation for why such graphs may be interesting will be given at the end of this section.)

We remark that all the blocks B of a graph satisfying (2) are countable. Indeed, by Lemma 2.4, B has a normal spanning tree T , and this must be locally finite: if t is a vertex of infinite degree in T , then its finite down-closure $\lceil t \rceil$ separates B into infinitely many components (Lemma 2.5), and by the 2-connectedness of B infinitely many of these send edges to the same two vertices in $\lceil t \rceil$. Then these two vertices are joined by infinitely many independent paths, contradicting (2).

Although all our results will technically be true for any graph satisfying (2), to make them interesting we may wish to impose the following stronger condition:

No two vertices are joined by infinitely many edge-disjoint paths. (3)

(Figure 5 shows that this is indeed stronger, ie. that (2) does not imply (3). Note that by the (straightforward) infinite analogue of Menger's theorem, condition (2) (resp. condition (3)) implies that every pair of non-adjacent vertices of G can be separated by finitely many vertices (resp. edges).)




Figure 5: A graph that satisfies (2) but violates (3)

The possible justification for imposing (3) lies in the fact that for graphs G satisfying (2) but not (3) our quotient space \tilde{G} may contain circles consisting only of ends and vertices—in which case (1) fails, different circles may have the same circuit, and there may be no topological spanning tree. In the

proof of the following lemma we exhibit such a graph, which is essentially the graph of Figure 5 with x and y identified. See the start of Section 4 for a formal definition of \tilde{G} if desired.

Proposition 3.4 *There exists a countable graph G satisfying (2) for which \tilde{G} contains a circle consisting only of vertices and ends.*

Proof. Consider the binary tree T_2 whose vertices are the finite 0–1 sequences and where each sequence is adjacent to its two one-digit extensions. The ends of T_2 correspond to the infinite 0–1 sequences, which we view as binary expansions of the reals in $[0, 1]$. Our aim is to turn this Cantor set into a copy of $[0, 1]$ by identifying the pairs of ends that correspond to the same rational $q \in [0, 1]$, ie. by identifying every two ends of the form $s1000\dots$ and $s0111\dots$ for some $s \in T_2$. To achieve this identification, we join the vertex s to every such pair of ends by a couple of fans, so that in \tilde{G} these ends will both get identified with s , and hence with each other.

Formally, we join each finite sequence $s \in T_2$ to all sequences of at least $|s| + 2$ digits that begin with $s1$ and thereafter contain only 0s, and to all sequences of at least $|s| + 2$ digits that begin with $s0$ and thereafter contain only 1s. Finally, we add a new vertex x joined to all sequences consisting only of 0s or only of 1s (Figure 6).




Figure 6: A graph whose ends form a circle in the identification topology

Any two vertices of this graph G are separated by the finite vertex set consisting of x and their common initial segments, so the graph satisfies (2). It is easily checked that mapping 0 and 1 to x and every other element of $[0, 1]$ to its corresponding end or identified pair of ends is a homeomorphism between $[0, 1]$ with 0 and 1 identified and the set of all vertices and ends in \tilde{G} (after identification). \square

Since any topological spanning tree of a graph must contain all its vertices and ends, every \tilde{G} as in Proposition 3.4 fails to have a topological spanning tree. Thus:

Corollary 3.5 *There exists a connected countable graph G satisfying (2) such that \tilde{G} has no topological spanning tree.*

We shall prove below that (3), unlike (2), suffices to imply (1) for \tilde{G} (Section 4), and that all such \tilde{G} have topological spanning trees (Section 5).

We close this section with an observation that may lend some unexpected relevance to graphs satisfying (3). In every infinite graph, being linked by infinitely many edge-disjoint paths is an equivalence relation on the vertex set. Now it may be of interest to study the quotient graph obtained by identifying each equivalence class into one point. (For example, this is a central tool in Nash-Williams's celebrated proof of his cycle decomposition theorem [12].) All such quotient graphs satisfy (3).

4 The identification topology

In this section, we first define the identification topology ITOP more formally, and then prove some basic facts about it.

Let $G = (V, E)$ be an infinite graph, fixed throughout this and the next section. Put $\Omega(G) =: \Omega$, and let $\Omega' \subseteq \Omega$ denote the subset of those ends that are not dominated by any vertex. (We remark that these are precisely the ends of G that correspond to its ends in the topological sense of Freudenthal [9]; see [6].) Given a vertex $v \in G$, we write Ω_v for the set of ends it dominates. As always in the rest of the paper when we consider the topology ITOP to be defined now, we assume that

every end of G is dominated by at most one vertex. (†)

Note that (2) implies (†). Let \tilde{G} be the quotient space obtained from \overline{G} endowed with TOP by identifying every vertex with all the ends it dominates. When two points $x, y \in \overline{G}$ are thus identified, we call them *equivalent*. As usual, we write $\pi : \overline{G} \rightarrow \tilde{G}$ for the canonical projection sending each point of \overline{G} to its equivalence class. The (identification) topology of \tilde{G} , which is the finest topology on the set \tilde{G} that makes π continuous, will be denoted by ITOP. Recall that a set $N \subseteq \tilde{G}$ is open if and only if $\pi^{-1}(N)$ is open in \overline{G} , or equivalently if and only if $N = \pi(U)$ for some open set $U \subseteq \overline{G}$ that is closed under equivalence.

By (†), no two vertices of G are equivalent, so the identification does not alter V . We may thus view $G \cup \Omega'$ as the point set of \tilde{G} , and in particular denote every equivalence class which is not a singleton by the unique vertex it contains. For the rest of this paper,

\overline{G} will always denote the space of $G \cup \Omega$ endowed with TOP;

\tilde{G} will always denote the space of $G \cup \Omega'$ endowed with ITOP.

Note that if no end of G is dominated (in particular, if G is locally finite) then \overline{G} and \tilde{G} coincide.

Let us collect some facts about \overline{G} and \tilde{G} . Recall that if G is locally finite, then \overline{G} ($= \tilde{G}$) is compact. When G has a vertex of infinite degree, however, then neither \overline{G} nor \tilde{G} is compact. For example, if G is obtained from a ray R by adding a new vertex x and infinitely many x - R edges, then the set $N \subseteq \tilde{G}$ obtained from \tilde{G} by deleting a closed interval I_e in the interior of every x - R edge e is open. But N can be combined with an open interval around each I_e to an open cover of \tilde{G} that has no finite subcover. However, if G is 2-connected and satisfies (2), then the subset $V \cup \Omega'$ of \tilde{G} is compact. (By the remark following (2), G has a locally finite normal spanning tree, which can be used to show that $V \cup \Omega$ is a compact subset of \overline{G} (see Section 7). But every open cover of $V \cup \Omega' \subset \tilde{G}$ corresponds via π^{-1} to an open cover of $V \cup \Omega \subset \overline{G}$, and therefore has a finite subcover.)

Using straightforward topological arguments one can show that, in \tilde{G} just as in \overline{G} , every arc whose endpoints are vertices or ends, and similarly every circle, includes every edge of which it contains an inner point.

Lemma 4.1 *For every vertex $v \in G$, the set Ω_v of ends dominated by v is closed in \overline{G} .*

Proof. Consider any point $x \in \text{cl}(\Omega_v)$; we show that $x \in \Omega_v$. Clearly x is an end; pick a ray $R \in x$. For every finite set $S \subseteq V(G - v)$, some end $\omega \in \Omega_v$ belongs to $C(S, x)$. Since v sends an infinite fan to ω it must lie in $C(S, x)$, and so $G - S$ contains a v - R path. Choosing as S the vertex sets of suitable initial segments R' of R together with any v - R' paths already found, we may thus construct an infinite v - R fan inductively. Hence $x \in \Omega_v$, as claimed. \square

Recall that a *cut* in G is the set F of all edges between the two (non-empty) classes of some bipartition of V .

Lemma 4.2 *Let $F \subseteq E$ be a finite cut in G and let D be a component of $G - F$. Then \overline{D} is closed under equivalence, ie. every set of the form $\pi(\hat{D})$*

is open in \tilde{G} . If $x, y \in V \cup \Omega'$ belong to different components of $G - F$, then every x - y arc A in \tilde{G} contains an edge from F .

Proof. Let S be the set of all endvertices outside D of edges from F . Then S is finite and D is a component of $G - S$. Since there are only finitely many edges between S and D , no end belonging to D is dominated by a vertex outside D , and all ends dominated by vertices in D belong to D . Thus \overline{D} is closed under equivalence, and hence $\pi(\hat{D})$ is open in \tilde{G} .

To prove the second claim, suppose that A does not contain an edge from F . Then A avoids \hat{F} . Let N_x be a set of the form $\hat{C}(F, x)$ and let N_y be a union of sets \hat{D} , one for every component $D \neq C(F, x)$ of $G - F$, such that N_x and N_y are disjoint. Then both N_x and N_y are closed under equivalence. Hence $\pi(N_x)$ and $\pi(N_y)$ are non-empty disjoint open subsets of \tilde{G} whose union contains A , a contradiction to the connectedness of A . \square

Lemma 4.3 *If G satisfies (3) and $x, y \in V \cup \Omega'$ are distinct, then x can be separated from y in G by finitely many edges.*

Proof. The lemma clearly holds when both x and y are vertices: just take as the separator the edges of any maximal set of edge-disjoint x - y paths. Now suppose that $x \in \Omega'$, and let S be a finite set of vertices separating x from y in G . As no vertex dominates x , there is for every $z \in S$ a finite set S_z of vertices in $C(S, x)$ that separates z from x in the subgraph of G induced by z and all the vertices in $C(S, x)$. Then $S' := \bigcup_{z \in S} S_z$ separates x from S in G . As G satisfies (3), there is a finite set $F \subseteq E$ separating S from S' in G . Then F also separates x from y , as desired. \square

In Proposition 3.4 we constructed a graph G that satisfies (2) but for which \tilde{G} contains a circle whose circuit is empty. The following result, which by continuity is an immediate consequence of Lemmas 4.2 and 4.3, shows that such circles cannot occur if G satisfies (3). Indeed, all graphs G satisfying (3) also satisfy (1) under ITOP:

Corollary 4.4 *Suppose that G satisfies (3), and let $x, y \in V \cup \Omega'$. Then for every x - y arc A in \tilde{G} the union of all edges contained in A is dense in A . Similarly, for every circle D in \tilde{G} the union $\bigcup C(D)$ of its edges is dense in D .* \square

The following lemma is an analogue of Lemma 4.2 for finite vertex separators. As a corollary of this lemma we obtain a weakening of Corollary 4.4 for arbitrary graphs G satisfying (†): for every arc A in \tilde{G} the set of its points in G , ie. of its non-ends, is dense in A .

Lemma 4.5 *Let S be a finite set of vertices of G . Then for every component D of $G - S$, every set of the form $\pi(\widehat{D}) \setminus S$ is open in \widetilde{G} . If $x, y \in V \cup \Omega'$ belong to distinct components of $G - S$ then every x - y arc in \widetilde{G} meets S .*

Proof. By Lemma 4.1, every set of the form $\widehat{D} \setminus \bigcup_{s \in S} \Omega_s$ is open in \overline{G} , and it is clearly closed under equivalence. So its image $\pi(\widehat{D}) \setminus S$ in \widetilde{G} is open in \widetilde{G} .

For the second part of the lemma one shows as in the proof of Lemma 4.2 that any x - y arc avoiding S cannot be connected (a contradiction). \square

Corollary 4.6 *For every arc A in \widetilde{G} the set $A \cap G$ is dense in A .* \square

Our next aim is to prove that \widetilde{G} is Hausdorff. In the proof of this result we will use a normal spanning tree of G , which we know to exist by (\dagger) and Lemma 2.4.

We need some more notation. Given a rooted tree T and a vertex $t \in T$, we write $[t]$ for the unique path in T that joins t to the root. We say that t' lies *above* t if $t \in [t']$, and denote by $[t]$ the subtree of T induced by all vertices above t (including t itself). Now suppose that T is a normal spanning tree of a graph G . Using Lemma 2.5, one easily shows that every end ω of G contains exactly one ray $R_\omega \subseteq T$ starting at the root of T ; in particular, disjoint rays of T belong to distinct ends of G . (So normal spanning trees are end-faithful; see Section 7.)

Given a vertex x on R_ω , we write xR_ω for the subray of R_ω induced by all its vertices above x , and $\hat{x}R_\omega$ for the ray $xR_\omega - x$. We say that an end ω of G lies *above* a vertex $x \in G$ if all its rays have a tail in $[x]$; by Lemma 2.5, this is the case if and only if x lies on R_ω . Note that if x and y are neighbours on R_ω and y lies above x , then $\overline{C}([x], \omega)$ consists of the subgraph of G induced by $[y]$ together with all the ends of G that lie above y in T . Note also that an end ω lies above any vertex that dominates it, and that a vertex dominates ω if and only if it has a neighbour above z for every vertex $z \in R_\omega$.

Given a vertex $x \in G$, any union of half-edges $[x, z) \subset e$, one for every edge e at x , will be called an *open star around x* .

Theorem 4.7 *\widetilde{G} is Hausdorff.*

Proof. We have to show that for every two distinct points of $x, y \in \widetilde{G}$ there are disjoint open neighbourhoods around x and y . We only consider the case that both x and y are vertices of G ; the other cases are trivial or similar.

So we have to find disjoint open sets N_x and N_y in \overline{G} such that $x \in N_x$, $y \in N_y$, and both N_x and N_y are closed under equivalence. (Then $\pi(N_x)$ and $\pi(N_y)$ are disjoint open neighbourhoods of x and y in \widetilde{G} .) Let us first construct N_x .

We may assume that G is connected; then by Lemma 2.4 it has a normal spanning tree T . Given an end ω of G , let R_ω denote the unique ray in T that belongs to ω and starts at the root t_0 of T .

For every end $\omega \in \Omega_x$ define vertices t_ω and s_ω on R_ω as follows. Since $\omega \notin \Omega_y$ and Ω_y is closed in \overline{G} (Lemma 4.1), there exists a vertex $t_\omega \in \overset{\circ}{x}R_\omega$ such that neither y nor an end from Ω_y lies above t_ω in T . Since no vertex in $[t_\omega] \setminus \{x\}$ dominates ω , there exists a vertex $s_\omega \in \overset{\circ}{t_\omega}R_\omega$ such that G contains no edge between $\lfloor s_\omega \rfloor$ and $\lceil t_\omega \rceil \setminus \{x\}$.

Let N^1 be the union of an open star around x and sets of the form $\widehat{C}(\lceil s_\omega \rceil, \omega)$, one for each $\omega \in \Omega_x$. Then N^1 is open in \overline{G} and contains $\Omega_x \cup \{x\}$. If N^1 is closed under equivalence we set $N_x := N^1$.




Figure 7: Constructing N_x in the proof of Theorem 4.7

So suppose that N^1 is not closed under equivalence. N^1 contains every end lying above any of its vertices other than x , and thus N^1 contains all ends dominated by vertices in N^1 . Hence there must be an end τ' in N^1 that is dominated by a vertex outside N^1 . For every such τ' there exists an $\omega \in \Omega_x$ such that τ' lies above s_ω . For every $\omega \in \Omega_x$ let Z_ω^1 be the

set of all vertices $z \notin N^1$ dominating some end in $\widehat{C}([s_\omega], \omega) \subseteq N^1$. Then $Z_\omega^1 \subseteq [s_\omega] \setminus [t_\omega]$, by the choice of s_ω . Let Ω_ω^1 be the set of all ends outside N^1 dominated by vertices in Z_ω^1 . Let τ be any end in Ω_ω^1 , dominated by $z \in Z_\omega^1$, say. Then τ lies above z , and z also dominates an end $\tau' \in N^1$ above s_ω . As z is the only vertex dominating τ , no vertex in $[t_\omega]$ dominates τ . Hence there is a vertex $s_\tau \in {}^z R_\tau \subseteq t_\omega R_\tau$ such that G contains no edge between $[s_\tau]$ and $[t_\omega]$. Let N^2 be the union of open stars around the vertices in $\bigcup_{\omega \in \Omega_x} Z_\omega^1$ and sets of the form $\widehat{C}([s_\tau], \tau)$, one for each $\tau \in \bigcup_{\omega \in \Omega_x} \Omega_\omega^1$.

Note that $N^1 \cup N^2$ contains all points that are equivalent to points in N^1 . Thus as before, $N^1 \cup N^2$ is closed under equivalence if no end in N^2 is dominated by a vertex z outside $N^1 \cup N^2$. If such vertices z exist, then we extend $N^1 \cup N^2$ further by adding an open set $N^3 \subseteq \overline{G}$ which contains all ends dominated by such vertices z and open stars around these vertices. Thus $N^1 \cup N^2 \cup N^3$ contains all points that are equivalent to points in $N^1 \cup N^2$. We continue in this fashion for (at most) ω steps and put $N_x := N^1 \cup N^2 \cup \dots$. Then N_x is closed under equivalence and it consists of an open star around x together with a subset of $\bigcup_{\omega \in \Omega_x} \widehat{C}([t_\omega], \omega)$.

Similarly we construct an open set $N_y \subseteq \overline{G}$ which contains y and is closed under equivalence, and which consists of an open star around y together with a subset of $\bigcup_{\omega' \in \Omega_y} \widehat{C}([t_{\omega'}], \omega')$ (where $t_{\omega'}$ is a vertex on $R_{\omega'}$ such that neither x nor an end from Ω_x lies above $t_{\omega'}$).

It remains to show that N_x and N_y are disjoint. First note that $x \notin N_y$ since for each $\omega' \in \Omega_y$, x does not lie above $t_{\omega'}$ in T . Similarly, $y \notin N_x$. Let us now show that for all $\omega \in \Omega_x$ and $\omega' \in \Omega_y$ we have $\widehat{C}([t_\omega], \omega) \cap \widehat{C}([t_{\omega'}], \omega') = \emptyset$. Every vertex and every end in $\widehat{C}([t_\omega], \omega)$ lies above t_ω in T and so does an endvertex of every (half) edge in $\widehat{C}([t_\omega], \omega)$; and the same is true for ω' . As by definition none of $t_\omega, t_{\omega'}$ lies above the other, it follows that $\widehat{C}([t_\omega], \omega)$ and $\widehat{C}([t_{\omega'}], \omega')$ are disjoint. Therefore N_x and N_y are disjoint if the open stars around x and y were chosen small enough. \square

We conclude this section by proving three simple lemmas to be used in Section 5.

Lemma 4.8 *Suppose that G satisfies (2) and has distinct vertices x_1, x_2, \dots any three of which lie on a common arc in \overline{G} . Then there exists a subsequence of x_1, x_2, \dots which converges in \overline{G} to an end of G .*

Proof. Suppose not. Since all the vertices x_i must lie in the same component of G , we may assume that G is connected. Then Lemma 2.2 implies that G contains a subdivided infinite star S with leaves in $\{x_1, x_2, \dots\}$. (Indeed, G cannot contain a ray R with infinitely many $\{x_1, x_2, \dots\}$ - R paths,

since then the starting vertices of these paths would converge to the end of G containing R .) Let X denote the set of leaves of S , and let s be the centre of S . Suppose first that one component C of $G - s$ contains infinitely many vertices from X . Applying Lemma 2.2 again to the graph C and the set $X \cap V(C) =: X'$ we find a subdivided infinite star S' in C whose leaves lie in X' . It is now easy to find infinitely many internally disjoint paths in G between s and the centre of S' , which contradicts our assumption that G satisfies (2). Therefore there are infinitely many components of $G - s$ each containing a vertex from X . Let $x, x', x'' \in X$ be vertices from different components of $G - s$, and let A be an arc in \tilde{G} containing them. Then A has a subarc which avoids s but joins two of the three points x, x', x'' , a contradiction to Lemma 4.5. \square

Lemma 4.9 *Suppose that G satisfies (2). Let x_1, x_2, \dots and y_1, y_2, \dots be sequences of distinct vertices of G which in \overline{G} converge to ends ω_x and ω_y , respectively. Suppose that for every $k \geq 1$ there exists an arc A_k in \tilde{G} containing all the points $x_1, y_1, \dots, x_k, y_k$ in that order. Then $\omega_x = \omega_y$.*

Proof. Suppose that $\omega_x \neq \omega_y$ and let S be a finite set of vertices of G separating ω_x from ω_y in G . By considering subsequences we may assume that $x_i \in C(S, \omega_x)$ and $y_i \in C(S, \omega_y)$ for all $i \geq 1$. Then the arc $A_{|S|+1}$ contains a subarc which avoids S but joins x_i to y_i for some $i \leq |S| + 1$, contradicting Lemma 4.5. \square

Our last lemma shows that if G satisfies (2) then every topological ray in \tilde{G} converges:

Lemma 4.10 *Suppose that G satisfies (2), and let $\sigma : [0, 1) \rightarrow R \subseteq \tilde{G}$ be a homeomorphism. Then σ can be extended to a continuous map $[0, 1] \rightarrow \tilde{G}$.*

Proof. Using Corollary 4.6, we can find a sequence $\Sigma = (x_1, x_2, \dots)$ of points in $R \cap G$ whose images under σ^{-1} converge to 1. Clearly the lemma holds if all but finitely many of the x_i lie on a common edge. We may therefore assume that every x_i is a vertex. By Lemma 4.8, some subsequence of Σ converges in \overline{G} to an end ω of G . We show that putting $\sigma(1) := \pi(\omega)$ makes σ continuous (at 1).

Let N be an open neighbourhood of $\sigma(1) = \pi(\omega)$ in \tilde{G} and S a finite set of vertices such that $\pi(\widehat{C}(S, \omega)) \subseteq N$. By Lemma 4.5, $\pi(\overline{C}(S, \omega) \cup E(S, \omega)) \setminus S$ is open in \tilde{G} , and the frontier of this set is contained in the finite set S . As R has arbitrarily late points x_i in $\pi(C(S, \omega))$, this implies that R has a final segment in $\pi(\widehat{C}(S, \omega))$, as required. \square

Lemma 4.10 implies that for every topological ray R in \tilde{G} there is a unique point $p \in \tilde{G}$ such that $R \cup \{p\}$ is a topological path in \tilde{G} . We will call p the *endpoint* of R .

5 Trees and paths in the identification topology

In Section 3 we saw that even if our graph $G = (V, E)$ satisfies (2), it may still happen that \tilde{G} has no topological spanning tree (Corollary 3.5). We now show that a topological spanning tree does exist if we strengthen our assumption of (2) to (3). To do so, we shall use Zorn's lemma to show that the set of path-connected subspaces of \tilde{G} has a minimal element with respect to edge-deletion (Lemma 5.1), which is then easily seen to be a topological spanning tree.

Let us recall some notation. A subsequence Σ' of a given (transfinite well-ordered) sequence Σ is *cofinal* in Σ if for every $s \in \Sigma$ there is an element of Σ' that does not strictly precede s . Given a (graph-theoretical) rooted tree T and $i \geq 0$, the *i th level* of T is the set of all its vertices at distance i from the root of T .

The following lemma has been abstracted from the proofs of Theorems 5.2, 5.3, 6.3 and will be used in all those proofs.

Lemma 5.1 *Assume that G satisfies (2), and let $x, y \in V \cup \Omega'$. Suppose that $(A_\alpha)_{\alpha < \gamma}$ is a (transfinite) sequence of x - y arcs in \tilde{G} . Then there is a topological x - y path P in \tilde{G} and a dense subset P^* of P such that $P^* \subseteq G$ and for all points $p \in P^*$ the arcs A_α containing p form a cofinal subsequence of $(A_\alpha)_{\alpha < \gamma}$. In particular, for every edge e whose interior meets P the arcs A_α containing e form a cofinal subsequence of $(A_\alpha)_{\alpha < \gamma}$.*

Proof. We may assume that G is connected and consider only the case that $x, y \in V$; the other cases are similar. If γ is a successor ordinal, say $\gamma = \beta + 1$, then by Corollary 4.6 we can set $P := A_\beta$. Thus we may assume that γ is a limit ordinal. By Lemma 2.4, G has a normal spanning tree T . Let us call a point $p \in G$ *good* if the arcs A_α containing p form a cofinal subsequence of $(A_\alpha)_{\alpha < \gamma}$. To construct our topological x - y path P , we shall first assign to every rational $r \in [0, 1]$ a good point $\sigma(r) \in G$. We then extend this map σ to a continuous map $[0, 1] \rightarrow \tilde{G}$, whose image will be the desired topological path P .

Put $\sigma(0) := x$ and $\sigma(1) := y$, and let r_1, r_2, \dots be an enumeration of $(0, 1) \cap \mathbb{Q}$. We define our partial mapping σ in at most ω steps, so that after step n its domain is a closed subset of $[0, 1]$ containing r_n .

If xy is an edge of G and the A_α consisting of xy form a cofinal subsequence of $\Sigma^0 := (A_\alpha)_{\alpha < \gamma}$, then let $\sigma : [0, 1] \rightarrow xy \subseteq G$ be a homeomorphism sending 0 to x and 1 to y . (So in this case we take xy for P .) Otherwise we define σ only at r_1 . Our candidates for $\sigma(r_1)$ are all the good vertices $z \in V \setminus \{x, y\}$. Since G satisfies (2), there is a finite set $S \subseteq V$ separating x from y in $G - xy$. By Lemma 4.5, S meets every arc A_α not consisting of xy , so there is at least one candidate for $\sigma(r_1)$. From amongst all the candidates we choose a vertex z_1 at the lowest possible level of T , set $\sigma(r_1) := z_1$, and define Σ^1 to be the cofinal subsequence of Σ^0 consisting of all A_α containing z_1 .

Next we consider r_2 . For example, let us assume that $r_2 \in (0, r_1)$. If xz_1 is an edge of G and if the subsequence $\Sigma_{xz_1}^2$ of Σ^1 consisting of all A_α with $xA_\alpha z_1 = xz_1$ is cofinal in Σ^1 , we define σ on $(0, r_1)$ so as to send $[0, r_1]$ continuously and bijectively onto xz_1 and put $\Sigma^2 := \Sigma_{xz_1}^2$. Otherwise we just choose a good vertex as $\sigma(r_2)$. This time our candidates for $\sigma(r_2)$ are the vertices z for which the subsequence $\Sigma_z^2 \subseteq \Sigma^1$ of all A_α with $z \in \mathring{x}A_\alpha\mathring{z}_1$ is cofinal in Σ^1 . (As Σ^1 is cofinal in Σ^0 , all these candidates for $\sigma(r_2)$ are good vertices.) As before there is at least one candidate. From amongst all the candidates we choose a vertex z_2 at the lowest possible level of T , and put $\sigma(r_2) := z_2$ and $\Sigma^2 := \Sigma_{z_2}^2$. As before, Σ^2 is cofinal in Σ^1 and hence in $\Sigma^0 = (A_\alpha)_{\alpha < \gamma}$.

Now consider the first rational r in r_1, r_2, \dots for which $\sigma(r)$ is not yet defined. Since the current domain of σ is closed in $[0, 1]$ and its frontier consists of rationals, there are rationals $q_1 < r < q_2$ such that σ is already defined on both q_1 and q_2 but not yet on any point in (q_1, q_2) . For all A_α in Σ^2 we consider the segments $\sigma(q_1)A_\alpha\sigma(q_2)$ and extend σ as before, either as a homeomorphism between $[q_1, q_2]$ and the edge $\sigma(q_1)\sigma(q_2)$ or by choosing a good vertex as $\sigma(r)$. We continue in this fashion for at most ω steps until we have defined σ on $[0, 1] \cap \mathbb{Q}$. Let X be the domain of σ . Then X contains all rationals in $[0, 1]$ and, for each irrational $q \in X$, $\sigma(q)$ is an inner point of an edge contained in $\sigma(X)$. Moreover, σ is injective on X . In what follows we will extend σ to a continuous map $[0, 1] \rightarrow \tilde{G}$ by sending the points of $[0, 1] \setminus X$ to suitable ends of G (or to vertices dominating them). As $\sigma(X) \supseteq \sigma([0, 1] \cap \mathbb{Q})$ consists of good points, $\sigma([0, 1])$ will then be a topological path P as desired.

Let I be the set of all points $p \in [0, 1]$ for which there exists a sequence $q_1 < q_2 < \dots$ of rationals converging to p such that each $\sigma(q_i)$ is a vertex. For $p \in I$, let \mathcal{Q}_p denote the set of sequences $\sigma(q_1), \sigma(q_2), \dots$ of vertices corresponding to such sequences $q_1 < q_2 < \dots$ of rationals. Similarly, let I' be the set of all points $p \in [0, 1]$ for which there exists a sequence

$q_1 > q_2 > \dots$ of rationals converging to p such that each $\sigma(q_i)$ is a vertex. For $p \in I'$, let \mathcal{Q}'_p denote the set of sequences $\sigma(q_1), \sigma(q_2), \dots$ corresponding to such sequences $q_1 > q_2 > \dots$. Note that $[0, 1] \setminus X \subseteq I \cap I'$.

Let us prove the following:

For every $p \in I$ all the sequences in \mathcal{Q}_p converge in \overline{G} to a single end ω_p . Similarly, for every $p \in I'$ all the sequences in \mathcal{Q}'_p converge in \overline{G} to a single end ω'_p . ()*

We only consider the case that $p \in I$; the other case is similar. Consider any finite set $U \subset (0, 1) \cap \mathbb{Q}$, and let r_i be the last element of U in our enumeration of $(0, 1) \cap \mathbb{Q}$. Then every A_α with $\alpha \in \Sigma^i$ contains all the points of $\sigma(U \cup \{0, 1\})$, in the order induced by $[0, 1]$ and σ . Lemma 4.8 therefore implies that every sequence in \mathcal{Q}_p has a subsequence which converges in \overline{G} to an end of G . Lemma 4.9 shows that these ends are the same for all such subsequences and all choices of sequences in \mathcal{Q}_p ; in particular, every sequence in \mathcal{Q}_p must itself converge to this single end. This completes the proof of (*).

We now extend σ to all of $[0, 1]$ by setting $\sigma(p) := \pi(\omega_p)$ for all $p \in [0, 1] \setminus X$. Thus if $\omega_p \in \Omega'$ then $\sigma(p) = \omega_p$, while otherwise $\sigma(p)$ is the unique vertex dominating ω_p . This completes the definition of σ .

For our proof that σ is continuous we need the following assertions about all $p \in [0, 1]$:

If $p \notin X$ then $\omega_p = \omega'_p$.

*If $p \in I \cap \mathbb{Q}$ (and thus $\sigma(p)$ is a vertex) then $\sigma(p)$ dominates ω_p . (**)*

If $p \in I' \cap \mathbb{Q}$ (and thus $\sigma(p)$ is a vertex) then $\sigma(p)$ dominates ω'_p .

We only consider the case that $p \notin X$; the other cases are similar. Suppose that $\omega_p \neq \omega'_p$. Since $p \notin X$, and after any finite number of steps in the inductive definition of σ the set of points in $[0, 1]$ for which σ was still undefined was open and the frontier of this set consisted of rationals, there is a sequence $\sigma(s_1), \sigma(s_2), \dots$ in \mathcal{Q}_p with the property that for every $i \geq 1$ there are rationals q_i^1, q_i^2 such that $q_i^1 < s_i < p < q_i^2$ and such that when σ was defined for s_i it had previously been defined for q_i^1 and q_i^2 but not for any point in (q_i^1, q_i^2) . Then $\sigma(q_i^1)$ and $\sigma(q_i^2)$ are vertices, the points $q_1^1 < q_2^1 < \dots$ converge to p from below, and $q_1^2 \geq q_2^2 \geq \dots$ converge to p from above. By choosing a subsequence if necessary we may further assume that $q_1^2 > q_2^2 > \dots$. Then the sequence $\sigma(q_1^1), \sigma(q_2^1), \dots$ lies in \mathcal{Q}_p while $\sigma(q_1^2), \sigma(q_2^2), \dots$ lies in \mathcal{Q}'_p . Now (*) implies that $\sigma(q_1^1), \sigma(q_2^1), \dots$ converges to ω_p , while $\sigma(q_1^2), \sigma(q_2^2), \dots$ converges to ω'_p . Let S be a finite set of vertices

separating ω_p from ω'_p in G . Then for all but finitely many i we have $\sigma(q_i^1) \in C(S, \omega_p)$ but $\sigma(q_i^2) \in C(S, \omega'_p)$, and hence every arc of the form $\sigma(q_i^1)A_\alpha\sigma(q_i^2)$ meets S (Lemma 4.5). Hence for all but finitely many i some vertex in S was a candidate for $\sigma(s_i)$. But since, by (*), the sequence $\sigma(s_1), \sigma(s_2), \dots$ converges to ω_p in \overline{G} , eventually the vertices $\sigma(s_i)$ lie at a higher level of T than all the vertices from S , contradicting the definition of $\sigma(s_i)$ for these s_i . This completes the proof of (**).

Let us now show that $\sigma : [0, 1] \rightarrow \widetilde{G}$ is continuous. This is clear at points $p \in [0, 1] \setminus (I \cup I')$. Indeed, for all such p , the point $\sigma(p)$ is either an inner point of an edge contained in $\sigma([0, 1])$ or a vertex incident with two edges contained in $\sigma([0, 1])$ or, if $p = 0, 1$, an endvertex of an edge contained in $\sigma([0, 1])$. So suppose that $p \in I \cup I'$. We will only consider the case that $p \notin X$, the remaining cases being similar. Let N be an open neighbourhood of $\sigma(p)$ in \widetilde{G} . Then there is a finite set S of vertices such that $\pi(\widehat{C}(S, \omega_p)) \subseteq N$. We have to show that there exist points $a < p < b$ such that $(a, b) \subseteq \sigma^{-1}(N)$. If not, then there exists a sequence r_1, r_2, \dots of reals converging to p whose images under σ all lie outside $\pi(\widehat{C}(S, \omega_p))$. We may assume that these $\sigma(r_i)$ are not inner points of edges. Indeed, as σ is injective on X , only finitely many such edges can have an endvertex in S , and so all but finitely many of the r_i whose image is an inner point of an edge can be replaced by a rational whose image is an endvertex of that edge lying outside $S \cup C(S, \omega_p)$ and hence outside $\pi(\widehat{C}(S, \omega_p))$. Thus in particular no r_i is an irrational contained in X . We may even assume that every r_i is rational: if $r_i \in [0, 1] \setminus X \subseteq I$ then $\omega_{r_i} \notin \overline{C}(S, \omega_p)$ (as $\sigma(r_i) \notin \pi(\widehat{C}(S, \omega_p))$), so by (*) we may replace r_i with a rational close to it whose image is a vertex outside $S \cup C(S, \omega_p)$. But now the sequence $\sigma(r_1), \sigma(r_2), \dots$ has a subsequence in \mathcal{Q}_p or \mathcal{Q}'_p not converging to $\omega_p = \omega'_p$ (cf. (**)), in violation of (*). \square

Theorem 5.2 *If G is connected and satisfies (3), then \widetilde{G} has a topological spanning tree.*

Proof. Let \mathcal{X} be the set of all path-connected subspaces of \widetilde{G} of the form $\widetilde{G} \setminus \overset{\circ}{F}$ with $F \subseteq E$. Then \mathcal{X} is non-empty since $\widetilde{G} \in \mathcal{X}$. Let \mathcal{X} be ordered by inclusion, and let us use Zorn's lemma to show that \mathcal{X} has a minimal element. Let $(X_\alpha)_{\alpha < \gamma}$ be a (well-ordered) descending chain in \mathcal{X} , say $X_\alpha = \widetilde{G} \setminus \overset{\circ}{F}_\alpha$. (Thus $(F_\alpha)_{\alpha < \gamma}$ is an ascending chain of subsets of E .)

Let us show that $X := \bigcap_{\alpha < \gamma} X_\alpha \in \mathcal{X}$. Clearly $X = \widetilde{G} \setminus \overset{\circ}{F}$ with $F := \bigcup_{\alpha < \gamma} F_\alpha$. In particular, $V \cup \Omega' \subseteq X$. To show that X is path-connected, let x, y be distinct points in $V \cup \Omega'$. In every X_α there is a topological x - y

path, which by Lemma 2.1 and Theorem 4.7 contains an x - y arc A_α . By Lemma 5.1 these yield a topological x - y path P in \tilde{G} that avoids \tilde{F} and hence lies in X . We have thus shown that every descending chain in \mathcal{X} has a lower bound, and hence that \mathcal{X} has a minimal element T .

It remains to show that T is a topological spanning tree of \tilde{G} . If not, then T contains a circle D . By Corollary 4.4, D contains an edge e . But then $T \setminus \dot{e}$ is still path-connected and hence contained in \mathcal{X} , contradicting the minimality of T . \square

Let us reapply Lemma 5.1 to prove the following:

Theorem 5.3 *If G is countable and satisfies (2), then every closed connected subset of \tilde{G} is path-connected.*

Proof. Suppose that $X \subseteq \tilde{G}$ is closed and connected, but not path-connected. It is easily seen that there are $x, y \in V \cup \Omega'$ lying in different path-components of X . Let e_1, e_2, \dots be an enumeration of all the edges $e \in G$ with $\dot{e} \not\subseteq X$, and let z_1, z_2, \dots be an enumeration of all the vertices of G outside X . Let $G_i := G - \{e_1, \dots, e_i\} - \{z_1, \dots, z_i\}$.

Suppose that x and y belong to the same component of G_i for all i . Then each \tilde{G}_i contains an x - y arc A_i that is a finite path or the closure of a ray or a double ray, and by Lemma 5.1 there is a topological x - y path P in \tilde{G} with a dense subset P^* such that $P^* \subseteq G$ and every point from P^* lies in A_i for infinitely many i . Then $P^* \subseteq X$ and, as X is closed, $P \subseteq X$. This contradicts the choice of x and y .

So there exists an i such that x and y belong to different components of G_i . We will show that this implies that X cannot be connected (a contradiction). Put $F := \{e_1, \dots, e_i\}$ and $S := \{z_1, \dots, z_i\}$. Let C_x and C_y denote the components of G_i with $x \in \overline{C_x}$ and $y \in \overline{C_y}$. By making F smaller (and replacing G_i with a supergraph) we may assume that every edge in F joins C_x to C_y . From the interior of every edge $e \in F$ pick a point a^e not in X . Let S_x be the union of S with the set of endvertices of edges from F outside C_x . Define S_y correspondingly. By Lemma 4.5, every set of the form $\pi(\widehat{C_x}) \setminus S_x$ is open in \tilde{G} . Since no end belonging to C_x is dominated by a vertex in $S_x \setminus S$, we have $\pi(\widehat{C_x}) \setminus S_x = \pi(\widehat{C_x}) \setminus S$. Let N_x be the set of the form $\pi(\widehat{C_x}) \setminus S$ which contains \dot{e} for every edge e joining C_x to S and which contains the half-edge $[c, a^e] \subseteq e$ with $c \in C_x$ for every $e \in F$. Define N_y correspondingly. By our assumption on F , every component C of G_i other than C_x, C_y is a component of $G - S$. For every such C let N_C be the set of the form $\pi(\widehat{C}) \setminus S$ which contains \dot{e} for every edge e joining C to S . Let N'_x be the union of N_x , all the N_C and the interiors of all the edges in $G[S]$. Then N'_x and N_y

are disjoint open subsets of \tilde{G} whose union contains X , contradicting the connectedness of X . \square

It is not hard to show that \tilde{G} is locally path-connected. This implies that every open connected subset of \tilde{G} is path-connected. We expect that Theorem 5.3 extends to connected subsets that are neither closed nor open, but have been unable to prove this.

6 Cycles in the identification topology

In this section we extend Theorems 3.1–3.3 to all graphs G satisfying (2) endowed with ITOP.

Theorem 6.1 *Let G be a graph satisfying (2). Then the fundamental circuits of \tilde{G} with respect to any fixed topological spanning tree span its cycle space $\mathcal{C}(\tilde{G})$.*

The proof of this theorem is similar to its analogue for locally finite graphs [4, Thm. 5.1]. The following lemma ensures that sums of distinct fundamental circuits are always well-defined.

Lemma 6.2 *Let G be a graph satisfying (2). Then the fundamental circuits of \tilde{G} with respect to any fixed topological spanning tree T form a thin family.*

Proof. Suppose not. Then there exists an edge $e = xy$ that lies in infinitely many fundamental circuits C_{e_i} ($i = 1, 2, \dots$). Clearly $e \in E(T)$. Let B_x and B_y be the path-components of $T \setminus \dot{e}$ containing x and y , respectively. Since T contains no circle, Lemma 2.1 and Theorem 4.7 imply that B_x and B_y are distinct. Clearly each e_i joins a vertex $x_i \in B_x$ to a vertex $y_i \in B_y$. As all the e_i are distinct, at least one of the sets $U_x := \{x_i \mid i \geq 1\}$ and $U_y := \{y_i \mid i \geq 1\}$ is infinite. Let us assume that U_x is infinite. Apply Lemma 2.3 to B_x and U_x to obtain an infinite set $U'_x \subseteq U_x$ and either a topological comb $C_x \subseteq B_x$ whose set of teeth is U'_x or a topological \aleph_0 -star $S_x \subseteq B_x$ whose set of leaves is U'_x . If Lemma 2.3 returns a topological comb C_x , let $p_x \in \tilde{G}_x$ be the endpoint of its back R_x (as defined in and after Lemma 4.10). By our definition of a comb, p_x is a limit of vertices or ends, and hence is either itself a vertex or an end; in particular, $p_x \in T$. Replacing C_x with a subcomb if necessary, we may assume that $p_x \notin C_x$; then $R_x \cup \{p_x\}$ is an arc in B_x . Let \mathcal{A}_x be the set of all arcs in $C_x \cup \{p_x\}$ joining p_x to a tooth of C_x . If Lemma 2.3 returns a topological \aleph_0 -star S_x ,

let p_x be its centre (which may be an end), and let \mathcal{A}_x be the set of all arcs in S_x joining p_x to a leaf of S_x .

Let $U'_y \subseteq U_y$ be the set of all y_i for which $x_i \in U'_x$. If U'_y is finite, let p_y be any point in U'_y such that $p_y = y_i$ for infinitely many i with $y_i \in U'_y$. If U'_y is infinite, then apply Lemma 2.3 again to B_y and U'_y to obtain an infinite set $U''_y \subseteq U'_y$ and either a topological comb C_y or a topological \aleph_0 -star S_y with teeth (resp. leaves) in U''_y . Define p_y and \mathcal{A}_y as earlier for x . Thus in each case we have $p_y \in B_y$ and, if U'_y is infinite, \mathcal{A}_y consists of arcs in B_y . Let \mathcal{A} be the (infinite) set of all p_x - p_y arcs with a first segment in \mathcal{A}_x , another segment equal to some e_i , and, if U'_y was infinite, a final segment in \mathcal{A}_y . Note that every arc in \mathcal{A} contains a vertex of B_x other than p_x ; hence if p_x and p_y are both vertices and G contains the edge $e_{xy} := p_x p_y$, then no arc in \mathcal{A} meets \mathring{e}_{xy} . Moreover, by construction of \mathcal{A}_x and \mathcal{A}_y no vertex other than p_x and p_y lies on more than finitely many arcs in \mathcal{A} .

By (2), there is a finite set S of vertices separating p_x from p_y in G (resp. in $G - e_{xy}$, if e_{xy} exists). By Lemma 4.5 every arc in \mathcal{A} meets S , and hence infinitely many arcs in \mathcal{A} share an inner vertex (a contradiction). \square

Proof of Theorem 6.1. Let T be a topological spanning tree of \tilde{G} . It suffices to prove the following claim.

Every circuit C is equal to the sum of all the fundamental circuits C_e with $e \in C \setminus E(T)$. (*)

Before proving (*), let us show how it implies Theorem 6.1. Let $Z \in \mathcal{C}(\tilde{G})$ be given. By definition, Z is a sum $Z = \sum_{i \in I} C_i$ of distinct circuits C_i . By (*) we have $C_i = \sum_{C' \in \mathcal{C}_i} C'$ with $\mathcal{C}_i := \{C_e \mid e \in C_i \setminus E(T)\}$ for all $i \in I$. Let \mathcal{C} be the family $\bigcup_{i \in I} \mathcal{C}_i$. (So a fundamental circuit lying in several C_i occurs more than once in \mathcal{C} .) Then every fundamental circuit C_e occurs only finitely often in \mathcal{C} : if C_e occurs in some \mathcal{C}_i then it does so only once, giving $e \in C_i$; as the family $(C_i)_{i \in I}$ is thin, this happens for only finitely many i . So by Lemma 6.2 the family \mathcal{C} is thin. Clearly, the circuits in \mathcal{C} sum to Z .

Let us now prove (*). By Lemma 6.2 the family of all fundamental circuits C_e with $e \in C \setminus E(T)$ is thin, so it suffices to show that it sums to C . Thus for every edge $f \in G$ we have to show that f lies in C if and only if it lies in an odd number of the circuits C_e in (*). This is clear if $f \notin E(T)$. So let us assume that $f \in E(T)$ and let B_1 and B_2 be the path-components of $T \setminus \mathring{f}$ containing the two endvertices of f , respectively. Then $B_1 \cup B_2 = T \setminus \mathring{f}$ (because T is path-connected), and $B_1 \neq B_2$ (because T contains no circle). By Lemma 6.2, the set E_f of all the edges of G between B_1 and B_2 is

finite, because $E_f \setminus \{f\}$ consists of precisely those edges $e \notin E(T)$ whose fundamental circuit C_e contains f . We will show that $|E_f \cap C|$ is even. This will imply that $f \in C$ if and only if C contains an odd number of other edges from E_f , ie. if and only if f lies in an odd number of the circuits C_e in (*).

So suppose that $|E_f \cap C|$ is odd. Let D be a circle in \tilde{G} whose circuit is C . Then the closure of $D \setminus \bigcup(E_f \cap C)$ consists of subarcs of D between the endpoints of edges in $E_f \cap C$. Since $|E_f \cap C|$ is odd, there must be at least one such arc A which joins a vertex $x_1 \in B_1$ to a vertex $x_2 \in B_2$. Since any x_1 - x_2 path in G contains a B_1 - B_2 edge, E_f separates x_1 from x_2 in G . But the x_1 - x_2 arc A avoids $\overset{\circ}{E}_f$, so this contradicts Lemma 4.2. \square

Theorem 6.3 *Let G be a graph satisfying (2). Then every element of $\mathcal{C}(\tilde{G})$ is a union of disjoint circuits.*

Proof. Let $Z = \sum_{\alpha < \gamma} C_\alpha$ be any element of $\mathcal{C}(\tilde{G})$, the C_α being circuits in \tilde{G} . Every C_α is a countable set (of edges), because every edge on a circle has an inner point that corresponds to a rational point on the unit circle. Consider the auxiliary graph H whose vertices are the C_α , and in which C_α and $C_{\alpha'}$ are joined by an edge whenever they are not disjoint. Since the C_α form a thin family and are countable, so are the components of H . For each component D of H let Z_D be the sum of all those C_α that are vertices of D . Thus Z_D is the sum of countably many circuits, and Z is the disjoint union of all the Z_D . Therefore, to prove the theorem for Z , it suffices to show that each Z_D is a union of disjoint circuits. So let us prove the following claim.

Let $Z' = \sum_{i \in I} C_i$ be the sum of countably many circuits and let $e = xy$ be any edge in Z' . Then \tilde{G} has an x - y arc A that contains only edges from $Z' \setminus \{e\}$. ()*

Before we prove (*), let us see how it implies that Z' is a union of disjoint circuits. Let x_1y_1, x_2y_2, \dots be an enumeration of all the edges in Z' . Apply (*) to obtain an x_1 - y_1 arc A'_1 in \tilde{G} that contains only edges from $Z' \setminus \{x_1y_1\}$. Then $A'_1 \cup x_1y_1$ is a circle in \tilde{G} whose circuit $C(A'_1 \cup x_1y_1) =: C'_1$ contains x_1y_1 , and $Z'' := Z' + C'_1$ is a subset of Z' that does not contain x_1y_1 . Let x_jy_j be the first edge from x_1y_1, x_2y_2, \dots contained in Z'' , and apply (*) to Z'' and x_jy_j to obtain an x_j - y_j arc A'_2 which contains only edges from $Z'' \setminus \{x_jy_j\}$. Again $A'_2 \cup x_jy_j$ is a circle in \tilde{G} whose circuit $C(A'_2 \cup x_jy_j) =: C'_2$ contains x_jy_j . Let $Z''' := Z'' + C'_2$ and continue in this fashion for at most ω steps to exhaust Z' . Then Z' is the union of the disjoint circuits C'_1, C'_2, \dots

To prove (*), let G' be the subgraph of G with the edge set $\bigcup_{i \in I} C_i$. Let $E' := (E(G') \setminus Z') \cup \{e\}$. Choose an enumeration e_0, e_1, \dots of the edges in E' , with $e_0 = e$. We shall show that for each $j \geq 1$ there is an x - y arc A_j in \tilde{G} that contains only edges in $E(G') \setminus \{e_0, \dots, e_j\}$. Lemma 5.1 then yields a topological x - y path P in \tilde{G} which meets only the interiors of edges that lie on A_j for infinitely many j , and hence lie in $E(G') \setminus E' = Z' \setminus \{e\}$. By Lemma 2.1 and Theorem 4.7, P contains an x - y arc A , which is as desired in (*).

So let us prove the existence of the arcs A_j . Since the family $(C_i)_{i \in I}$ is thin, we can choose a sequence $X_0 \subset X_1 \subset \dots$ of finite subsets of $\{C_i \mid i \in I\}$ such that each X_j contains all the circuits C_i containing e_j . With every circuit $C \in X_j$ we associate a finite auxiliary cycle C' , as follows. Let D be a circle in \tilde{G} whose circuit is C . To form C' , we first take all the edges in $C \cap \{e_0, \dots, e_j\} =: E_C$, in the same cyclic order as on D . The closure of $D \setminus \bigcup E_C$ is a disjoint union of closed segments S of D , and we form C' by replacing in D each of these segments S by a new vertex x_S joined to the endpoints of S . These new vertices shall differ for distinct segments S and distinct circuits $C \in X_j$. Now let H_j be the finite graph consisting of the sum of all the C' with $C \in X_j$. The definition of X_j implies that $\sum_{C \in X_j} C$, and hence also $E(H_j)$, agrees with Z' on the set $\{e_0, \dots, e_j\}$, ie. contains precisely $e = e_0$ from this set. As H_j is a finite sum of finite cycles and hence an edge-disjoint union of finite cycles, it contains a finite path P'_j that joins the endvertices of e but does not contain e . Replacing in P'_j the vertices x_S and their incident edges with the corresponding circle segments S , we obtain a topological x - y path P_j in \tilde{G} that contains only edges from $E(G')$ and avoids the interiors of all of e_0, \dots, e_j . Lemma 2.1 and Theorem 4.7 imply that P_j contains the desired x - y arc A_j . \square

Corollary 6.4 *For every graph G satisfying (2), its cycle space $\mathcal{C}(\tilde{G})$ is closed under infinite sums.*

Proof. Any sum of (a thin family of) elements of $\mathcal{C}(\tilde{G})$ that are each a union of disjoint circuits can be rewritten as the sum of all these circuits, since these again form a thin family. \square

Theorem 6.5 *Let G be a graph satisfying (3). Then its cycle space $\mathcal{C}(\tilde{G})$ consists of precisely those sets of edges that meet every finite cut in an even number of edges.*

Proof. Let $F \subseteq E(G)$ be any finite cut in G . As in the proof of (*) in Theorem 6.1 it can be shown that every circuit meets F in an even number of edges. Since for every $Z = \sum_{i \in I} C_i$ in $\mathcal{C}(\tilde{G})$ only finitely many of the circuits C_i meet F , and since finite sums (mod 2) of even sets are even, it follows that Z meets F in an even number of edges.

For the converse implication suppose that $Z \subseteq E(G)$ meets every finite cut in an even number of edges, and assume without loss of generality that G is connected. By Theorem 5.2, \tilde{G} has a topological spanning tree T . We show that Z is equal to the sum of all the fundamental circuits C_e with $e \in Z \setminus E(T)$. Let f be an edge of G . We have to show that $f \in Z$ if and only if f lies in an odd number of C_e with $e \in Z \setminus E(T)$. This is clear if $f \notin E(T)$. So suppose that $f \in E(T)$ and let E_f be the set of all edges in G joining the two path-components of $T \setminus \overset{\circ}{f}$. Then E_f is a cut in G , and the fundamental circuits containing f are precisely the C_e with $e \in E_f \setminus \{f\}$. Hence by Lemma 6.2, E_f is finite. Thus by assumption Z meets E_f in an even number of edges, ie. Z contains f if and only if it contains an odd number of the other edges from E_f . This is the case if and only if f lies on an odd number of fundamental circuits C_e with $e \in Z \setminus E(T)$, as required. \square

We remark that Theorem 6.5 also holds for graphs G which only satisfy (2). Indeed, as in the proof of Theorem 6.5 it can be shown that every element of $\mathcal{C}(\tilde{G})$ meets every finite cut in an even number of edges. The difference now is that in the proof of the converse implication we can no longer assume that \tilde{G} has a topological spanning tree. Instead of using such a tree, we consider a *pre-tree* T of \tilde{G} : a path-connected subspace of G that contains no circle *with a non-empty circuit* and which is obtained from \tilde{G} by deleting $\overset{\circ}{F}$ for some $F \subseteq E(G)$. Thus, every topological spanning tree of \tilde{G} is a pre-tree. Corollary 4.4 implies that for graphs satisfying (3) the converse is true, while the graph constructed in Proposition 3.4 shows that the converse need not hold if we only assume (2). But as in the proof of Theorem 5.2 one can show that \tilde{G} has a pre-tree T whenever G satisfies (2). Moreover, adding an edge $e \in E(G) \setminus E(T)$ to T yields a circle in $T \cup e$ containing e , but there may be several such circles. However, they must all have the same circuit C_e . So we may think of these circuits C_e as fundamental circuits. As before it can be shown that the fundamental circuits with respect to a pre-tree T of \tilde{G} form a thin family, so every sum of distinct such fundamental circuits is well-defined. This again implies that for every edge f the set E_f of all edges in G joining the two path-components of $T \setminus \overset{\circ}{f}$ is finite. ($T \setminus \overset{\circ}{f}$ consists of two path-components, because f is not contained

in a circle in T .) Finally, as in the proof of Theorem 6.5 one shows that if $Z \subseteq E(G)$ meets each finite cut in G in an even number of edges then Z is equal to the sum of all C_e with $e \in Z \setminus E(T)$, and hence lies in $\mathcal{C}(\tilde{G})$.

Note that the discussion in the previous paragraph also shows that Theorem 6.1 remains true for the fundamental circuits with respect to any fixed pre-tree.

7 Topological vs. end-faithful spanning trees, and their general existence problem

Our treatment of topological spanning trees has so far been motivated by the role they can play for the study of the cycle space, which is why we considered the problem of their existence only for the relatively narrow class of graphs satisfying (3). In this section we consider the existence problem more generally. Unless otherwise mentioned, we assume that any graph with ends considered is endowed with the standard topology TOP.

Topological spanning trees are closely related to ‘end-faithful’ spanning trees (see below), which have been widely studied in the literature. In order to put the existence problem for topological spanning trees into context, we start by pointing out this relationship.

Let G be any graph. Given any subgraph $H \subseteq G$, there is a canonical map $\eta : \Omega(H) \rightarrow \Omega(G)$ taking every end of H to the end of G that contains it as a subset (of rays). H is called *end-faithful* in G if this map η is a bijection, and *topologically end-faithful* if it is a homeomorphism of the subspaces $\Omega(H) \subset \overline{H}$ and $\Omega(G) \subset \overline{G}$. (By definition of TOP, η is always continuous.) If H is locally finite and end-faithful, it is also topologically end-faithful (because \overline{H} is compact), but in general the latter is a stronger property.

Proposition 7.1 *If T is a topological spanning tree of \overline{G} and the graph $H := T \cap G$ is connected, then H is an end-faithful (ordinary) spanning tree of G .*

Proof. H is clearly a spanning tree of G ; we show that H is end-faithful. If an end ω of G contains rays R and R' from two distinct ends of H , we can choose them so that $R \cup R'$ is a double ray. Then $R \cup R' \cup \{\omega\}$ is a circle in T , a contradiction. Hence every end of G contains at most one end of H .

Now suppose that some end ω of G contains no end of H . Let x be any vertex of G , and let A be an x - ω arc in T (which exists by Lemma 2.1). It is easy to see that A starts with a ray $R \subseteq H$ as an initial segment [5,

Lemma 2.3]; let ω' be the end of G containing R . By assumption $\omega' \neq \omega$, so $R \not\subseteq A$. Pick a vertex $y \in A \setminus R$ (which again exists by [5, Lemma 2.3]). Then xAy is not equal to the finite x - y path in H , and hence T contains a circle (contradiction). \square

If the spanning tree H in Proposition 7.1 is not locally finite, it need not be topologically end-faithful. For example, consider the graph G obtained from K_{\aleph_1} by adding for every vertex v a new ray that starts at v but is otherwise disjoint from the K_{\aleph_1} and from the other new rays. Let ω denote the end of the K_{\aleph_1} in G . Let H be any end-faithful spanning tree of G (which is easily found), and let T be its closure in \overline{G} ; this is easily seen to be a topological spanning tree of \overline{G} . Then $\eta^{-1}(\omega)$ will have an open neighbourhood O in \overline{H} that excludes infinitely many ends. But in \overline{G} every neighbourhood of ω contains all but finitely many ends, so η^{-1} cannot map it into O .

Thus, although \overline{H} and T in Proposition 7.1 coincide as point sets (up to the bijection η) and in the topologies they induce on H , the topology of \overline{H} on this set (ie. TOP for \overline{H}) may be finer than that of T (the subspace topology from \overline{G}). This can have curious effects; see Proposition 7.4.

Let us consider the converse problem to Proposition 7.1. Given an end-faithful spanning tree H of G , let us refer to its closure in \overline{G} as the subspace of \overline{G} induced by H . This subspace contains all the vertices and ends of G (because H spans G and every neighbourhood of an end contains a vertex), it is path-connected (because H is path-connected, every end is the limit point of all its rays, and H contains a ray from every end), and it contains every edge of which it contains an inner point (because H does). So the only reason why this space might fail to be a topological spanning tree of \overline{G} is that it might contain a circle—which can indeed happen (see below).

Problem 7.2 *For which graphs G does every end-faithful spanning tree induce a topological spanning tree in \overline{G} ?*

For locally finite graphs this is always the case:

Theorem 7.3 *If G is locally finite, then a spanning tree of G is end-faithful if and only if it induces a topological spanning tree in \overline{G} .*

Proof. We only have to show that if H is an end-faithful spanning tree of G then its closure in \overline{G} contains no circle. If it did, then by [4, Lemma 4.3] (or by Corollary 4.4) this would be the closure of a circuit C in \overline{G} . By [4, Thm. 5.1] (or by Theorem 6.1), C would be a sum of fundamental circuits of H and hence contain a chord of H , a contradiction. \square

(The reader may wonder whether it is necessary in the proof of Theorem 7.3 to use the result of [4, Thm. 5.1]. Indeed, if we extend η to all of \overline{H} by the identity on H , then $\eta : \overline{H} \rightarrow \overline{G}$ is continuous and injective, and hence a topological embedding (since H is locally finite and hence \overline{H} compact). So all we need to show is that \overline{H} itself contains no circle. But the proof of this ‘obvious’ fact, though straightforward, is already about half of the short proof of [4, Thm. 5.1] (which is just like the proof of $(*)$ in Theorem 6.1).)

In general, however, the converse of Proposition 7.1 can fail:

Proposition 7.4 *There is a countable graph G that has an end-faithful spanning tree whose closure in \overline{G} contains a circle.*

Proof. Consider the binary tree T_2 whose vertices are the finite 0–1 sequences and where each sequence is adjacent to its two one-digit extensions. The ends of T_2 correspond to the infinite 0–1 sequences, which we view as binary expansions of the reals in $[0, 1]$. Let J be the set of all those rationals in $(0, 1)$ that have a finite binary expansion. Every number in $[0, 1] \setminus J$ corresponds to exactly one end of T_2 , while every $q \in J$ has the form $q = 0.s1$ and corresponds to the two ends $s1000\dots$ and $s0111\dots$. Let G be the graph obtained from T_2 by adding for each $s \in T_2$ a new edge e_s between the vertices $s100$ and $s011$; then T_2 is an end-faithful subgraph of G . For every $q = 0.s1$ in J let D_q denote the double ray consisting of the new edge e_s and the two rays of T_2 in $s1000\dots$ and $s01111\dots$ starting at the endpoints of e_s . Let D_0 denote the double ray that is the union of the two rays of T_2 starting at the empty sequence \emptyset and corresponding to the numbers 0 and 1. In [4, Section 5] we showed that the closure D in \overline{G} of all the D_q with $q \in J \cup \{0\}$ is a circle containing all the ends of G .

Let G' be the graph obtained from G by subdividing every edge of the form e_s once. Let v_s denote the subdividing vertex and, for each $q \in J \cup \{0\}$, let D'_q be the subdivision of D_q contained in G' . (Thus $D'_0 = D_0$.) Clearly, the set D' obtained from D by replacing each D_q with D'_q and each end of G by its corresponding end of G' is a circle in $\overline{G'}$.

Our aim now is to add edges to G' in order to obtain a graph G^* in which G' is end-faithful, and which has an end-faithful spanning tree H containing all these double rays D'_q . Then the closure of H in $\overline{G^*}$ will contain D' (replace the ends of G' in D' by the corresponding ends of G^*), and D' will still be a circle in $\overline{G^*}$.

To do this, first join in G' the vertex \emptyset to all those vertices of the form v_s that have distance 3 from D'_0 in G' (Fig. 8). Let N_\emptyset denote the set of all these neighbours. Then each component C of $G' - (D'_0 \cup N_\emptyset)$ is a copy of G' , where




Figure 8: G' , and the edges of G^* at \emptyset

the unique neighbour t_C of D'_0 in C plays the role of \emptyset in G' . Similarly as above, we join t_C to all those vertices in C of the form v_s that have distance 3 from the ‘outer double ray’ of C —the double ray that is union of the two rays in T_2 starting at t_C and belonging to the ends $t_C 000\dots$ and $t_C 111\dots$ of T_2 . We continue in this fashion and denote the resulting graph by G^* .

Since G^* is obtained from the locally finite graph G' by adding edges, Lemma 2.2 implies that every end of G^* contains a ray in G' . Moreover, it is easy to verify the following claim.

For every vertex t of T_2 the set X_t of all vertices of T_2 above t together with all those vertices of the form v_s for which s lies above t in T_2 has only finitely many neighbours in G^ (outside X_t).* (*)

Thus in particular, every two distinct rays in T_2 starting at \emptyset can be separated in G^* by finitely many vertices. Since every end of G' contains such a ray $R \subseteq T_2$ (because T_2 is end-faithful in G'), it follows that no end of G^* contains distinct ends of G' as subsets. Thus G' is end-faithful in G^* , and hence D' is still a circle in $\overline{G^*}$.

Let H be the subgraph of G^* that consists of all the double rays D'_q with

$q \in J \cup \{0\}$, all the edges of G^* not in G' and, for each finite sequence s , the two edges joining the endvertices $s011$ and $s100$ of e_s to their respective predecessors $s01$ and $s10$ in T_2 . It is easy to check that H is a spanning tree of G^* . So it remains to show that H is end-faithful in G^* . For this, first note that for every two distinct rays R and R' of H starting at \emptyset there are incomparable vertices t and t' of T_2 (ie. none of these vertices lies above the other in T_2) such that $G^*[X_t]$ contains a tail of R and $G^*[X_{t'}]$ contains a tail of R' . Since $G[X_t]$ and $G[X_{t'}]$ are disjoint, $(*)$ implies that R and R' belong to distinct ends of G^* . Thus no end of G^* contains distinct ends of H as subsets. Furthermore, it is easily seen that for every ray R in T_2 there exists a ray in H which is equivalent to R in G^* . As T_2 is end-faithful in G^* , it follows that every end of G^* contains a ray in H . \square

Normal spanning trees, however, do induce topological spanning trees:

Proposition 7.5 *Every normal spanning tree of a graph G induces a topological spanning tree of \overline{G} .*

Proof. Rewrite the proof of Theorem 7.3 with [5, Lemma 4.1] replacing [4, Thm. 5.1]. \square

Proposition 7.5 suggests that, in search of a converse to Proposition 7.1, instead of focussing on the structure of G we might try to characterize the spanning trees that induce topological spanning trees directly (although, of course, in terms of their position within G):

Problem 7.6 *For which end-faithful spanning trees H of an arbitrary infinite graph G is the closure of H in \overline{G} a topological spanning tree of \overline{G} ?*

Here is another question that we have been unable to decide:

Problem 7.7 *Are there connected graphs G such that \overline{G} has no topological spanning tree?*

By Proposition 7.5 and the results of [8], any graph G as in Problem 7.7 must contain certain substructures; in particular, G must be uncountable.

Seymour & Thomas [14] and Thomassen [16] have constructed connected graphs that have no end-faithful spanning tree. From Proposition 7.1 we know that for such graphs G there can be no topological spanning tree T of \overline{G} such that $T \cap \overline{G}$ is connected. But there might be other topological spanning trees, and in all the cases we looked at we managed to find one.

In particular, all the known connected graphs without end-faithful spanning trees have only one end (or contain a one-ended such graph), and for these we do have topological spanning trees:

Proposition 7.8 *If G is a connected graph with only one end, then \overline{G} has a topological spanning tree.*

Proof. We shall construct a spanning forest of G whose components each contain a ray but no double ray. Together with the unique end ω of G , this forest will form a path-connected subspace of \overline{G} that contains no circle, because every circle in \overline{G} is finite or consists of ω together with a double ray.

Such a forest H is easily constructed inductively, as the union of a well-ordered chain of subforests. We start by well-ordering the vertices of G . Then in the induction step we consider the least vertex x not yet covered by our current subforest F . If $G - F$ contains a ray starting at x , we add this ray to F ; if not, we add a finite x - F path.

It is easily checked that H has the desired properties. Indeed, every component C of H contains the ray R that came with its first vertex. And every component of $C - R$ is rayless: otherwise its first vertex should have started a new component of H rather than become part of C . Therefore C , being a tree, contains no double ray. \square

Finally, one might ask whether the topological spanning trees T that we found to exist for \overline{G} under ITOP can always be chosen with $T \cap G$ connected:

Problem 7.9 *When G is a connected graph satisfying (3), does \tilde{G} always have a topological spanning tree whose intersection with G is connected?*

References

- [1] R. Diestel, *Graph Theory* (2nd edition), Springer-Verlag 2000.
<http://www.math.uni-hamburg.de/home/diestel/books/graph.theory/download.html>
- [2] R. Diestel, The end structure of a graph: recent results and open problems, *Discrete Math.* **100** (1992), 313–327.
- [3] R. Diestel, Spanning trees and k -connectedness, *J. Combin. Theory B* **56** (1992), 263–277.
- [4] R. Diestel and D. Kühn, On infinite cycles I, *Combinatorica* (to appear).

- [5] R. Diestel and D. Kühn, On infinite cycles II, *Combinatorica* (to appear).
- [6] R. Diestel and D. Kühn, Graph-theoretical versus topological ends, *J. Combin. Theory B* (to appear).
- [7] R. Diestel and I. Leader, A proof of the bounded graph conjecture, *Invent. math.* **108** (1992), 131–162.
- [8] R. Diestel and I. Leader, Normal spanning trees, Aronszajn trees and excluded minors, *J. London Math. Soc.* (2) **63** (2001), 16–32.
- [9] H. Freudenthal, Über die Enden topologischer Räume und Gruppen, *Math. Zeitschr.* **33** (1931), 692–713.
- [10] R. Halin, Simplicial decompositions of infinite graphs, in: (B. Bollobás, ed.) *Advances in Graph Theory (Annals of Discrete Mathematics 3)*, North-Holland, Amsterdam/London 1978.
- [11] D.W. Hall and G.L. Spencer, *Elementary Topology*, John Wiley, New York 1955.
- [12] C.St.J.A. Nash-Williams, Decompositions of graphs into closed and endless chains, *Proc. London Math. Soc.* (3) **10** (1960), 221–238.
- [13] N. Polat, Topological aspects of infinite graphs, in: (G. Hahn et al., Eds.) *Cycles and Rays*, NATO ASI Ser. C, Kluwer Academic Publishers, Dordrecht 1990.
- [14] P.D. Seymour and R. Thomas, An end-faithful spanning tree counterexample, *Proc. Am. Math. Soc.* **113** (1991), 1163–1171.
- [15] J. Širáň, End-faithful forests and spanning trees in infinite graphs, *Discrete Math.* **95** (1991), 331–340.
- [16] C. Thomassen, Infinite connected graphs with no end-preserving spanning trees, *J. Combin. Theory B* **54** (1992), 322–324.

Reinhard Diestel & Daniela Kühn
 Mathematisches Seminar
 Universität Hamburg
 Bundesstraße 55
 D - 20146 Hamburg
 Germany

Version 24.1.2003

diestel@math.uni-hamburg.de
 kuehn@math.uni-hamburg.de