On End-faithful Spanning Trees in Infinite Graphs

Reinhard Diestel

1. Introduction

Let G be an infinite connected graph. A *ray* (*from* v) in G is a 1-way infinite path in G (with initial vertex v). An infinite connected subgraph of a ray $R \subset G$ is called a *tail* of R. If $X \subset G$ is finite, the infinite component of $R \backslash X$ will be called *the* tail of R in $G \backslash X$.

The following assertions are equivalent for rays $P, Q \subset G$:

(i) There exists a ray $R \subset G$ which meets each of P and Q infinitely often.

(ii) For every finite $X \subset G$, the tails of P and Q in $G \backslash X$ lie in the same component of $G \backslash X$.

(iii) G contains infinitely many disjoint paths connecting a vertex of P with one of Q.

If two rays $P, Q \subset G$ satisfy (i)–(iii), we call them *end-equivalent* (or briefly *equivalent*) in G and write $P \underset{G}{\sim} Q$. An *end* of G is an equivalence class under $\underset{G}{\sim}$, and $\mathcal{E}(G)$ denotes the set of ends of G. For example, the 2-way infinite ladder has two ends, the infinite grid $\mathbb{Z} \times \mathbb{Z}$ and every infinite complete graph have one end, and the dyadic tree has $2^{\aleph_0}$ ends.

This paper is concerned with the relationship between the ends of a connected infinite graph G and the ends of its spanning trees. If T is a spanning tree of G and P, Q are end-equivalent rays in T, then clearly P and Q are also equivalent in G. We therefore have a natural map $\eta : \mathcal{E}(T) \to \mathcal{E}(G)$ mapping each end of T to the end of G containing it. In general, η need be neither 1–1 nor onto. For example, the 2-way infinite ladder has a spanning tree with 4 ends (the tree consisting of its two sides together with one rung), and every infinite complete graph is spanned by a star, which has no ends at all. A spanning tree T of G for which η is 1–1 is said to *respect* the ends of G or called *end-respecting*, and a spanning tree T for which η is onto is called *end-complete*. An end-respecting and end-complete spanning tree is *end-faithful*.

The concept of an end was introduced for graphs by Halin [5] in 1964. It has since inspired some profound work in infinite graph theory; see for example Halin [6], Polat [10, 11], Seifter [12], Watkins [14], or any of several articles in [4]. The problem which Halin originally addressed in [5] is this:

Problem. *Does every infinite connected graph have an end-faithful spanning tree?*

Very recently, Seymour and Thomas [13] have been able to construct graphs which are infinitely connected—and hence have precisely one end—but in which every spanning tree has uncountably many ends. (One of their examples is $TK_{\aleph_1}$-free, a fact which will lend additional relevance to the main result of this paper; see Theorem 1.3 below.) Thus, the general answer to Halin's question is no; the problem remains to understand what makes a graph contain an end-faithful spanning tree, and how these graphs can be recognized.

For countable graphs, a construction of an end-faithful spanning tree was already given as the main result in Halin [5]:

Theorem 1.1. [5] *Every countable connected graph has an end-faithful spanning tree.*

In 1969, Jung [8] investigated end-faithful spanning trees of a particularly intuitive kind: he characterized the graphs G containing a *normal* rooted spanning tree T, one for which every pair of adjacent vertices of G is comparable in the induced partial order $\underset{T}{\leq}$ on $V(G)$ (see below for an exact definition of $\underset{T}{\leq}$). Jung's characterization implies the following sharpening of Theorem 1.1:

Theorem 1.2. [8] *Every countable connected graph has a normal rooted spanning tree.*

The purpose of this paper is to construct an end-faithful spanning tree for any graph, irrespective of its cardinality, that does not contain a subdivided infinite complete graph as a subgraph:

Theorem 1.3. *If G is a connected graph not containing a $TK_{\aleph_0}$, then G has an end-faithful spanning tree T.*

Our construction of the tree T employs a certain decomposition of G into countable factors, which enables us to use the end-faithful spanning trees constructed in the proof of Theorem 1.1. It should be emphasized that the mere existence of the spanning tree T can be shown with considerably less effort by using the stronger Theorem 1.2 instead of Theorem 1.1 (Halin [7]).

The decomposition results needed for our construction of the tree T are presented in Section 2. Most of these results have fairly straightforward proofs, found in [2]. Our key decomposition theorem however, Theorem 2.2, is proved in a separate paper [3]. Section 3 contains the construction of T. In Sections 4 and 5, T is shown to be end-faithful.

The terminology used in this paper is mostly standard, see e.g. [1]. In addition, we shall use the following notations.

If P is a path with vertices x and y, then $P_{x,y}$ denotes the subpath of P from x to y. If $P = x_1 \dots x_n$, then $\mathring{P}$ is the *interior* $x_2 \dots x_{n-1}$ of P. For $X, Y \subset G$, we call a path $P \subset G$ an X–Y path if its endvertices are in X and Y, respectively, and its interior vertices are in $G \setminus (X \cup Y)$.

If T is a rooted tree, with root v_0 say, then T induces a natural partial order $\underset{T}{\leq}$ on its vertices: $v \underset{T}{\leq} w$ if v lies on the unique v_0–w path in T. A ray, for the purpose of this definition, will be assumed to be rooted at its initial vertex.

A graph $H \subset G$ is called *convex* in G if H contains every induced path in G whose endvertices are in H. Equivalently, H is convex in G if and only if the endvertices of every H–H path in G are adjacent in H.

G is *locally finite* if every vertex of G has finite degree. By a well-known theorem of König [9], every infinite but locally finite connected graph contains a ray from each of its vertices.

And finally, if G is a graph and a is a cardinal, the *a-closure* of G is obtained from G by adding all edges $xy \notin E(G)$ for which G contains a independent x–y paths (see [3]).

In our construction of the tree T for Theorem 1.3, we shall use Theorem 1.1 in the following slightly sharper version:

Theorem 1.1′. *If G is a countable connected graph and $F \subset G$ is a finite forest, then G has an end-faithful spanning tree T which contains F.*

Proof. Use Theorem 1.1 to find an end-faithful spanning tree in every component of $G \backslash F$. Extend the union of these trees with F to a spanning tree T of G. T is end-faithful in G.

2. Simplicial decompositions and tree-decompositions.

Let G be a graph, $\sigma > 0$ an ordinal, and let B_λ be an induced subgraph of G for every $\lambda < \sigma$. The family $F = (B_\lambda)_{\lambda<\sigma}$ is called a *simplicial decomposition* of G if the following three conditions hold:

(S1) $G = \bigcup_{\lambda<\sigma} B_\lambda$;

(S2) $\left(\bigcup_{\lambda<\mu} B_\lambda\right) \cap B_\mu =: S_\mu$ is a complete graph for each μ $\qquad (0 < \mu < \sigma)$;

(S3) no S_μ contains B_μ or any other B_λ $\qquad (0 \leq \lambda < \mu < \sigma)$.

For $v \in V(G)$ and $H \subset G$, we denote by $\lambda(v)$ the minimal $\lambda < \sigma$ for which $v \in B_\lambda$, and set $\Lambda(H) := \{\, \lambda(v) \mid v \in V(H)\,\}$. Then $\lambda(v) = \mu$ if and only if $v \in B_\mu \backslash S_\mu$, and $\lambda(v) < \mu$ for all $v \in S_\mu$. For $\mu \leq \sigma$, we write $G|_\mu := \bigcup_{\lambda<\mu} B_\lambda$.

We shall usually refer to a complete graph as a *simplex*, as is the custom in the field. The graphs $S_\mu = G|_\mu \cap B_\mu$ in (S2) will be called simplices *of attachment*.

In a simplicial decomposition, each simplex of attachment S_μ is by definition contained in the union of the factors B_λ, $\lambda < \mu$. In many simplicial decompositions, including all those of finite graphs, each S_μ is even contained in just one of the earlier factors [2]:

(S4) each S_μ is contained in B_λ for some $\lambda < \mu$ $\qquad (\mu < \sigma)$.

When this happens, we denote by $\tau(\mu)$ the minimal $\lambda < \mu$ for which $S_\mu \subset B_\lambda$, and inductively define $\tau^k(\mu) := \tau(\tau^{k-1}(\mu))$, where $\tau^0(\mu) = \mu$.

A family $F = (B_\lambda)_{\lambda<\sigma}$ which satisfies (S1) and (S4) (but not necessarily (S2) or (S3)) is called a *tree-decomposition* of G, and if F satisfies all of (S1)–(S4), it is called a *simplicial tree-decomposition* of G. The reason for this is that we can associate with F a *decomposition tree* $T_F = T_F(G)$, as follows:

$$\begin{aligned} V(T_F) &:= \{ B_\lambda \mid \lambda < \sigma \}, \\ E(T_F) &:= \{ B_\mu B_{\tau(\mu)} \mid \mu < \sigma \}. \end{aligned}$$

The first factor in F, B_0, is taken to be the root of T_F. Thus $B_\mu \underset{T_F}{\leq} B_\nu$ if and only if $\mu = \tau^k(\nu)$ for some $k \geq 0$. Furthermore, $\Lambda(S_\mu) \subset \{ \tau^k(\mu) \mid k \in \mathbb{N} \}$ (induction on μ), so in particular $B_{\lambda(v)} \underset{T_F}{\leq} B_\mu$ for all $v \in B_\mu$ (see [2] for details). If a graph G contains no infinite simplex, its rays are closely related to the rays in its decomposition tree. This fact will be central to our construction of the tree T.

The decompositions we shall use will have another property: they are coherent. A decomposition $(B_\lambda)_{\lambda<\sigma}$ is *coherent* if, for every $\lambda < \sigma$, each vertex of S_λ has a neighbour in $B_\lambda \backslash S_\lambda$, and $B_\lambda \backslash S_\lambda$ is connected.

We now list a number of facts about simplicial decompositions and tree-decompositions that will be used later. The first of these facts is a fundamental property of the factors in a simplicial decomposition.

Proposition 2.1. [2] *If $(B_\lambda)_{\lambda<\sigma}$ is a simplicial decomposition of G, then every B_μ is a convex subgraph of G.*

The next theorem will be our main tool. Its proof is given in [3].

Theorem 2.2. *Let G be an uncountable graph not containing any subdivided infinite simplex. Then the $\aleph_1$-closure G' of G admits a coherent simplicial tree-decomposition $F = (B_\lambda)_{\lambda<\sigma}$ into countable factors and with finite simplices of attachment, which has the following property: for every $\mu < \sigma$ and every edge $xy \in \big(E(B_\mu)\backslash E(S_\mu)\big) \setminus E(G)$, there are uncountably many ordinals ν, with $\tau(\nu) = \mu$, such that B_ν contains an S_ν–S_ν path P with endvertices x, y and $E(P) \subset E(G)$.*

The remaining two propositions concern separation properties in tree-decompositions. Their proofs are straightforward throughout. [2]

Proposition 2.3. *If B, B', B'' are factors in a tree-decomposition F of G and B lies on the B'–B'' path in $T_F(G)$, then B separates $B'\backslash B$ from $B''\backslash B$ in G.*

Proposition 2.4. *Let $F = (B_\lambda)_{\lambda<\sigma}$ be a tree-decomposition of G. Let $\lambda, \mu, \nu < \sigma$ be such that $B_\nu \underset{T_F}{\geq} B_\mu$ but $B_\lambda \underset{T_F}{\not\geq} B_\mu$. Then S_μ separates $B_\lambda \backslash S_\mu$ from $B_\nu \backslash S_\mu$ in G.*

3. The construction of T

Let G be a given uncountable graph, $G \not\supset \mathrm{T}K_{\aleph_0}$. We shall construct an end-faithful spanning tree T of G.

The basic idea for the construction of T is as follows. Using Theorem 2.2, we decompose the $\aleph_1$-closure G' of G into countable factors B_μ, $\mu < \sigma$. By Halin's theorem, we can then find an end-faithful spanning tree T_μ in each of the factors B_μ. Essentially, our task will be to choose the trees T_μ in such a way that they can be pieced together inductively to form T, our desired end-faithful spanning tree of G.

There are various problems we have to be aware of during the construction of T. One of them lies in the fact that the trees T_μ will in general contain edges from $E(G') \setminus E(G)$, which must be replaced by paths in G before T_μ can be incorporated into T. In replacing these edges, we have to ensure that their replacement paths are pairwise independent and avoid the part of T already constructed. Conversely, we shall not be entirely free in choosing T_μ, because replacement paths corresponding to earlier trees T_λ ($\lambda < \mu$) may have spilled over into B_μ and have to be accommodated into T_μ. This problem will be taken care of by Theorem 2.2, which was tailored specifically for this purpose.

Another problem deserving attention is that of stringing the T_μ's together in the right way, so that T does indeed emerge as end-faithful when the construction is complete. To deal with this problem, we shall rely on the close relationship between the ends of G and those of the decomposition tree $T_F(G')$ belonging to our simplicial tree-decomposition F of G'. This relationship is based on the fact that all the simplices of attachment in F are finite: since a ray in G can pass only finitely often through any given S_μ, it must either be 'centred on' (have infinitely many vertices in) one factor B_μ, or follow the course of a ray in T_F. Moreover, equivalent rays in G must follow the same ray in T_F, because their tails cannot be separated by a finite set of vertices (cf. Proposition 2.4). In this way, each end of G induces an end of T_F (or collapses to one vertex B_μ of T_F).

Since equivalent rays in G follow a unique ray in T_F (ore none at all), T can only be end-faithful if it contains, for each end E of T_F, a unique ray Q (from a fixed root v_0) that induces E. The uniqueness of these rays $Q \subset T$ will be ensured by specifying a single vertex s_μ in every S_μ, to serve as a bottle-neck for all paths in T passing from $G'|_\mu$ into $B_\mu \backslash S_\mu$. The existence of the rays Q will be guaranteed by the specific choice of s_μ in S_μ. Figure 1 shows examples of how failure to select bottle-neck vertices $s_\mu \in S_\mu$ at all or a wrong choice of s_μ's may result in a spanning tree T that fails to be end-respecting or end-complete, respectively.

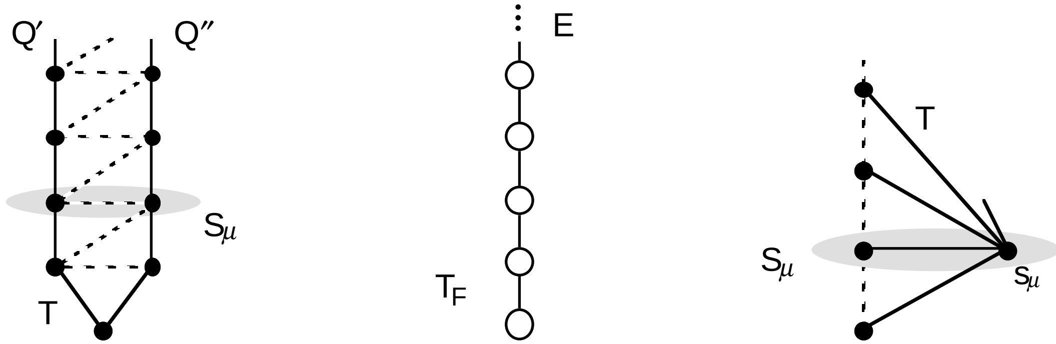

FIGURE 1. End-unfaithful spanning trees in graphs with simplicial decompositions into triangles ∎

Before we begin our formal construction of T, let us run through some of the terms that will be used. $F = (B_\lambda)_{\lambda<\sigma}$ will be a fixed simplicial tree-decomposition of the $\aleph_1$-closure G' of G into countable factors B_λ. The tree T will be obtained as the union of a nested sequence of graphs $T|_\mu$, $\mu < \sigma$. Each $T|_\mu$ will have the following properties:

A$_1(\mu)$: *$T|_\mu \subset G$, $T|_\mu$ is a tree, and $V(T|_\mu) \supset V(G'|_\mu)$. Moreover, $T|_\mu \supset T|_\lambda$ for every $\lambda < \mu$.*

For each $\mu < \sigma$, the definition of $T|_{\mu+1}$ will depend on $T|_\mu$, on another graph T_μ, and on the choice of a certain edge e_μ. Here,

A$_2(\mu)$: *T_μ is an end-faithful spanning tree of $B_\mu \backslash S_\mu$; T_μ may have edges that are not edges of G.*

A$_3(\mu)$: *e_μ is an edge of G' joining T_μ to $T|_\mu \cap S_\mu$ in such a way that $(T_\mu \cup T|_\mu) + e_\mu$ is a tree.*

The term E_μ will denote the set of edges $e \in E(T_\mu) \cup \{ e_\mu \}$ that are not edges of G. When $T|_{\mu+1}$ is formed from the union of $T|_\mu$, T_μ and $\{ e_\mu \}$, these edges e are replaced with independent paths $P(e) \subset G$. By Theorem 2.2, these paths will be chosen in such a way that they run through different factors $B_{\nu(e)}$, $\nu(e) > \mu$, for different edges $e \in E_\mu$. This scattering of the paths $P(e)$ will have the desired effect that any spillover of $T|_\mu$ into $B_\nu \backslash S_\nu$ (for some fixed $\nu \geq \mu$) remains unchanged as μ grows towards ν, and therefore keeps its original form of $P(e)$:

A$_4(\mu)$: *If $\mu \leq \nu < \sigma$ and $T|_\mu \cap B_\nu \not\subset S_\nu$, then $T|_\mu \cap B_\nu = (T|_\mu \cap S_\nu) \cup P$, where P is an S_ν–S_ν path in B_ν. Moreover, there exist $\lambda < \mu$ and an edge $e \in E_\lambda$, such that $\lambda = \tau(\nu)$, $\nu = \nu(e)$ and $P = P(e)$.*

We are now ready to begin the formal construction of T, an end-faithful spanning tree of our graph G.

Let G' be the $\aleph_1$-closure of G, and let $F = (B_\lambda)_{\lambda<\sigma}$ be a simplicial tree-decomposition of G' as provided by Theorem 2.2. Let v_0 be a vertex of B_0; v_0 will be kept fixed throughout the proof and serve as the root of T. Since F is a tree-decomposition, it has a decomposition tree T_F, whose vertices are the factors B_λ, $\lambda < \sigma$. For simplicity, we shall normally use $\leq$ rather than $\underset{T_F}{\leq}$ to denote the natural partial order on $V(T_F)$.

Since any two vertices $x, y \in G'$ with $xy \in E(G')\backslash E(G)$ are joined in G by uncountably many independent paths (recall the definition of the $\aleph_1$-closure), the replacement of such edges with suitable paths in G yields the following:

Lemma 3.1. *If $P, Q \subset G$ are rays with $P \underset{G'}{\sim} Q$, then $P \underset{G}{\sim} Q$.* □

Let $0 \leq \mu \leq \sigma$, and suppose that for all $\lambda < \mu$ we have defined $T|_\lambda$, $T|_{\lambda+1}$, T_λ, e_λ, s_λ and E_λ, together with $\nu(e)$ and $P(e)$ for all $e \in E_\lambda$, such that $\mathbf{A}_1(\lambda)$, $\mathbf{A}_1(\lambda+1)$, $\mathbf{A}_4(\lambda)$ and $\mathbf{A}_4(\lambda+1)$ hold. In the following we shall define the above terms for $\lambda = \mu$ and prove $\mathbf{A}_1$ and $\mathbf{A}_4$ for μ and $\mu+1$, provided that $\mu, \mu+1 < \sigma$ (except that for the definition of $T|_\mu$ and the proofs of $\mathbf{A}_1(\mu)$ and $\mathbf{A}_4(\mu)$ we include the case of $\mu = \sigma$, and that e_μ will only be defined for $\mu > 0$).

We first define $T|_\mu$ and prove $\mathbf{A}_1(\mu)$ and $\mathbf{A}_4(\mu)$. If μ is a successor ordinal, $\mu = \lambda+1$ say, then $T|_\mu$ is already defined, and $\mathbf{A}_1(\mu)$ and $\mathbf{A}_4(\mu)$ hold by assumption. Suppose therefore that μ is not a successor ordinal. If $\mu = 0$, set $T|_\mu := \emptyset$; then $\mathbf{A}_1(\mu)$ and $\mathbf{A}_4(\mu)$ hold trivially. Otherwise, i.e. if μ is a limit ordinal, let

$$T|_\mu := \bigcup_{\lambda<\mu} T|_\lambda \,.$$

$\mathbf{A}_1(\mu)$ is easily seen to follow from our assumption that $\mathbf{A}_1(\lambda)$ holds for all $\lambda < \mu$. To verify $\mathbf{A}_4(\mu)$, notice that if $\mu \leq \nu < \sigma$ and $T|_\mu \cap B_\nu \not\subset S_\nu$, then also $T|_{\lambda'} \cap B_\nu \not\subset S_\nu$ for some $\lambda' < \mu$, since μ is a limit ordinal. By $\mathbf{A}_4(\lambda')$,

$$T|_{\lambda'} \cap B_\nu = (T|_{\lambda'} \cap S_\nu) \cup P\,,$$

where $P = P(e)$ for some $e \in E_\lambda$ with $\lambda < \lambda'$, $\lambda = \tau(\nu)$ and $\nu = \nu(e)$. Let λ'' be any ordinal with $\lambda' \leq \lambda'' < \mu$. Clearly again $T|_{\lambda''} \cap B_\nu \not\subset S_\nu$, because $T|_{\lambda''} \supset T|_{\lambda'}$. Thus, by $\mathbf{A}_4(\lambda'')$, $T|_{\lambda''} \cap B_\nu$ has the form $(T|_{\lambda''} \cap S_\nu) \cup \widetilde{P}$, where $\widetilde{P}$ is an S_ν–S_ν path in B_ν. But $\widetilde{P} \supset P$ (again by $T|_{\lambda''} \supset T|_{\lambda'}$), and P is also an S_ν–S_ν path. Therefore $\widetilde{P} = P = P(e)$. Thus

$$\begin{aligned} T|_\mu \cap B_\nu &= \bigcup_{\lambda' \leq \lambda'' < \mu} (T|_{\lambda''} \cap B_\nu) \\ &= \bigcup_{\lambda' \leq \lambda'' < \mu} (T|_{\lambda''} \cap S_\nu) \cup P(e) \\ &= (T|_\mu \cap S_\nu) \cup P(e)\,, \end{aligned}$$

where $P(e)$ is an S_ν–S_ν path in B_ν and $e \in E_\lambda$ with $\lambda < \lambda' < \mu$, $\lambda = \tau(\nu)$ and $\nu = \nu(e)$. This completes the proof of $\mathbf{A}_4(\mu)$.

For the remaining definitions and the proofs of $\mathbf{A}_1(\mu+1)$ and $\mathbf{A}_4(\mu+1)$, we shall assume that $\mu \neq \sigma$. Let us say that μ is of *type 1* if $T|_\mu \cap B_\mu \subset S_\mu$, and of *type 2* otherwise.

For the definition of T_μ, recall that since F is coherent, $B_\mu \backslash S_\mu$ is connected. If μ is of type 1, then $T|_\mu \cap (B_\mu \backslash S_\mu) = \emptyset$, and we let T_μ be any end-faithful spanning tree of $B_\mu \backslash S_\mu$. If μ is of type 2, then $T|_\mu \cap (B_\mu \backslash S_\mu) = \mathring{P}(e)$ for some $e \in E_\lambda$ and $\lambda < \mu$ (by $\mathbf{A}_4(\mu)$), and we let T_μ be an end-faithful spanning tree of $B_\mu \backslash S_\mu$ containing $\mathring{P}(e)$. These choices of T_μ are possible by Theorems 1.1 and 1.1′, and they satisfy $\mathbf{A}_2(\mu)$.

Next we define e_μ and s_μ. Set $s_0 := v_0$, and assume in the sequel that $\mu > 0$. If μ is of type 1, we let e_μ be any edge xy of G' with $\lambda(x) = \tau(\mu)$ and $\lambda(y) = \mu$. Notice that this choice of e_μ is always possible: by definition of $\tau(\mu)$, S_μ has a vertex x with $\lambda(x) = \tau(\mu)$, and x has a neighbour y in $B_\mu \backslash S_\mu$, because F is coherent. If μ is of type 2 on the other hand, we let e_μ be the unique edge that lies on the v–v_0 path in $T|_\mu$ for every $v \in \mathring{P}(e)$ ($e \in E_\lambda$ as earlier). The existence and uniqueness of such an edge follow from the fact that $T|_\mu$ is a tree ($\mathbf{A}_1(\mu)$) and that every $v \in \mathring{P}(e)$ has degree 2 in $T|_\mu$ ($\mathbf{A}_4(\mu)$). In each case, $(T_\mu \cup T|_\mu) + e_\mu$ is a tree ($\mathbf{A}_3(\mu)$).

Notice also that in both cases the definition of e_μ is such that e_μ has one endvertex in S_μ and one in $B_\mu \backslash S_\mu$; we let s_μ be the endvertex of e_μ in S_μ. Then

$\mathbf{B}(\mu)$: *For every $v \in B_\mu \backslash S_\mu$, the vertex s_μ lies on the v–v_0 path in the tree $(T_\mu \cup T|_\mu) + e_\mu$.*

Let

$$E_\mu := \Big(E(T_\mu) \cup \{ e_\mu \} \Big) \setminus E(G) \,.$$

Using the property of F given by Theorem 2.2, we now choose for each edge $xy \in E_\mu$ an ordinal $\nu =: \nu(e)$ with $\tau(\nu) = \mu$, such that B_ν contains an S_ν–S_ν path P with endvertices x, y and $E(P) \subset E(G)$; the path P will be denoted by $P(e)$. Moreover, we choose the ordinals $\nu(e)$ in such a way that $\nu(e) \neq \nu(e')$ for distinct $e, e' \in E_\mu$; this is again possible by Theorem 2.2, because $|E_\mu| \leq |B_\mu|^2 \leq \aleph_0$. Since $\lambda(v) = \nu(e) > \mu$ for every $v \in \mathring{P}(e)$ with $e \in E_\mu$, the following holds:

$\mathbf{C}(\mu)$: *If $e, e' \in E_\mu$ and $e \neq e'$, then $\mathring{P}(e) \cap \mathring{P}(e') = \emptyset$ and $\mathring{P}(e) \cap G'|_{\mu+1} = \emptyset$.*

It remains to define $T|_{\mu+1}$ and to prove $\mathbf{A}_1(\mu+1)$ and $\mathbf{A}_4(\mu+1)$. Let us set

$$T|_{\mu+1} := \Big(\big((T_\mu \cup T|_\mu) + e_\mu \big) \cup \bigcup_{e \in E_\mu} P(e) \Big) - E_\mu \,.$$

In order to prove $\mathbf{A}_1$ and $\mathbf{A}_4$ for $\mu+1$, observe first that the sets

$$N_\lambda := \{ \nu(e) \mid e \in E_\lambda \}$$

are disjoint for distinct values of $\lambda \leq \mu$, because $\tau(\nu) = \lambda$ for all $\nu \in N_\lambda$ (by definition of $\nu(e)$). In particular,

$$e \in E_\lambda,\ e' \in E_{\lambda'},\ \lambda \neq \lambda' \quad \Rightarrow \quad \mathring{P}(e) \cap \mathring{P}(e') = \emptyset \,. \tag{1}$$

By $\mathbf{A}_4(\mu)$, any vertex v of $T|_\mu \setminus G'|_{\mu+1}$ must be on some $P(e)$ with $e \in E_\lambda,\ \lambda < \mu$. Hence (1) implies that $(T|_\mu \setminus G'|_{\mu+1}) \cap \mathring{P}(e) = \emptyset$ for all $e \in E_\mu$. Combining this with $\mathbf{C}(\mu)$, we obtain

$\mathbf{D}(\mu)$: $\quad T|_\mu \cap \mathring{P}(e) = \emptyset$, *for all* $e \in E_\mu$.

As $(T_\mu \cup T|_\mu) + e_\mu$ is a tree, and the paths $\mathring{P}(e)$, $e \in E_\mu$, are pairwise disjoint and avoid T_μ, $\mathbf{D}(\mu)$ implies that $T|_{\mu+1}$ is a tree. This establishes $\mathbf{A}_1(\mu+1)$, the other assertions being obvious.

For the proof of $\mathbf{A}_4(\mu+1)$, let ν with $\mu+1 \leq \nu < \sigma$ and $T|_{\mu+1} \cap B_\nu \not\subset S_\nu$ be given. If $\nu \in N_\mu$, say $\nu = \nu(e)$ with $e \in E_\mu$, then $\nu \notin N_\lambda$ for all $\lambda < \mu$, and hence $T|_\mu \cap B_\nu \subset S_\nu$ by $\mathbf{A}_4(\mu)$. Thus $T|_\mu \cap (B_\nu \setminus S_\nu) = \emptyset$, and therefore $T|_{\mu+1} \cap (B_\nu \setminus S_\nu) = \mathring{P}(e)$. This implies

$$T|_{\mu+1} \cap B_\nu = (T|_{\mu+1} \cap S_\nu) \cup P(e)\,,$$

as desired. On the other hand if $\nu \notin N_\mu$, then $T|_{\mu+1} \cap (B_\nu \setminus S_\nu) \subset T|_\mu$, so

$$T|_{\mu+1} \cap (B_\nu \setminus S_\nu) = T|_\mu \cap (B_\nu \setminus S_\nu) = \mathring{P}(e)$$

for some $e \in E_\lambda$ and $\lambda < \mu$, again by $\mathbf{A}_4(\mu)$. Thus again

$$T|_{\mu+1} \cap B_\nu = (T|_{\mu+1} \cap S_\nu) \cup P(e)\,,$$

completing the proof of $\mathbf{A}_4(\mu+1)$.

Let us finally set

$$T := T|_\sigma\,.$$

By $\mathbf{A}_1(\sigma)$, T is a spanning tree of G.

The proof that T is end-faithful with respect to G will be given in Sections 4 and 5. In the remainder of this section we shall extract a few facts from the construction of T for later use. Unless otherwise stated, each of these facts holds for every $\mu < \sigma$.

The first fact concerns the edges in the sets E_μ.

$\mathbf{E}(\mu)$: *If* $e = xy \in E_\mu$ *and* $\lambda(x) \leq \lambda(y)$, *then either* $\lambda(x) = \lambda(y) = \mu$, *or* $\lambda(x) = \tau(\mu)$, $\lambda(y) = \mu$ *and* $x = s_\mu$.

For the proof of $\mathbf{E}(\mu)$, notice first that if $e \in E(T_\mu)$, then $\lambda(x) = \lambda(y) = \mu$ by $T_\mu \subset B_\mu \setminus S_\mu$. Suppose therefore that $e \notin E(T_\mu)$, i.e. that $e = e_\mu$. Then μ must be of type 1, since otherwise e_μ would be in $E(T|_\mu) \subset E(G)$, and hence not in E_μ. Therefore $\lambda(x) = \tau(\mu)$, $\lambda(y) = \mu$ and $x = s_\mu$ by definition of e_μ and s_μ.

The second fact contains the information ensuring that the definition of s_μ achieves its purpose; see our earlier informal discussion.

$\mathbf{F}(\mu)$: $\quad \lambda(s_\mu) \in \{\,\tau(\mu), \tau^2(\mu)\,\}$, *and if* $\lambda(s_\mu) = \tau^2(\mu)$, *then* $s_\mu = s_{\tau(\mu)}$ *(for* $\mu > 0$*).*

The proof of $\mathbf{F}(\mu)$ is clear by definition of e_μ if μ is of type 1 (and hence $\lambda(s_\mu) = \tau(\mu)$). If μ is of type 2, then s_μ is an endvertex of $P(e)$ and hence of e for some $e \in E_\lambda$ with

$\lambda = \tau(\mu)$. By $\mathbf{E}(\lambda)$, this implies that either $\lambda(s_\mu) = \lambda = \tau(\mu)$, or else $\lambda(s_\mu) = \tau(\lambda) = \tau^2(\mu)$ and $s_\mu = s_\lambda = s_{\tau(\mu)}$.

Let us note the following immediate consequence of (F).

$$\text{If } B_{\lambda_0} B_{\lambda_1} \dots \text{ is a ray from } B_0 \text{ in } T_F\text{, then } \{\, s_{\lambda_i} \mid i \in \mathbb{N} \,\} \text{ is infinite.} \tag{2}$$

Indeed, as $B_{\lambda_0} = B_0$, we have $\tau(\lambda_n) = \lambda_{n-1}$ for all $n \in \mathbb{N}$. Therefore $\lambda(s_{\lambda_n}) \in \{\, \lambda_{n-1}, \lambda_{n-2} \,\}$ (by $\mathbf{F}(\lambda_n)$), for all $n \in \mathbb{N}$. This implies (2).

Next, we show that T contains essentially no B_μ–B_μ paths other than those of the form $P(e)$, $e \in E_\mu$.

G(μ): *If P is a B_μ–B_μ path in T, with endvertices x and y where $y \in B_\mu \backslash S_\mu$, then $xy \in E_\mu$ and $P = P(xy)$.*

To prove $\mathbf{G}(\mu)$, we first show that $P \subset T|_{\mu+1}$. As $x, y \in V(T|_{\mu+1})$ and $T|_{\mu+1}$ is a tree, x and y are joined by a path in $T|_{\mu+1}$. Since $T|_{\mu+1} \subset T$ and T contains only one x–y path, this path must be P. Notice that S_μ separates y from every vertex $v \notin S_\mu$ with $\lambda(v) < \mu$ (Proposition 2.4), while S_μ does not separate y from any $v \in \mathring{P}$. Therefore $\lambda(v) \geq \mu$ and hence $\lambda(v) > \mu$ for all $v \in \mathring{P}$. By definition of $T|_{\mu+1}$, this implies that every $v \in \mathring{P}$ is contained in $\mathring{P}(e)$ for some $e \in E_\mu$. But inner vertices of different paths $P(e)$ cannot be adjacent, since $\nu(e) \neq \nu(e')$ and $\tau(\nu(e)) = \tau(\nu(e')) = \mu$ for distinct $e, e' \in E_\mu$. Therefore all $v \in \mathring{P}$ are on the same path $P(e)$, i.e. $\mathring{P} \subset P(e)$ for some $e \in E_\mu$. Since P and $P(e)$ are both B_μ–B_μ paths, this means that $P = P(e)$ and $xy = e$.

$\mathbf{G}(\mu)$ has the following useful consequence: if $P \subset T$ is a path that meets B_μ infinitely often but avoids S_μ, we can turn P into a path $P' \subset T_\mu$ by replacing each B_μ–B_μ path $P_{x,y} \subset P$ with the edge $xy \in E_\mu$. Similarly, we can contract *any* path $P \subset G$ onto a path $P' \subset B_\mu$—recall that since B_μ is a convex subgraph of G (Proposition 2.1), the endvertices of P will be adjacent in B_μ:

H(μ): *If $P = v_1 v_2 \dots$ is a (finite or infinite) path in G, then B_μ contains a path $P' = v_{k_1} v_{k_2} \dots$, where $V(P') = V(P) \cap V(B_\mu)$ and $k_i < k_j$ if and only if $i < j$. Moreover, if $P \subset T \backslash S_\mu$, then $P' \subset T_\mu$.*

Finally, we prove what was earlier desribed as the 'bottle-neck' property of the vertices s_μ.

I(μ): *If $\mu = \tau^k(\nu)$, $k \geq 0$, and $v \in G$ with $\lambda(v) = \nu$, then s_μ separates v from v_0 in T (for $\mu > 0$).*

We prove $\mathbf{I}(\mu)$ be induction on k. If $k = 0$, then $\nu = \mu$, so $v \in T_\mu$. Let P be the v–v_0 path in the tree $(T_\mu \cup T|_\mu) + e_\mu$. By $\mathbf{B}(\mu)$, $s_\mu \in P$. Replacing every edge $e \in E(P) \cap E_\mu$ with $P(e)$, we obtain a path in $T|_{\mu+1}$ that joins v to v_0 and contains s_μ (cf. $\mathbf{D}(\mu)$). Since $T|_{\mu+1} \subset T$ and T is a tree, this implies the assertion.

Suppose now that $k > 0$, and that $\mathbf{I}(\mu)$ holds for all smaller values of k. Let P be a v–v_0 path in T; we have to show that $s_\mu \in P$. By the case of $k = 0$, v is separated from v_0 by s_ν in T, so P contains an s_ν–v_0 path P'. If $\lambda(s_\nu) = \tau(\nu)$, then $\mu = \tau^{k-1}(\lambda(s_\nu))$, so

$s_\mu \in P'$ by the induction hypothesis. Suppose therefore that $\lambda(s_\nu) \neq \tau(\nu)$. Then, by $\mathbf{F}(\nu)$, $\lambda(s_\nu) = \tau^2(\nu)$ and $s_\nu = s_{\tau(\nu)}$. If $\tau^2(\nu) \geq \mu$, the assertion again follows by the induction hypothesis (as above). But if $\tau^2(\nu) < \mu$, then $k = 1$ and $\mu = \tau(\nu)$, giving $s_\nu = s_{\tau(\nu)} = s_\mu$. This completes the proof of $\mathbf{I}(\mu)$.

4. *T* is end-complete

It will be convenient in this and the next section to call a ray $P \subset G$ *centred on* B_μ if P meets B_μ infinitely often, and *uncentred* if $V(P) \cap V(B_\mu)$ is finite for every $\mu < \sigma$.

Let P be a given ray in G; we have to find a ray $Q \subset T$ such that $P \underset{G}{\sim} Q$.

We first suppose that P is centred on some B_μ, $\mu < \sigma$. Since any ray Q is equivalent to P as soon as it is equivalent to some tail of P, we may assume that $P \cap S_\mu = \emptyset$. Let $P' \subset B_\mu$ be the path obtained from P by $\mathbf{H}(\mu)$; then P' is a ray in $B_\mu \backslash S_\mu$. As T_μ is an end-faithful spanning tree of $B_\mu \backslash S_\mu$, it contains a ray Q' with $P' \underset{B_\mu \backslash S_\mu}{\sim} Q'$, and hence $P' \underset{G'}{\sim} Q'$. Let Q be obtained from Q' by replacing each edge $e \in E(Q') \cap E_\mu$ with the path $P(e)$. Then $Q \subset T|_{\mu+1} \subset T$, and Q is a ray (by $\mathbf{C}(\mu)$). As $P' \underset{G'}{\sim} Q'$, clearly $P \underset{G'}{\sim} Q$. By Lemma 3.1, this implies that $P \underset{G}{\sim} Q$.

Let us from now on suppose that P is uncentred. Recall that the factors B_λ in our decomposition F of G' are the vertices of the decomposition tree T_F. For given rays $P \subset G$ and $R = B_{\lambda_0} B_{\lambda_1} \ldots \subset T_F$, let us set

$$I(P,R) := \{\, i \in \mathbb{N} \mid \lambda_i \in \Lambda(P) \,\}.$$

The following two lemmas relate the uncentred rays in G to rays in T_F.

Lemma 4.1. *For every uncentred ray $P \subset G$, there exists a ray R from B_0 in T_F such that $I(P,R)$ is infinite.*

Proof. Let $T(P)$ be the union of all B_λ–B_0 paths in T_F with $\lambda \in \Lambda(P)$. $T(P)$ is a subtree of T_F. For vertices B', B'' of T_F, let us say that B' *precedes* B'' if $B' < B''$ but there is no vertex $B \in T_F$ of the form $B = B_\lambda$, $\lambda \in \Lambda(P)$, such that $B' < B < B''$. Let us prove the following:

$$\text{Each vertex of } T(P) \text{ precedes at most finitely many vertices } B_\lambda \in T(P) \text{ with } \lambda \in \Lambda(P). \tag{3}$$

Suppose (3) fails, and let B be a vertex of $T(P)$ that precedes every vertex in some infinite set $U \subset \{\, B_\lambda \mid \lambda \in \Lambda(P) \,\}$. We show that whenever B', B'' are distinct elements of U, any subpath $P' = v \ldots w$ of P with $v \in B' \backslash B$ and $w \in B'' \backslash B$ passes through B. Since P has a vertex in $B_\lambda \backslash B$ for every $B_\lambda \in U$ (by $B_\lambda > B$ and the definition of U), this means that P meets B infinitely often, contrary to our assumption that P is uncentred.

Let B', B'' and P' be given as stated. B' and B'' are incomparable in T_F, because B precedes both of them. Let $\widetilde{B}$ be the maximal vertex of T_F that satisfies $\widetilde{B} < B'$ as well as $\widetilde{B} < B''$. Then $\widetilde{B} \geq B$. Moreover, $\widetilde{B}$ lies on the B'–B'' path in T_F, so $P' \cap \widetilde{B} \neq \emptyset$ by Proposition 2.3. Let $u \in P' \cap \widetilde{B}$. Clearly $B_{\lambda(u)} \leq \widetilde{B}$, so $B_{\lambda(u)}$ and B are comparable, because also $B \leq \widetilde{B}$. Since $B_{\lambda(u)} \leq \widetilde{B} < B'$ and $B_{\lambda(u)} \leq \widetilde{B} < B''$ but B precedes B' and B'', we have $B \not< B_{\lambda(u)}$, and therefore $B \geq B_{\lambda(u)}$ (see the first part of Figure 2). Thus B lies on the B'–$B_{\lambda(u)}$ path in T_F, which implies that $P_{u,v} \cap B \neq \emptyset$ (by Proposition 2.3). This completes the proof of (3).

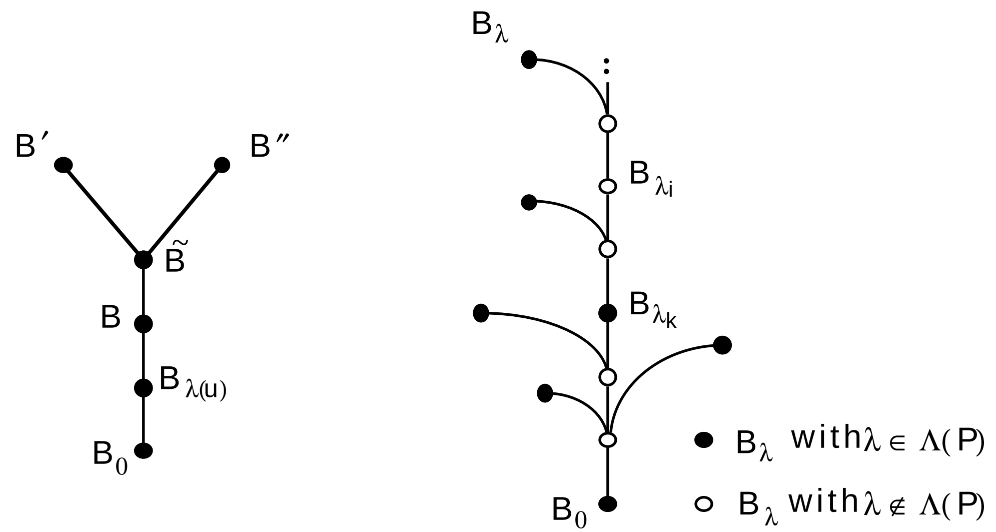

FIGURE 2. Vertices of $T(P)$ in the proof of (3)

As a first consequence of (3), we note that $T(P)$ is locally finite: if $B \in T(P)$ has infinitely many neighbours $B' > B$, and B'_λ is the smallest vertex of $T(P)$ with $B' \leq B'_\lambda$ and $\lambda \in \Lambda(P)$ (recall the construction of $T(P)$), then B precedes B'_λ, and all these B'_λ's are distinct. By König's theorem, $T(P)$ therefore contains a ray $R =: R(P)$, say $R = B_{\lambda_0} B_{\lambda_1} \ldots$. Since $B_0 \in T(P)$, we may assume that $\lambda_0 = 0$.

It remains to show that infinitely many of the indices λ_i are in $\Lambda(P)$. Suppose not, and let k be maximal with $\lambda_k \in \Lambda(P)$. Then no B_{λ_i} with $i > k$ is of the form B_λ, $\lambda \in \Lambda(P)$, but, by construction of $T(P)$, each of them precedes such a B_λ, and is in turn preceded by B_{λ_k} (see the second part of Figure 2). Since for every $B_\lambda \in T_F$ there are only finitely many vertices $B \in T_F$ with $B_\lambda > B$ but $\{ B_{\lambda_i} \mid i > k \}$ is infinite, this means that B_{λ_k} precedes infinitely many vertices B_λ with $\lambda \in \Lambda(P)$, contrary to (3). $\square$

Lemma 4.2. *If P is an uncentred ray in G and $R = B_{\lambda_0} B_{\lambda_1} \ldots$ is a ray from B_0 in T_F such that $I = I(P, R)$ is infinite, then every tail of P meets every S_{λ_n} with sufficiently large $n \in \mathbb{N}$.*

Proof. Let P' be a given tail of P. Pick $k \in I'$, where

$$I' := I \setminus \{\, i \in I \mid \lambda_i \in \Lambda(P \setminus P')\,\}.$$

Since $P \setminus P'$ is finite but I is infinite, I' is also infinite, and $\lambda_i \in \Lambda(P')$ for every $i \in I'$. We show that P' meets every S_{λ_n} with $k < n \in \mathbb{N}$. Let such n be given, and let $i \geq n$ with $i \in I'$. Let $u, v \in P'$ with $\lambda(u) = \lambda_k$ and $\lambda(v) = \lambda_i$. Since $\lambda_k < \lambda_n$, we have $u \in G|_{\lambda_n}$, so v and u are separated by S_{λ_n} in G' unless u is itself in S_{λ_n} (by Proposition 2.4). Hence $P'_{u,v} \cap S_{\lambda_n} \neq \emptyset$ as claimed. $\square$

Equipped with Lemmas 4.1 and 4.2, we can now tackle the second case in our proof of the end-completeness of T in G. Let P be a given uncentred ray in G, and let $R = B_{\lambda_0} B_{\lambda_1} \ldots$ be as provided by Lemma 4.1. Define $Q := \bigcup_{n=0}^{\infty} Q_n$, where Q_n is the s_{λ_n}–$s_{\lambda_{n+1}}$ path in T (and $s_{\lambda_0} := v_0$). We shall first prove that Q is a ray, and then show that $Q \underset{G}{\sim} P$.

In order to prove that Q is a path, we show that for different $n \geq 0$ the vertices $v \in Q_n - s_{\lambda_n}$ have distinct $\lambda(v)$, so the paths $Q_n - s_{\lambda_n}$ must be disjoint for different n. For every $\lambda < \sigma$, let λ' be such that $B_{\lambda'}$ is the maximal vertex on R with $B_{\lambda'} \leq B_\lambda$. (To see that λ' exists, recall that R is a ray from B_0, which is the root of T_F.) We prove the following:

$$v \in Q_n - s_{\lambda_n},\ n \geq 0 \quad \Rightarrow \quad \lambda'(v) = \lambda_n \tag{4}$$

We apply induction on n. Let $n \geq 0$, and suppose the assertion holds for all $i < n$. Then $\bigcup_{i=0}^{n-1} Q_i$ is (empty or) the v_0–s_{λ_n} path in T. By $\mathbf{F}(\lambda_{n+1})$, we have either $\lambda(s_{\lambda_{n+1}}) = \lambda_n$ or $s_{\lambda_{n+1}} = s_{\lambda_n}$; recall that $\tau(\lambda_{n+1}) = \lambda_n$. In the latter case $Q_n - s_{\lambda_n} = \emptyset$, so there is nothing to prove; we shall therefore assume that $\lambda(s_{\lambda_{n+1}}) = \lambda_n$. Then $s_{\lambda_{n+1}} \in B_{\lambda_n} \setminus S_{\lambda_n}$. Recall that e_{λ_n} is an edge of G' joining s_{λ_n} to a vertex of $B_{\lambda_n} \setminus S_{\lambda_n}$, say to x. Let Q'_n be the s_{λ_n}–$s_{\lambda_{n+1}}$ path in B_{λ_n} consisting of e_{λ_n} followed by the x–$s_{\lambda_{n+1}}$ path in T_{λ_n}. (Recall that T_{λ_n} is a spanning tree of $B_{\lambda_n} \setminus S_{\lambda_n}$.) If we replace every edge $e \in E(Q'_n) \cap E_{\lambda_n}$ with the path $P(e) \subset T$, we obtain an s_{λ_n}–$s_{\lambda_{n+1}}$ path in T (cf. $\mathbf{C}(\lambda_n)$). Since T is a tree, this is the unique s_{λ_n}–$s_{\lambda_{n+1}}$ path in T, and therefore equal to Q_n.

To complete the proof of (4), it remains to show that $\lambda'(v) = \lambda_n$ for every $v \in \mathring{P}(e)$ with $e \in E(Q'_n) \cap E_{\lambda_n}$. By definition of $P(e)$, we have $\lambda(v) = \nu(e)$ for $v \in \mathring{P}(e)$ and $\tau(\nu(e)) = \lambda_n$. Therefore $\lambda'(v) = \lambda_n$ unless $\lambda'(v) = \lambda(v) = \lambda_{n+1}$.

Before we show that $\lambda'(v)$ must be λ_n rather than λ_{n+1}, let us note that certainly $\lambda'(v) > \lambda_i$ for all $i < n$, and therefore $v \notin Q_0 \cup \ldots \cup Q_{n-1}$ by the induction hypothesis. Thus $\mathring{Q}_n \cap \bigcup_{i=0}^{n-1} Q_i = \emptyset$, and $Q_0 \cup \ldots \cup Q_n$ is the v_0–$s_{\lambda_{n+1}}$ path in T.

Let us now resume our proof that $\lambda'(v) = \lambda_n$ for any given $v \in \mathring{P}(e)$ and $e \in E(Q'_n) \cap E_{\lambda_n}$. As shown above, all we have to check is that $\lambda(v) \neq \lambda_{n+1}$. This, however, follows from the definition of $e_{\lambda_{n+1}}$ and $s_{\lambda_{n+1}}$: if $\lambda(v) = \lambda_{n+1}$, then λ_{n+1} is of type 2 (because $\mathring{P} \subset T|_{\lambda_{n+1}} \cap (B_{\lambda_{n+1}} \setminus S_{\lambda_{n+1}})$) and $s_{\lambda_{n+1}}$ lies on the v–v_0 path in $T|_{\lambda_{n+1}} \subset T$—which contradicts the fact that $v \in \mathring{Q}_n$ and $Q_0 \cup \ldots \cup Q_n$ is the v_0–$s_{\lambda_{n+1}}$ path in T. This completes the proof of (4), showing that Q is a path.

The proof that Q is infinite, and therefore a ray, is now straightforward. As $V(Q) \supset \{ s_{\lambda_n} \mid n \in \mathbb{N} \}$, Q can only be finite if infinitely many s_{λ_n}'s coincide. This however is ruled out by (2).

We have shown that Q is a ray in T that passes through every vertex s_{λ_i}, $i \in \mathbb{N}$. In order to prove that Q is equivalent to P in G, let U be a given finite set of vertices of G, and let P' and Q' be the tails of P and Q in $G - U$, respectively. We have to show that $G - U$ contains a P'–Q' path. By Lemma 4.2 and the definition of Q, we can find an $n \in \mathbb{N}$ such that $P' \cap S_{\lambda_n} \neq \emptyset$, say $s \in P' \cap S_{\lambda_n}$, and $s_{\lambda_n} \in Q'$. If $s = s_{\lambda_n}$ or $ss_{\lambda_n} \in E(G)$, we are done. But otherwise $ss_{\lambda_n} \in E(G') \backslash E(G)$, so G contains uncountably many independent s–s_{λ_n} paths, one of which avoids U.

5. *T* is end-respecting

Lemma 5.1. *Suppose P_1, P_2 are rays in G, P_1 is centred on B_μ ($\mu < \sigma$), and $P_1 \underset{G}{\sim} P_2$. Then P_2 is also centred on B_μ.*

Proof. If P_2 is centred at all, say on ν, then clearly $\nu = \mu$: if $\nu \neq \mu$, without loss of generality $\nu > \mu$, then S_ν separates infinitely many vertices of P_1 from infinitely many vertices of P_2 in G' (Proposition 2.4), which contradicts our assumption that $P_1 \underset{G}{\sim} P_2$.

Suppose therefore that P_2 is not centred. Let $R = B_{\lambda_0} B_{\lambda_1} \ldots$ be a ray in T_F such that $I(P_2, R)$ is infinite (Lemma 4.1). For every $i \in I(P_2, R)$, let v_i be a vertex on P_2 with $\lambda(v_i) = \lambda_i$. As at most finitely many B_{λ_n} can be such that $B_\mu \geq B_{\lambda_n}$, there exists $k \in \mathbb{N}$ with $B_\mu \not\geq B_{\lambda_k}$. By Proposition 2.4, S_{λ_k} separates every v_i with $i \geq k$ from every vertex $u \in B_\mu \backslash S_{\lambda_k}$. Thus the finite set $V(S_{\lambda_k})$ separates infinitely many vertices of P_1 from infinitely many vertices of P_2 in G', again contradicting $P_1 \underset{G}{\sim} P_2$. $\square$

For our proof that T respects the ends of G, let rays $P_1, P_2 \subset T$ with $P_1 \underset{G}{\sim} P_2$ be given. We have to show that $P_1 \underset{T}{\sim} P_2$. We shall distinguish two cases: that P_1 and P_2 are both centred on the same B_μ, and that P_1 and P_2 are both uncentred. By Lemma 5.1 this distinction is exhaustive.

Case 1: *P_1 and P_2 are centred on B_μ, $\mu < \sigma$.*

Since S_μ is finite, we may assume without loss of generality that $P_i \cap S_\mu = \emptyset$, $i = 1, 2$. Let us choose an infinite sequence $Q_1, Q_2, \ldots$ of disjoint paths in G as follows. Having definied $Q_1, \ldots, Q_{n-1}$ for some $n \in \mathbb{N}$, consider the tails P_1^n and P_2^n of P_1 and P_2 in $G \backslash \bigcup_{i=1}^{n-1} Q_i$. If $P_1^n \cap P_2^n \neq \emptyset$, set $Q_n := \{ q_n \}$ for some vertex $q_n \in P_1^n \cap P_2^n$. If $P_1^n \cap P_2^n = \emptyset$, let

$Q_n = q_n^1 \dots q_n^2$ be a path in $G \setminus (S_\mu \cup \bigcup_{i=1}^{n-1} Q_i)$ with $q_n^i \in P_i \cap B_\mu$, $i = 1, 2$. Note that Q_n exists, because by assumption P_1 and P_2 are equivalent in G and centred on B_μ.

Let P_1' and P_2' correspond to P_1 and P_2 as in $\mathbf{H}(\mu)$, and let Q_n' correspond to Q_n, $n \in \mathbb{N}$. Then $Q_1', Q_2', \dots$ is an infinite sequence of disjoint paths in $B_\mu \setminus S_\mu$, each joining a vertex of P_1' to one of P_2'. Therefore $P_1' \underset{B_\mu \setminus S_\mu}{\sim} P_2'$. As $P_i \subset T \setminus S_\mu$ by assumption, $\mathbf{H}(\mu)$ implies that $P_i' \subset T_\mu$, $i = 1, 2$. But T_μ is an end-respecting spanning tree of $B_\mu \setminus S_\mu$, so $P_1' \underset{T_\mu}{\sim} P_2'$. As T_μ is a tree, this equivalence means that $V(P_1') \cap V(P_2')$ is infinite. Therefore $V(P_1) \cap V(P_2)$ is infinite, too (recall that $V(P_i') \subset V(P_i)$, $i = 1, 2$), so $P_1 \underset{T}{\sim} P_2$ as claimed.

Case 2: *P_1 and P_2 are both uncentred.*

We shall assume, without loss of generality, that P_1 and P_2 are rays from v_0. For $i = 1, 2$, let R_i be a ray from B_0 in T_F such that $I_i := I(P_i, R_i)$ is infinite (by Lemma 4.1).

Let us use the equivalence of P_1 and P_2 in G to show that $R_1 = R_2$. Suppose $R_1 \neq R_2$, and let B_μ be the first (= minimal) vertex of R_1 that is not on R_2. Then $B \not\geq B_\mu$ for every $B \in V(R_2)$, because R_2 is a ray from B_0. By Proposition 2.4, therefore, and the fact that I_1 and I_2 are infinite, S_μ separates infinitely many vertices of P_1 from infinitely many vertices of P_2 in G'. As S_μ is finite, this contradicts our assumption that $P_1 \underset{G}{\sim} P_2$.

Thus $R_1 = R_2 =: R$, say $R = B_{\lambda_0} B_{\lambda_1} \dots$. As I_1 and I_2 are infinite, $\mathbf{I}(\lambda_n)$ implies that P_1 and P_2 contain s_{λ_n}, for every $n \in \mathbb{N}$. (Recall that P_1 and P_2 are rays from v_0.) By (2) therefore, P_1 and P_2 have infinitely many vertices in common, giving $P_1 \underset{T}{\sim} P_2$ as desired.

References

[1] B. Bollobás, *Graph Theory, An Introductory Course*, Springer-Verlag, New York 1979.

[2] R. Diestel, *Graph Decompositions—a study in infinite graph theory*, Oxford University Press, Oxford 1990.

[3] R. Diestel, The structure of $\mathrm{T}K_a$-free graphs, submitted.

[4] R. Diestel (Ed.), *Directions in Infinite Graph Theory*, Annals of Discrete Mathematics, in preparation.

[5] R. Halin, Über unendliche Wege in Graphen, *Math. Ann.* **157** (1964), 125–137.

[6] R. Halin, Automorphisms and endomorphisms of infinite locally finite graphs, *Abh. Math. Sem. Univ. Hamburg* **39** (1973), 251–283.

[7] R. Halin, Simplicial decompositions of infinite graphs, in: (B. Bollobás, Ed.) *Advances in Graph Theory* (Annals of Discrete Mathematics **3**), North-Holland Publ. Co., Amsterdam/London 1978.

[8] H. A. Jung, Wurzelbäume und unendliche Wege in Graphen, *Math. Nachr.* **41** (1969), 1–22.

[9] D. König, *Theorie der endlichen und unendlichen Graphen*, Leipzig 1936.

[10] 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.

[11] N. Polat, Spanning trees of infinite graphs, preprint 1989.

[12] N. Seifter, The action of nilpotent groups on infinite graphs, *Mh. Math.* **99** (1985), 323–333.

[13] P. D. Seymour and R. Thomas, An end-faithful counterexample, preprint 1989.

[14] M. E. Watkins, Infinite paths that contain only shortest paths, *J. Combin. Theory B* **41** (1986), 341–355.

Reinhard Diestel, SFB 343, Bielefeld University