The Classification of Finitely Spreading Graphs

Reinhard Diestel

Thomassen introduced the concept of a finitely spreading graph: an infinite graph whose edges can be oriented, each in one, both, or neither direction, so that every vertex has finite out-degree and every ray has a forward oriented tail. He conjectured that a graph is finitely spreading if and only if it is bounded in the sense of Halin—equivalently (see [3]), if it contains none of three specified infinitely spreading graphs.

We prove Thomassen's conjecture in amended form, adding a fourth minimal obstruction to the three conjectured ones.

1. Introduction

There are two basic structural phenomena that make an infinite graph 'truly infinite': vertices of infinite degree, and infinite paths. These features are in a sense dual: one or the other occurs in every infinite connected graph. If only one occurs, i.e. if the graph is rayless or locally finite, its structure tends to be much simpler and easier to describe. This makes it desirable to identify, if possible, in an arbitrary infinite graph a (say) locally finite substructure that captures the essence of certain aspects of the structure of the whole graph. An important such aspect for an infinite graph, as suggested above, lies in its rays—whose 'essence' is their *eventual* route, their behaviour at infinity.

The subject of this paper is a concept recently introduced by Thomassen: call a graph 'finitely spreading' if its edges can be oriented, each in one, both, or neither direction, so that every vertex has finite out-degree and every ray has a forward oriented tail. This concept fits the bill exactly: it determines when the rays in an infinite graph eventually follow a certain locally finite substructure. We shall characterize the finitely spreading graphs by way of excluding four infinitely spreading graphs as topological minors. This proves a conjecture of Thomassen in amended form.

Let us be more precise. We wish to assign to some of the edges xy one or both of their two possible orientations $x \to y$ and $x \leftarrow y$. If the orientation $x \to y$ has been assigned to xy, we say that xy is oriented *from x to y*, or *forward*; note that this does not preclude the possibility that xy has the orientation $x \leftarrow y$ assigned to it as well. (The expression 'forward' here is a reference not to the edge itself but to our present notation of it as xy rather than yx.) The aim, then, is twofold. First, every vertex x should have only finitely many neighbours y such that the edge xy is oriented from x to y. Second, every ray (one-way infinite path) $x_0x_1\dots$ should have a tail (subray) $x_nx_{n+1}\dots$ such that every edge $x_{n+i}x_{n+i+1}$ of this tail is oriented forward; such a tail will be called a *forward oriented tail*. If it is possible to orient the edges of a graph in this way, the graph is called *finitely spreading*, and its orientation a *finitely spreading orientation*; if not, the graph is *infinitely spreading*.

Thus, a finitely spreading orientation provides a graph with a locally finite substructure mapping out the essential directions of its rays: every ray will eventually coincide with a ray indicated by the orientation.

Before we state Thomassen's conjecture formally, let us see some examples. Any locally finite graph is trivially finitely spreading: just orient every edge both ways. Let us remark at this point that the provision for two-way orientations of edges is necessary if we want the locally finite graphs to be included among the finitely spreading ones (as seems clearly desirable): the infinite ladder is an example of a locally finite graph which has no finitely spreading orientation using only one-way orientations of edges.

The regular tree of countably infinite degree, which we denote by T_ω, is a simple example of an infinitely spreading graph. However, a graph need not be 'wide' to be infinitely spreading: a 'long and narrow' example is given by the graph B_0 shown in Fig. 1. (Formally, B_0 is obtained from a ray $x_0x_1\dots$ by adding, for each $n \in \mathbb{N}$, a countably infinite set of independent x_{3n+1}–x_{3n+3} paths, disjoint for different n.)

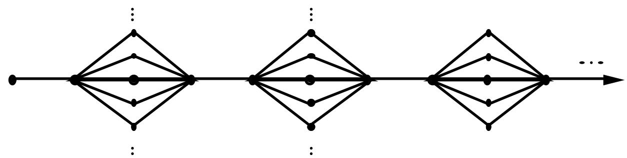

FIGURE 1. The 'bundle graph' B_0

It is clear that B_0 does not admit a finitely spreading orientation: each of the vertices x_{3n+1} would be incident with an edge on its right not oriented away from it, and one would readily find a ray through B_0 using all these edges from left to right, i.e. against their possible orientation.

The graph F_0 of Fig. 2, obtained from a ray $R = x_0x_1\dots$ by adding disjoint rays $Q_2, Q_4, Q_6, \dots$ with $Q_k \cap R = \{x_k\}$, and joining x_{2n+1} to all the new vertices of Q_{2n+2} for every $n \in \mathbb{N}$, is a similar example of an infinitely spreading graph. Interestingly, this graph becomes finitely spreading (and our first non-trivial example of such a graph) if its 'fans' are flipped horizontally, i.e. if the vertices x_{2n+1} are joined to all the vertices of Q_{2n} rather than to those of Q_{2n+2} (add a ray Q_0).

We are now ready to state Thomassen's conjecture. In its original form, the conjecture says that a countable graph is finitely spreading if and only if it is bounded; a graph is called *bounded* if for every integer labelling of its vertices there exists an infinite sequence of natural numbers which eventually exceeds the labelling along any ray in the graph. (A more formal definition

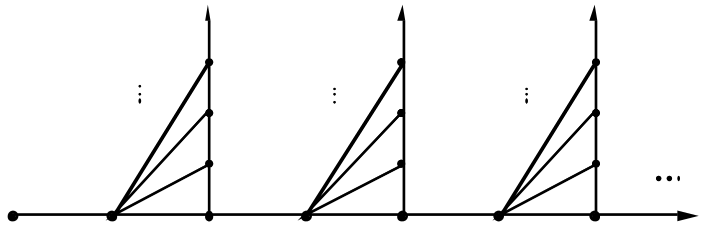

FIGURE 2. The ‘fan graph’ F_0

of boundedness can be found in any of [3,4,5].) It is not difficult to prove that a finitely spreading countable graph must indeed be bounded [5]; the conjecture’s claim is that bounded graphs are finitely spreading.

Confirming a long-standing conjecture of Halin, it was proved in [3] that a countable graph is bounded if and only if it contains no subdivision of T_ω, B_0 or F_0. (It is again easy to see that these graphs are themselves unbounded. An exposition of the main ideas from the proof in [3] of the converse, as well as of related results and problems, has been given in [4].) Thus, the following is an equivalent version of Thomassen’s original conjecture [10]:

Conjecture. (Thomassen 1990)
A countable graph is finitely spreading if and only if it contains no subdivision of T_ω, B_0 or F_0.

It is not difficult to show that any topological subgraph of a finitely spreading graph is again finitely spreading (with the ‘induced’ orientation). Thomassen’s conjecture therefore makes sense, in that the finitely spreading graphs can in principle be characterized by forbidding ‘minimal’ infinitely spreading ones as topological subgraphs. In particular, we see that the ‘only if’ implication of the conjecture is true: since T_ω, B_0 and F_0 are themselves infinitely spreading, they cannot be topological subgraphs of a finitely spreading graph.

Despite the natural appeal especially of its original version (equating finite spreading with boundedness), Thomassen’s conjecture was recently disproved [5]: the graph S_0 shown in Fig. 3 contains no subdivision of T_ω, B_0 or F_0, but is infinitely spreading.

The aim of this paper, then, is to prove that this counterexample to Thomassen’s conjecture is essentially the only one, that the conjecture becomes true if we add S_0 as a fourth forbidden graph:

Theorem 1. *An infinite graph is finitely spreading if and only if it has no subgraph isomorphic to a subdivision of any of the graphs T_ω, B_0, F_0 or S_0.*

Note that, in Theorem 1, we no longer require the graph to be countable. We remark, however, that the equivalence to boundedness is restricted to the

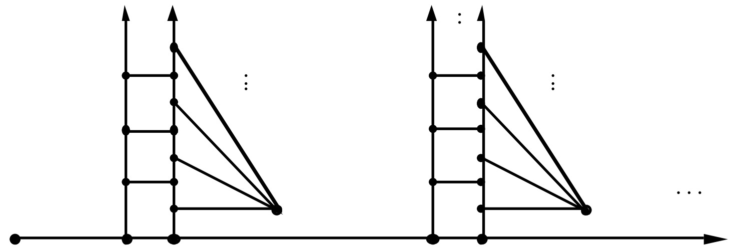

FIGURE 3. The ‘sail graph’ S_0

countable case: there are uncountable finitely spreading graphs, such as the disjoint union of continuum many rays, which are unbounded. Theorem 1, therefore, is somewhat stronger than what remains of Thomassen's original conjecture:

Corollary. *Every infinitely spreading bounded graph contains a subdivision of S_0. Thus, a countable graph not containing a subdivision of S_0 is finitely spreading if and only if it is bounded.*

Our proof of Theorem 1 owes much to techniques developed in [3] for the proof of Halin's bounded graph conjecture. After a short section on basic teminology, these techniques will be introduced and developed further in Sections 3 and 4. Section 5 contains the formal proof of Theorem 1.

2. Terminology

In this section we briefly review some of the basic notation used in this paper. (See [1] for standard definitions not given here.) All the graphs we consider are undirected and simple, i.e. have neither loops nor multiple edges. The axiom of choice is assumed throughout the paper. The set of natural numbers, $\mathbb{N}$, includes 0.

Let G be a graph. The edge set of G will be denoted by $E(G)$, its vertex set by $V(G)$, and its number of vertices by $|G|$. When H is a subgraph of G, denoted by $H \subset G$, we may simply say that G *contains* H. The *neighbourhood* of a vertex is the set of its neighbours. G is *locally finite* if the *degree* of each vertex, the number of its neighbours, is finite. The complete graph on a countably infinite vertex set is denoted by K_ω.

A *path* may be finite or infinite; its *length* is its number of edges. A path of length 0 is *trivial*. A one-way infinite path is a *ray*. Any infinite connected subgraph of a ray R is a *tail* of R. We usually write a path as the sequence of its vertices, which gives these a natural order. It then makes sense to say that a path *starts* or *ends* at some particular vertex, that one vertex *precedes* another on it, and so on. Rays are thought of as having a starting vertex but no last vertex.

The following theorem of König [9], a simpler version of his well-known Infinity Lemma, is a standard tool which we shall use frequently:

König's Theorem. *Every infinite connected graph has a vertex of infinite degree or contains a ray.*

If $P = x_1 \dots x_n$ is a path, we write $\mathring{P}$ for the *interior* of P, the subpath $x_2 \dots x_{n-1}$. Two or more paths are *independent* if their interiors are disjoint. The vertices of $\mathring{P}$ are the *inner* vertices of P. Similarly, if $1 \leqslant i \leqslant j \leqslant n$, we set $Px_i := x_1 \dots x_i$, $P\mathring{x}_i := x_1 \dots x_{i-1}$, $x_i P x_j := x_i \dots x_j$, $x_j P x_i := x_j \dots x_i$, $x_j P := x_j \dots x_n$ and $\mathring{x}_j P := x_{j+1} \dots x_n$ for subpaths of P. Analogous notation will be used for rays and for trees (so that xTy is the unique path from x to y in the tree T), and for the concatenation of paths. For example, if x is a common vertex of two paths P and Q, then PxQ denotes the 'walk' $Px \cup xQ$. The path Px_i defined above is an *initial segment* of P.

For $X, Y \subset G$, we call a finite path $P \subset G$ an X–Y *path* if its endvertices are in X and Y, respectively, and its inner vertices lie in $G - (X \cup Y)$. When X or Y consists of only one vertex, we speak of (say) x–Y or X–y paths rather than $\{x\}$–Y or X–$\{y\}$ paths.

If $X \subseteq G$ or $X \subset V(G)$, we write $G[X]$ for the subgraph of G induced by the vertices in X. For $H \subset G$ and $x \in V(G)$, we write $H[x]$ for the subgraph of H induced by those of its vertices y for which G contains an x–y path that has no vertices in H other than y. In particular, if x is a vertex of H then $H[x]$ is just the singleton $\{x\}$. On the other hand, if $x \notin H$ and C is the component of $G - H$ containing x, then $H[x]$ contains precisely those vertices of H that have a neighbour in C.

When T is a tree, we shall often pick a *root* $r \in V(T)$, and call T *rooted at* r. This induces a natural partial order on $V(T)$, in which $x \leqslant y$ iff x lies on the unique r–y path rTy in T. Informally, we think of this order as expressing the vertices' heights in T; we may then speak of y being *above* x if $y > x$, call it an *upper neighbour* of x if it is above x and adjacent to it, and so on.

Note that, for each vertex $z \in T$, the set of all vertices $x \leqslant z$ is linearly ordered; the subgraph it spans in T will be denoted by $\lfloor z \rfloor$. For any vertex v of T, the subtree of T spanned by all the vertices $x \geqslant v$ will be denoted by $T^{\geqslant v}$. Similarly, we write $T^{>v}$ for $T^{\geqslant v} - v$.

If $T \subset G$ and $V(T) = V(G)$, then T is said to *span* the graph G. The order $\leqslant$ on $V(T)$ thus becomes an order on $V(G)$; if we wish to be more specific, we may denote it by $\leqslant_T$.

A *subdivision* of a graph H is any graph obtained from H by replacing its edges with independent paths of lengths $\geqslant 1$. H is a *topological subgraph* of any graph containing a subdivision of H.

3. Sails, fans, and bundles

In this section, we introduce some specific terminology and techniques needed to handle the proof of Theorem 1. Just for the purpose of the next few definitions, let us call a path P' a *modification* of a non-trivial path $P = x \dots y$ if P' is a path from x to y (not from y to x!). If P is a ray, a ray P' will be called a *modification* of P if P' shares its starting vertex and a tail with P.

Let a and b be distinct vertices. The union of an infinite set of independent a–b paths of lengths $\geqslant 2$ will be called a *bundle from a to b*, or an *a–b bundle*. The interior of any of those paths is a *fibre* of the bundle; note that these fibres are pairwise disjoint. An a–b bundle B is said to be *on* a path P if $B \cap P = \{a, b\}$, a precedes b on P, and aPb has length at least 2. If Q is a (possibly trivial) path from b to some vertex c such that $Q \cap B = \{b\}$, then $B \cup Q$ is an *extended bundle from a to c*. Note that, formally, any ordinary bundle is also an extended bundle (with $b = c$).

Let $B \cup Q$ be an a–b bundle extended to c (as above), let P be a path containing both a and c, and assume (if P is infinite) that P meets only finitely many fibres of B. Then, by the following construction, $B \cup Q \cup P$ contains a bundle B' and a modification P' of P such that B' is on P'. Let B'' be the bundle obtained from B by deleting any fibres that meet P, let F be any a–b path in B'', and put $B' := B'' - \mathring{F}$. Let d be the vertex of P closest to b on the extension $Q = b \dots c$. If $d \in aP$, then B' is a bundle on $P' := PaFbQdP$. (In other words, we replace the segment aPd of P with the bundle B' followed by the path bQd.) Similarly, if $d \in Pa$, then B' is a b–a bundle on $P' := PdQbFaP$.

We shall refer to the above act of replacing P and B with a modification P' of P and a bundle B' on P', where both P' and B' are subgraphs of $P \cup B$, as the *placing* of B on P. To avoid cluttering our notation, we would normally continue to use the terms P and B for their replacements. We have thus shown the following.

(3.1) *An extended bundle from a to c can be placed on any path that contains a and c and meets only finitely many of the bundle's fibres.*

Let Q be a ray starting at a vertex b, and let P_n, $n \in \mathbb{N}$, be paths from a common starting vertex $a \notin Q$ to vertices $q_n \in \mathring{b}Q$. Assume that q_n precedes q_{n+1} on Q for all n. (Informally, we shall think of Q as pointing upwards; thus, q_{n+1} is *above* q_n on Q.) The union F of Q with the paths P_n is called a *fan from a to b*, or an *a–b fan*, if the paths P_n are pairwise disjoint except for a and have no inner vertices on Q. The ray Q in such a fan is called its *spine*; the paths aP_nq_n are its *spokes*.

A *segment* of F is any path of the form $P_nq_nQq_{n-1}$ (if $n > 0$) or P_0q_0Qb. Thus, every vertex $v \in F - \{a, b\}$ is in the interior of a unique segment $S(v)$, and F is the union of its segments $S(q_0), S(q_1), \dots$. Similarly, we let $S(b) := S(q_0) = P_0q_0Qb$. The order of the vertices q_n along Q induces a total order on the segments of F: we write $S(v) <_F S(w)$ if $S(v) = S(q_n)$ and $S(w) = S(q_m)$ with $n < m$. This in turn induces a partial order on the vertices of $F - a$, where $v <_F w$ if $S(v) <_F S(w)$, and v, w are incomparable (denoted as $v \sim_F w$) if $S(v) = S(w)$. For "$v <_F w$ or $v \sim_F w$" we write $v \leqslant_F w$.

An a–b fan F is said to be *on* a path P if $F \cap P = \{ a, b \}$ and a precedes b on P. *Placing* a fan is defined analogously to placing a bundle. However, it is not always possible to place an a–b fan F on a path P containing a and b, even if P meets only a finite part of F. For example, if P meets F in exactly a and b but b precedes a on P, then $F \cup P$ does not contain a fan F' and a modification P' of P such that F' is on P'. However:

(3.2) *An a–b fan F can be placed on any path P such that $P \cap F$ is finite, $a, b \in P$, and Pa avoids the spine of F.*

Indeed, let c be the vertex of P highest on the spine Q of F (so that $P \cap \mathring{c}Q = \emptyset$), and let F' be the a–c fan obtained from F by deleting $Q\mathring{c}$ and any spoke's interior that meets P or attaches to Qc. F' is an a–c fan on P (unmodified).

Let a, b, c be distinct vertices. A *sail from c to b*, or a *c–b sail*, is any union S of a ray Q_1 starting at c, an a–b fan F (with spine Q_2, say) such that $F \cap Q_1 = \emptyset$, and infinitely many disjoint Q_1–Q_2 paths avoiding $F - Q_2$. The ray Q_1 is called the *mast* of this sail; its *spine* and *spokes* are those of F. S is said to be *on* a path P if $S \cap P = \{ c, b \}$ and c precedes b on P.

We say that a sail S *can be placed* on P if $S \cup P$ contains a modification P' of P and either a sail *or a fan* on P'. Placing sails is more complicated than placing bundles or fans; the following sufficient conditions will be enough for our purposes.

Lemma 3.3. *Let S be a sail (with a, b, c, Q_1, Q_2, F as above), and let P be a path such that $P \cap S$ is finite and $a \in P$. Assume further that $aP \cap Q_2 \neq \emptyset$, and let q_2^+ be the vertex of aP highest on Q_2. S can be placed on P if at least one of the following conditions is satisfied:*

(i) *$Pa \cap q_2^+ Q_2 = \emptyset$;*

(ii) *condition (i) fails and $Pq_2^- \cap Q_1 \neq \emptyset$, where q_2^- is the first vertex of Pa on $q_2^+ Q_2$;*

(iii) *$aP \cap Q_1 \neq \emptyset$.*

Proof. If (i) holds, then F contains an a–q_2^+ fan which can be placed on P, by (3.2). Suppose now that (i) fails, and let q_2^- be defined as in (ii). By assumption, (ii) or (iii) holds, so $Q_1 \cap (Pq_2^- \cup aP) \neq \emptyset$; let q_1 be the highest vertex on Q_1 in this set.

If $q_1 \in Pq_2^-$, choose a $(q_1 Q_1)$–$(q_2^+ Q_2)$ path $U = u_1 \ldots u_2$ in $S - P$. Then

$$P' := P q_1 Q_1 u_1 U u_2 Q_2 q_2^+ P$$

is a modification of P, and S contains a u_1–u_2 sail placed on P': note that $P\mathring{q}_1$ avoids q_1Q_1 by definition of q_1, and it avoids $q_2^+Q_2$ by definition of q_2^- and the fact that $q_1 \in Pq_2^-$; q_2^+P avoids q_1Q_1 and $q_2^+Q_2$ by definition of q_1 and q_2^+, respectively; both Pq_1 and $\mathring{q}_2^+P$ avoid a, because $a \in \mathring{q}_1P\mathring{q}_2^+$.

If $q_1 \notin Pq_2^-$, then $q_1 \in aP$. Let U be a q_2^-–a path consisting of an initial segment of $q_2^-Q_2$ followed by a spoke avoiding $P - a$. Let $F' \subseteq S$ be an a–q_1 fan with spine q_1Q_1 such that $F' \cap U = \{ a \}$ and no spokes of F' meet $P - a$. By definition of q_1, the path Pq_2^-, and hence Pq_2^-U, avoids the spine of this fan. Thus if W is any a–q_1 path in F', then by (3.2) F' can be placed on the path $Pq_2^-UaWq_1P$, which is a modification of P. □

Let $\mathcal{H} = \{ H_1, H_2, \ldots \}$ be a finite or infinite, and possibly empty, set of disjoint sails, fans and/or bundles, all on a fixed path P. Assume that the H_i are arranged in such a way that the first vertex of P lies in none of them, and both vertices of $H_i \cap P$ precede both vertices of $H_{i+1} \cap P$ on P for every i. If $\mathcal{H}$ is infinite, then $P \cup \bigcup \mathcal{H}$ is called a *sail-fan-bundle graph*, or *SFB-graph*, on P; if $\mathcal{H}$ is finite, $P \cup \bigcup \mathcal{H}$ is called a *partial SFB-graph* on P.

Note that every SFB-graph contains a subdivision of one of the graphs B_0, F_0 and S_0 from Theorem 1. To prove the non-trivial direction of the theorem, it will thus be sufficient to find an SFB-graph or a subdivision of T_ω in every infinitely spreading graph.

We conclude this section with one of the basic lemmas from [3]. The construction used in its proof will be referred to later, so we include the proof for the reader's convenience.

Lemma 3.4. *Let a and b be distinct vertices, and let $\mathcal{P}$ be an infinite set of a–b paths whose second vertices are pairwise distinct. The union $\bigcup \mathcal{P}$ of these paths then contains either an extended bundle from a to b, or an a–b fan all whose segments are initial segments of paths in $\mathcal{P}$.*

Proof. Let P_n, $n \in \mathbb{N}$, be distinct paths from $\mathcal{P}$ such that the second vertex of P_n is not on P_k, for all $k < n$. Choose initial segments S_n of the P_n, as follows. Let $S_0 := P_0$. Having defined $S_0, \ldots, S_n$ for some $n \in \mathbb{N}$, let S_{n+1} be the initial segment of P_{n+1} which ends at the first vertex of P_{n+1} in $S_0 \cup \ldots \cup S_n$. Let this vertex be called w_{n+1}, and let $k(n+1)$ be the minimal k for which $w_{n+1} \in S_k$. Call $S_{k(n+1)}$ the *predecessor* of S_{n+1}.

Let K be the tree with vertex set $\mathbb{N}$ and edge set $\{ nk(n) \mid n > 0 \}$. By König's theorem, K contains a ray or has a vertex of infinite degree. If K has a vertex of infinite degree, k say, then S_k has a vertex w such that $w_n = w$ for infinitely many $n > k$. The corresponding paths S_n are pairwise independent, and they form an a–w bundle which extends to b along P_k (delete any fibres hit by w_kP_k; Fig. 4).

On the other hand if K contains a ray N, without loss of generality starting at 0, then the paths S_n with $n \in N$ are easily seen to form an a–b fan (Fig. 4;

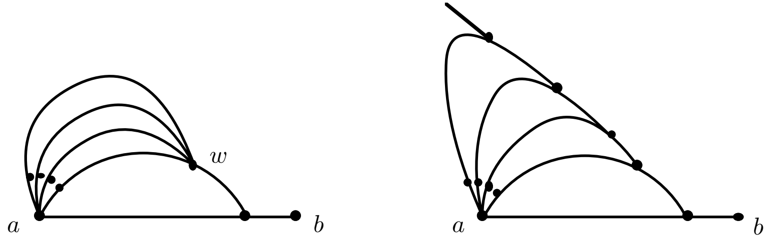

FIGURE 4. An extended bundle or a fan from a to b

note that, by the construction of K, if m succeeds n on N then $m > n$, so S_n is the predecessor of S_m). □

4. Normal spanning trees

In this section we introduce the main structural tool for our proof of Theorem 1, the concept of a normal spanning tree. A rooted spanning tree T of a graph G is called *normal* if the endvertices of every edge of G are comparable in the tree order $\leqslant_T$ which T induces on $V(G)$. Thus, if r is the root of T and $xy \in E(G)$, then either x lies on the r–y path in T or y lies on the r–x path in T. Intuitively, all the edges of G are 'along' branches of T, never 'across'.

It is easy to see that all finite connected graphs have normal spanning trees (obtained, for example, by a 'depth-first search'). Jung [8] characterized the infinite graphs having normal spanning trees; his characterization implies that every countable connected graph contains such a tree. For more on normal trees, especially for uncountable graphs, see [6].

Starting from Jung's result and using the theory of simplicial decompositions of graphs (see [2]), Halin [7] was able to prove the following:

Theorem 4.1. (Halin)
If G is connected and contains no subdivision of K_ω, then G has a normal spanning tree.

When we prove later that an infinitely spreading connected graph G must contain one of the four topological subgraphs listed in Theorem 1, we may thus assume that G has a normal spanning tree: if not, it contains a subdivision of K_ω, and hence all four of the desired types of subgraph.

We may already note one interesting consequence of Theorem 4.1 which is somewhat less immediate from first principles. If G is finitely spreading, it contains no subdivision of K_ω, and so each of its components has a normal spanning tree. Then all the edges of G are between vertices that are comparable in their respective tree order, and each vertex is adjacent to at most finitely many vertices below it. By orienting every edge in G downwards (in addition to any previous orientation), we may therefore turn any finitely spreading orientation of G into one in which every edge is oriented in at least one direction. The provision in the definition of finite spreading for edges to remain undirected thus turns out to be redundant.

We now collect together a number of simple facts about normal spanning trees to be used later. All these are easy to prove, and they should help the reader develop an intuition for the properties of a normal spanning tree. (Explicit proofs can be found in [3; §3].)

Let G be a fixed graph, and assume that G contains a normal spanning tree T with root r. Any reference to an order on the vertices of G (such as 'above', 'below', down-closures $\lfloor x \rfloor$ and so on) will be assumed to refer to the order $\leqslant_T$ induced by T. A ray in G will be called a *normal ray* if it starts at r and is contained in T. The union of all normal rays, clearly a subtree of T, will be denoted by T'.

Our first lemma translates the local defining property for T (that every edge of G runs vertically along T) into a more global separation property of G:

Lemma 4.2. *Let $x, y \in V(G)$. Then $\lfloor x \rfloor \cap \lfloor y \rfloor = \lfloor \inf\{x, y\} \rfloor$ separates x from y in G, i.e. every x–y path in G contains a vertex of $\lfloor x \rfloor \cap \lfloor y \rfloor$.*

Lemma 4.2 implies that, in terms of vertex sets, separators of the form $\lfloor x \rfloor$ leave the same components in G as in T:

(4.3) *If $x \in V(G)$ and C is an induced subgraph of G, then C is a component of $G - \lfloor x \rfloor$ if and only if $C \cap T$ is a component of $T - \lfloor x \rfloor$.*

Note that Lemma 4.2 does *not* imply that the *interior* of any x–y path P meets $\lfloor x \rfloor \cap \lfloor y \rfloor$. This is true only if x and y are incomparable; if $x < y$, then P may well have a non-empty interior lying somewhere above x.

If $x \in T'$ then $\lfloor x \rfloor \subset T'$. Lemma 4.2 thus has another immediate consequence:

(4.4) *$T'[x] \subset \lfloor x \rfloor$ for every $x \in V(G)$. Similarly, if R is a normal ray, then $R[x] \subset \lfloor x \rfloor$.*

(4.4) implies that, if x is a vertex and R is a normal ray in G, the vertex set of $R[x]$ is a finite chain with respect to $\leqslant_T$. Let us define the *R-height* of x to be the unique maximal vertex of $R[x]$. Thus, the R-height of x is the first vertex of R on the descending path xTr. Note that the R-height of a vertex $h \in R$ is h itself; the other vertices of R-height h are precisely those vertices which lie above h but not above any vertex of $\mathring{h}R$. In particular, $\lfloor h \rfloor$ separates (in T and in G; cf. (4.3)) the vertices of R-height h from any vertices of different R-height:

(4.5) *Let $x, y \in V(G)$, and let R be a normal ray.*

(i) *If x and y are in a common component of $G - R$, then their R-heights coincide.*

(ii) *If $y > x$, then the R-height of y is at least that of x.*

If $H \subset G$ or $H \subset V(G)$, and R is a normal ray, we shall say that the R-height of (the vertices in) H is *bounded* if R has a vertex h above the R-heights of all the vertices in H. Another subgraph $H' \subset G$ is *H-clear* with respect to R if the R-height of each vertex of H' is strictly greater than the R-heights of all the vertices in H. The R-height of H *tends to infinity* if H is infinite and for each $h \in R$ only finitely many vertices of H have R-height $\leqslant h$.

The concept of R-height will help us to organize our construction of an SFB-graph in Section 5. For example, if H is a partial SFB-graph already constructed at some point (on the path P, say) and H' is H-clear, we shall aim to extend P into H' and continue our construction there; we can then be sure that sails, fans or bundles constructed at later stages will not interfere with earlier ones.

If R is any ray in G and R' is a normal ray, let us say that R *follows* R' if $|R \cap R'| = \infty$.

(4.6) *Every ray in G follows a unique normal ray.*

If R follows R' then, in terms of R'-height, the vertices of R behave like those of R' itself:

(4.7) *Let $R \subset G$ be a ray, and let R' be the normal ray it follows. Then the R'-height of R tends to infinity.*

Our standard application of (4.7) for the construction of an SFB-graph will be that if $H \subseteq G$ has bounded R'-height and R follows R', then R has an H-clear tail.

And a last lemma from [3; §3]:

Lemma 4.8. *Let R be a normal ray. If G has infinitely many vertices with neighbourhoods of unbounded R-height, then G contains a subdivision of K_ω.*

5. Proof of the classification theorem

The following special case of Theorem 1 was first observed by Thomassen [10]; for completeness, we sketch a proof.

Lemma 5.1. *A tree is finitely spreading if and only if it contains no subdivision of T_ω.*

Proof. Let G be a tree not containing a subdivision of T_ω; we show that G is finitely spreading. Pick a root r in G. Proceeding by ordinal recursion, we shall orient some of the edges upwards and delete some of the vertices. For each ordinal α, let G_α denote the subgraph of G induced by those vertices that have not been deleted by time α. (If G_α is non-empty, it will be a subtree of G rooted at r.) Find a vertex $v \in G_\alpha$ such that $G_\alpha^{\geqslant v}$ is locally finite; such a vertex v exists, as otherwise G_α contains a subdivision of T_ω. Orient every edge in $G_\alpha^{\geqslant v}$ upwards, and delete $G_\alpha^{\geqslant v}$. This completes the recursion step. The recursion ends only when all the vertices have been deleted; since at least one vertex is deleted at each step, this happens after no more than $|G|^+$ steps.

For each vertex x of G, the only time an edge incident with x gets oriented is when x and this edge belong to the same locally finite tree $G_\alpha^{\geqslant v}$. Since these trees are vertex disjoint, x has finite out-degree.

Now consider any ray in G. Let R be a tail of this ray going up in G. Let α be minimal such that some vertex x of R was deleted at time α. Then the entire tail xR of R was deleted at that time, and hence all its edges were oriented upwards. $\square$

As we observed in the Introduction, no finitely spreading graph can contain a subdivision of B_0, F_0, S_0 or T_ω. We now prove the other direction of Theorem 1, that every infinite graph G not containing any of these four types of subgraph is finitely spreading. Obviously, we may assume that G is connected. By Halin's theorem (4.1) and our assumption that G contains no subdivision of T_ω (say), G has a normal spanning tree T. As before, we denote the union of all normal rays by T', and denote the order $\leqslant_T$ on the vertices of G simply by $\leqslant$. (Again, any unspecified reference to an order on vertices will refer to this order.)

As T' contains no subdivision of T_ω, Lemma 5.1 implies that T' has a finitely spreading orientation. Let us fix such an orientation, and consider the edges of T' that are oriented upwards as *labelled*. Every vertex of T' is thus incident with finitely many labelled edges, and every normal ray has a tail whose edges are all labelled. We shall use this labelling as guidance when we come to orient the edges of G; this does not mean, however, that that orientation will necessarily be an extension of the above orientation of T'.

To define our desired orientation of the edges of G, we need some more definitions.

Let us call a pair (v, w) of vertices of T' *good* if G contains only finitely many v–w paths whose interiors avoid T'. For each good pair (v, w) let G_{vw} be the union of all those v–w paths; this is a finite subgraph of G. Note that if vw is an edge of T', then G_{vw} contains this edge. Let G^+ be the union of the graphs G_{vw}, taken over all good pairs (v, w). Then

(5.2) *any vertex of infinite degree in G^+ lies on T'.*

Indeed, recall that any subgraph of the form $T'\lceil x \rceil$ is finite (4.4). As $v, w \in T'\lceil x \rceil$ whenever $x \in G_{vw} - T'$, any $x \in G^+ - T'$ is in only finitely many graphs G_{vw}, and thus has finite degree in G^+.

It was shown in [3; pp. 147–149]* that if G has a ray without a tail in G^+, then it contains a subdivision of B_0. Hence,

(5.3) *every ray in G has a tail in G^+.*

It will thus suffice to orient the edges of G^+.

Recall that, in terms of vertex sets, the deletion of a subgraph of the form $\lfloor x \rfloor$ from G or from T leaves exactly the same components (4.3). For each $a \in T'$, let $\mathcal{C}_a$ be the set of those components of $G - \lfloor a \rfloor$ above a whose unique minimal vertex is in T' and joined to a by a labelled edge.

Now consider a finite partition $\mathcal{D}$ of the graph $G - \lfloor a \rfloor$ into unions of its components, with $\emptyset \in \mathcal{D}$. (Thus, every component of $G - \lfloor a \rfloor$ is either itself an element of $\mathcal{D}$ or else a component of an element of $\mathcal{D}$.) Note that a is uniquely determined by $\mathcal{D}$, as the maximal vertex of $G - \bigcup \mathcal{D}$. A $\mathcal{D}$-*trace* is any $(2k)$-tuple (for some integer $k > 0$) of the form $(v_1, D_1, \ldots, v_k, D_k)$, where the v_i are distinct vertices in $\lfloor a \rfloor$ and the D_i are elements of $\mathcal{D}$. Note that for any fixed $\mathcal{D}$ there are only finitely many $\mathcal{D}$-traces.

Since $\lfloor a \rfloor$ separates the elements of $\mathcal{D}$ pairwise in G, any finite or infinite path $P \subseteq G$ starting in $\lfloor a \rfloor$ leaves a unique $\mathcal{D}$-trace in a natural way: for each $i \leqslant k := |P \cap \lfloor a \rfloor|$, the vertex v_i in this trace is the i'th vertex of P in $\lfloor a \rfloor$, while D_i is either empty (if $v_i v_{i+1}$ is an edge of P or $i = k$ and P ends in v_k) or else is the unique element of $\mathcal{D}$ containing $\mathring{v}_i P \mathring{v}_{i+1}$ (or $\mathring{v}_i P$, in the case of $i = k$).

The idea behind this definition is that such traces may help us in the construction of an SFB-graph along a ray R, in the following way. If F is a fan (say) not on R itself but on some other path P, we may still hope to use F for our SFB-graph by substituting a segment of P carrying F for a suitable segment of R. If P has the same $\mathcal{D}$-trace τ as R for some $\mathcal{D}$ and F is contained in a single $D \in \mathcal{D}$, we may replace all the segments of R through D with the corresponding segments of P. Unless D is the last element of $\mathcal{D}$ in τ, the resulting modification of R will then again be a ray, and F will be properly placed on it in the same way as it was placed on P.

---

* This is non-trivial. The relevant part of [3], however, uses the same terminology as here and can be copied almost verbatim.

We are now ready to orient the edges of G. We proceed in three steps; observe that, in each step, every vertex has only finitely many edges oriented away from it, so all out-degrees remain finite.

Step 1. *Orient every edge downwards.* (Recall that adjacent vertices of G are comparable in the tree order of T, so one is below the other. Each vertex has only finitely many vertices below it.)

Step 2. *Orient any edge $xy \in G^+$ from x towards y if $x \notin T'$.* (Recall that such vertices x have finite degree in G^+, by (5.2).)

Step 3. *Consider every vertex $a \in T'$ in turn. For each a, let*

$$\mathcal{D} := \mathcal{C}_a \cup \{ D^+, D^-, \emptyset \},$$

where D^+ is the union of all components of $G - \lfloor a \rfloor$ above a that are not in $\mathcal{C}_a$, and D^- is the union of all other components of $G - \lfloor a \rfloor$ (i.e. of all components whose minimal vertex is not an upper neighbour of a in T). Consider every $\mathcal{D}$-trace $\tau = (v_1, D_1, \dots, v_k, D_k)$ such that $a = v_i$ for some $i \in \{ 1, \dots, k \}$. If a has only finitely many neighbours x in D_i such that a, x are consecutive vertices (in this order) on some path in G with $\mathcal{D}$-trace τ, then orient the edges ax from a towards x for all these x.

It remains to prove that, with the above orientation of the edges of G, every ray in G has a forward oriented tail. So let R be an arbitrary ray in G, and let R' be the normal ray it follows (4.6). *Whenever we use the term '$(\cdot)$-clear' in the rest of this paper, it will be with respect to R'.*

Replacing R with a suitable tail of R if necessary, we may assume the following four statements as true.

(5.4) $R \subset G^+$.

(Recall that, by (5.3), R has a tail in G^+.)

(5.5) *The neighbourhood (in G) of each vertex on R has bounded R'-height.*

(By Lemma 4.8, only finitely many vertices of G have neighbourhoods of unbounded R'-height, since G contains no subdivision of K_ω.)

(5.6) *The starting vertex r_0 of R is on R', and minimal in $R \cap R'$.*

(Recall that $R \cap R'$ is infinite by definition of R'.) Note that, since R' is normal and hence $\lfloor r_0 \rfloor \subseteq R'$, Lemma 4.2 and (5.6) imply that r_0 is in fact below *all* the other vertices of R.

(5.7) *All the edges of $r_0 R'$ are labelled.*

(Since the labelling of T' comes from a finitely spreading orientation, R' has only finitely many unlabelled edges. If (5.7) requires replacing R with a tail, this may undo any previous adjustment to (5.6). However, this is easily mended by another replacement according to (5.6).)

Suppose that R has infinitely many edges not oriented forward. We shall inductively construct a sail-fan-bundle graph along R, adding one sail, fan or bundle at a time.

More formally, let us define a sequence $(G_i)_{i\in\mathbb{N}}$ of partial SFB-graphs on paths R_i, with the following properties:

(5.8) (i) G_i is a partial SFB-graph on a path R_i, with sails, fans or bundles $H_1, \dots, H_i$ in this order (the H_j with $j < i$ being the same as for G_j);

(ii) R_i starts at r_0 and ends at a vertex $s_i \in R \cap R'$;

(iii) R_i has a vertex $r_i \neq s_i$ such that all the sails, fans or bundles H_j of G_i are on $R_i r_i$;

(iv) $r_i R_i \subset R$, and r_i precedes s_i on R (so $r_i R_i = r_i R_i s_i = r_i R s_i$);

(v) $s_i R$ is G_i'-clear, where $G_i' := G_i - \mathring{r}_i R_i$;

(vi) if $i > 0$, then $R_i r_i \supseteq R_{i-1} r_{i-1}$ and $G_i' \supseteq G_{i-1}'$.

(See Figure 5.)

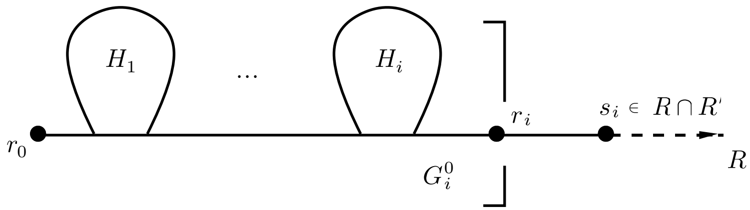

FIGURE 5. The partial SFB-graph G_i

Condition (v) means that every vertex on R from s_i onwards has R'-height strictly greater than the R'-height of any vertex in G_i up to r_i. So in particular, the R'-height of G_i', and hence that of $G_i = G_i' \cup r_i R s_i$, is bounded:

(5.9) *G_i has bounded R'-height.*

Condition (vi) ensures that the essential parts G_i' of G_i (which contain all the sail, fans or bundles), and their principal paths $R_i r_i$, are nested. When we have completed the induction step, it will therefore be clear that the union of all the G_i' is an SFB-graph, and the proof of Theorem 1 will be complete.

It is clear that the induction starts: just let s_0 be any vertex in $R \cap R'$ such that $s_0 R$ is $\{ r_0 \}$-clear (cf. (4.7)), and put $G_0 := R_0 := R s_0$. Turning to the induction step, let us assume that graphs $G_0, \dots, G_i$ have been defined in accordance with (i)–(vi). Our aim is to find a new sail, fan or bundle of bounded R'-height, which is disjoint from G'_i and can be placed on $\mathring{r}_i R$.

To this end, choose a vertex a on R, far enough ahead that aR is $(G_i \cup Rs_i)$-clear, and so that the first edge ab of aR is not oriented forward. As both the first and the second orientation step failed to orient ab from a to b, we have $a < b$ and $a \in T'$ (5.4). Let c be the unique minimal vertex of the component of $T^{>a}$ that contains b. Thus, c is an upper neighbour of a in T.

By the remark following (5.6) we have $r_0 \in \lfloor a \rfloor$, so R has a well-defined $\mathcal{D}$-trace τ (where $\mathcal{D}$ is defined as in Step 3 above). Since ab did not get oriented forward in Step 3, a must have infinitely many neighbours x each following a on some (finite) path $P(x)$ with $\mathcal{D}$-trace τ; let X be the set of these neighbours. (We shall later replace X with an infinite subset of itself but continue to call this subset X. Any properties of X established in the meantime and needed later will be invariant under this replacement.) By (5.5),

(5.10) *X has bounded R'-height.*

Let C denote the element of $\mathcal{D}$ following a in τ. Thus, $X \subseteq V(C)$, and either $C \in \mathcal{C}_a$ (if the edge ac is labelled) or $C = D^+$ (if ac is unlabelled). As C lies above a, the R'-height of any vertex in C is at least the R'-height of a (4.5). Our assumption that aR is G_i-clear thus implies that C is G_i-clear, and hence that

(5.11) $C \cap G_i = \emptyset$.

Let us deal first with the case that ac is an edge of R'. Then $a \in \mathring{s}_i R'$, by the choice of a. By (5.7), ac is labelled. Thus $C \in \mathcal{C}_a$, so C is connected, $C \cap T$ is a tree (4.3), and c is its minimal element. By Lemma 3.4, the graph $\bigcup_{x \in X} ax(C \cap T)b$ contains a fan or extended bundle H from a to b. By (5.10), H has bounded R'-height, so $R \cap H$ is finite (4.7). We may thus try to place H on $\mathring{r}_i R$.

By assumption (v) in the induction hypothesis, $\{ s_i \}$, and hence $s_i R'$, is G'_i-clear. In particular, $G'_i \cap s_i R' a = \emptyset$. Let r be the first vertex of $r_i R_i$ in $s_i R' a$; r exists, since s_i is a candidate (Fig. 6).

As $c \in R'$, we have $cR' \subset C$. Since R follows R' and meets $\lfloor a \rfloor$ only finitely often, R has a tail $\hat{R}$ in C (cf. Lemma 4.2). If $\hat{R} \cap H = \emptyset$, let r_{i+1} be the unique vertex of $\hat{R}$ on some b–$\hat{R}$ path $P \subseteq C$. If $\hat{R} \cap H \neq \emptyset$, let r_{i+1} be the last vertex of $\hat{R}$ in H, and let P be any b–r_{i+1} path in H. In both cases let s_{i+1} be chosen from $R' \cap \mathring{r}_{i+1} R$ so that $s_{i+1} R$ is $(G_i \cup H \cup P)$-clear, and define

$$R_{i+1} := R_i r R' a b P r_{i+1} R s_{i+1} \, .$$

Note that R_{i+1} is indeed a path: since $R_i \cap C \subset G_i \cap C = \emptyset$ (5.11), we have $R_{i+1} a \cap C = \emptyset$, while $b R_{i+1} \subseteq C$.

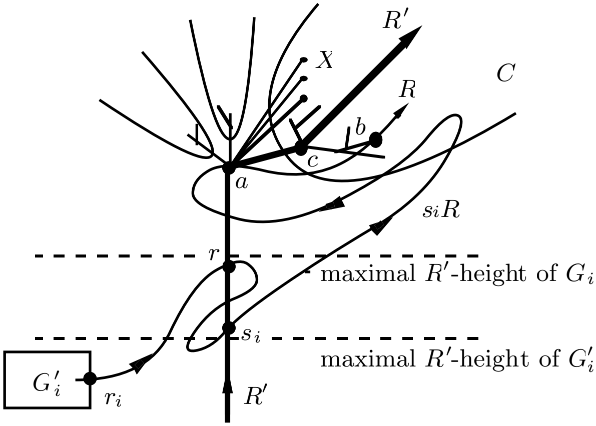

FIGURE 6. Constructing G_{i+1} when $ac \in R'$

Since $R_{i+1}\mathring{a} \cap H = \emptyset$, we may place H on $\mathring{r}_i R_{i+1} r_{i+1}$, no matter whether H is a bundle or a fan. Let H_{i+1} be the resulting bundle or fan. By the definition of placing, H_{i+1} and the new $\mathring{r}_i R_{i+1} r_{i+1}$ will be contained in the union of H and the old $\mathring{r}_i R_{i+1} r_{i+1}$. Since this union meets $r_{i+1}R$ only in r_{i+1}, the placing of H will leave $r_{i+1}R_{i+1}$ unaltered and hence satisfy (iv). Similarly, the choice of s_{i+1} ensures that $s_{i+1}R$ will be G'_{i+1}-clear, as required by condition (v).

Verification of the other conditions in the induction hypothesis for $i+1$ is straightforward.

Having dealt with the case when ac is an edge of R', we shall from now on assume that ac is not an edge of R'. Then $R' \cap C = \emptyset$, even if $C \in \mathcal{C}_a$. (We already know from (5.7) that R' cannot meet C if $C = D^+$.) In particular:

(5.12) *C has constant R'-height, namely that of a.*

The fact that $R' \cap C = \emptyset$ further implies that the last entry in τ cannot be equal to C: since R follows R', it cannot have a tail in C. Thus, R must return to $\lfloor a \rfloor$ after b, say first at the vertex d. Then d is the vertex following a in τ, and so each of the paths $P(x)$ hits d after a, with $\mathring{a}P(x)\mathring{d} \subseteq C$.

Let C^- be the subgraph of G induced by C together with those vertices $v \in R \cap \lfloor a \rfloor$ for which R contains an edge between v and C. These are precisely the vertices occuring in τ either before or after an occurrence of C (such as a and d). Let us prove the following.

(5.13) *To complete the induction step for (5.8), it suffices to find a sail, fan or extended bundle H in C^- that can be placed on some $P(x)$, $x \in X$.*

For a proof of (5.13), observe first that all the vertices of $C^- \cap \lfloor a \rfloor$ are on s_iR. Indeed, $Rs_i \cap C = \emptyset$, by (5.12) and the choice of a. As the vertices of $C^- \cap \lfloor a \rfloor$ are on R and adjacent (in R) to a vertex in $R \cap C$, they must be on s_iR.

Let H and x be given as in (5.13). As $H \subseteq C^-$, the above considerations and (5.11) imply that $H \cap G_i' = \emptyset$. Using the fact that $P(x)$ has the same $\mathcal{D}$-trace τ as R, we may obtain a ray $\widetilde{R}$ from R by replacing its segments through C with the corresponding segments of $P(x)$. As all the above segments (including their endvertices) lie in C^-, this will only affect $s_i R$. Moreover, the replacement will only affect a finite part of $s_i R$: since the last entry in τ cannot be equal to C, the tail R^+ of R following its last vertex below a is also a tail of $\widetilde{R}$. As H is contained in C^-, it can be placed on $\widetilde{R}$ in the same way as it could be placed on $P(x)$; let H_{i+1} be the resulting sail, fan or bundle on $\widetilde{R}$. Again, the placing will affect neither Rs_i nor R^+.

To complete the induction step, we pick $r_{i+1} \in R^+$, choose $s_{i+1} \in R' \cap \mathring{r}_{i+1} R$ so that $s_{i+1} R$ is $(G_i \cup Rr_{i+1})$-clear, set $R_{i+1} := R_i s_i \widetilde{R} s_{i+1}$, and define G_{i+1} accordingly. Since the R'-height of $H \cup (\widetilde{R} - R)$ is bounded by that of a (5.12), our choice of s_{i+1} will satisfy (5.8)(v); verification of the other conditions in (5.8) is again straightforward. This completes the proof of (5.13).

By Lemma 3.4, the subgraph $\bigcup_{x \in X} aP(x)d$ of C^- contains a fan or extended bundle H from a to d. If H is an extended bundle, it can be placed on any $P(x)$ and we are done by (5.13). We may therefore assume that H is an a–d fan; let us call this fan F, and its spine Q. By Lemma 3.4, we may assume that every segment of F is an initial segment of some path $aP(x)d$, whose second vertex is x. The neighbours of a in F are therefore all in X; let us replace X with the set of these neighbours. X is thus totally ordered by $<_F$, and $\{ S(x) \mid x \in X \}$ is the set of the segments of F.

For each $x \in X$, F contains a subfan $F(x)$ on $S(x) \subseteq P(x)$. By (5.13), however, we may assume that $F(x)$ cannot be placed on $P(x)$. By (3.2) this means that, for each x, the path $P(x)a$ hits the spine of $F(x)$, which is the part of Q above $S(x)$ in F. For each x, let $v(x)$ be the first vertex of $P(x)a$ strictly above $S(x)$ in F (though not necessarily on Q), and let $u(x)$ be the last vertex of $P(x)v(x)$ below a (Fig. 7). Note that $\mathring{u}(x)P(x)v(x) \subseteq C$; thus, $u(x) \in C^-$.

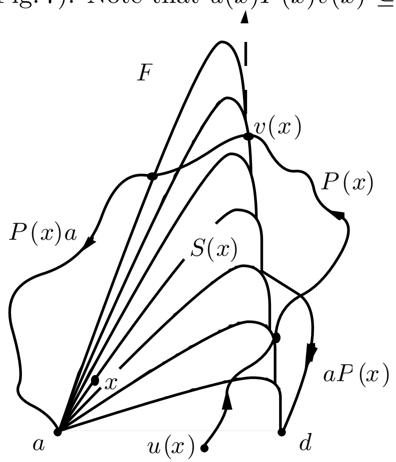

FIGURE 7. The fan F

Our aim is to use the paths $u(x)P(x)v(x)$ to extend F to a sail, in a similar way as we constructed F from the paths $aP(x)d$. We proceed in three steps. The first two steps are as in the construction of a fan (see the proof of Lemma 3.4): we shall define a graph K on a set of vertices corresponding to internally disjoint segments of the paths $u(x)P(x)v(x)$, and then use König's theorem to select a sequence of these segments that defines a ray in K. In the third step, we shall try to arrange the segments from this sequence to form, together with F, a $u(x)$–d sail on $P(x)$ for some x.

For the first step, we shall select an infinite increasing sequence $x_1 <_F x_2 <_F \ldots$ in X, choose for each n a path $Q(x_n) \subseteq P(x_n)$ starting at $v(x_n)$ and ending at a vertex $w(x_n) \in u(x_n)P(x_n)\mathring{v}(x_n)$ (thus, $Q(x_n)$ is a subpath of $P(x_n)$ in reverse direction), and define a graph K with vertex set $\mathbb{N}$ and coloured edges, so that the following conditions hold for all $n > 0$:

(5.14) (i) *If $n > 1$, then $x_n \sim_F v(x_{n-1})$.*

(ii) $\mathring{Q}(x_n) \cap \Big(\lfloor a \rfloor \cup F \cup Q(x_1) \cup \ldots \cup Q(x_{n-1}) \Big) = \emptyset$.

(iii) *Exactly one of the following three statements holds:*

(a) *$w(x_n) \in (C^- \cap \lfloor a \rfloor) - d$;*

(b) *$w(x_n) \in F - a$, with $x_m \leqslant_F w(x_n) <_F x_n$ for some $m < n$;*

(c) *there exists an $m < n$ with $w(x_n) \in \mathring{Q}(x_m)$* (by (ii), this m will be unique).

(iv) *In K, the vertex n has exactly one neighbour in $\{0, \ldots, n-1\}$.*

If (a) of (iii) holds, then this neighbour is 0 and the edge $n0$ is coloured blue.

If (b) of (iii) holds, then this neighbour is the maximal $m < n$ with $x_m \leqslant_F w(x_n)$, and the edge nm is coloured green.

If (c) of (iii) holds, then this neighbour is m as defined there, and the edge nm is coloured red.

(v) *If nm is a green edge of K with $n > m$, then $w(x_n) <_F v(x_m)$.*

Let us now choose such a sequence $x_1, x_2, \ldots$ and paths $Q(x_n)$. Let x_1 be the $<_F$-minimal element of X, let $Q(x_1)$ be $u(x_1)P(x_1)v(x_1)$ in reverse direction (i.e. from $v(x_1)$ to $u(x_1)$), and set $w(x_1) = u(x_1)$. For K, join 1 to 0 by a blue edge. Conditions (i)–(v) are then satisfied for $n = 1$. (Note that $\mathring{Q}(x_1)$ cannot hit F, by the choice of x_1 and definition of $v(x_1)$.)

Assume now that $n > 1$, and that $x_1, \ldots, x_{n-1}$, with the corresponding paths $Q(x_i)$ and edges of $K[0, \ldots, n-1]$, have been chosen according to (i)–(v). Let x_n be the unique vertex in X with $v(x_{n-1}) \in \mathring{S}(x_n)$. Note that (i) holds for n. Let $w(x_n)$ be the last vertex of $u(x_n)P(x_n)\mathring{v}(x_n)$ in

$$\{ u(x_n) \} \cup F \cup Q(x_1) \cup \ldots \cup Q(x_{n-1}),$$

and set $Q(x_n) := v(x_n)P(x_n)w(x_n)$. This satisfies (ii) for n.

Note that $w(x_n)$ cannot be on $S(x_n)$, because $w(x_n) \in P(x_n)\mathring{a}$ and $S(x_n) \subseteq aP(x_n)$. Thus, at least one of the following statements is true:

(1) $w(x_n) = u(x_n)$;

(2) $w(x_n) \in F - a$, and $w(x_n) <_F x_n$ (compare the definition of $v(x_n)$);

(3) there is an $m < n$ such that $w(x_n) \in \mathring{v}(x_m)Q(x_m)$.

If (1) holds, we join n to 0 in K and colour this edge blue. Then (iii)–(v) are true for n. Suppose now that (2) holds. Then (iii) is true for n, with $m = 1$ in (b). Let $m < n$ be maximal such that $x_m \leqslant_F w(x_n)$, and join n to m in K by a green edge. This satisfies (iv). By the maximality of m and (i) for $m+1$ ($\leqslant n$; recall that we verified (i) for n above),

$$w(x_n) <_F x_{m+1} \sim_F v(x_m)$$

as required by (v).

Suppose finally that (3) holds but neither (1) nor (2). Then $w(x_n) \notin F \cup \lfloor a \rfloor$; recall that $w(x_n)$ cannot be (on or) above $S(x_n)$ in F, by definition of $v(x_n)$. Let $m < n$ be minimal with $w(x_n) \in \mathring{v}(x_m)Q(x_m)$. Then $w(x_n) \neq w(x_m)$, by (iii) for m, so $w(x_n) \in \mathring{Q}(x_m)$ and (iii) holds for n. Join n to m in K by a red edge; this satisfies (iv) and makes (v) vacuous.

This completes the inductive definition of $x_1, x_2, \ldots$ and the corresponding paths $Q(x_n)$.

We now turn to the second step in the construction of our sail. We shall use the assertions of (5.14) freely. K is an infinite tree. By König's theorem, therefore, K either contains a ray or has a vertex of infinite degree. If m is a vertex of infinite degree, and M is an infinite set of neighbours $n > m$ of m, then the paths $Q(x_n)$ with $n \in M$ are pairwise disjoint except possibly for their last vertices $w(x_n)$, and these last vertices are all in the union of the three finite vertex sets of $\lfloor a \rfloor \cap C^-$ (if the edge nm is blue, i.e. if $m = 0$), of $\bigcup_{x <_F v(x_m)} S(x)$ (if nm is green), and of $\mathring{Q}(x_m)$ (if nm is red). Hence infinitely many of the paths $Q(x_n)$ have the same endvertex, w say. The corresponding paths $S(v(x_n))v(x_n)Q(x_n)$ then form an a–w bundle, which can be placed on any $P(x)$ that contains w. We are thus done by (5.13).

Thus, K contains a ray $N = n_0 n_1 \ldots$. We may assume that $n_0 = 0$, so by (iv) we have $n_i < n_j$ whenever $i < j$. For each $i > 0$, let us rename $S(x_{n_i})$, $Q(x_{n_i})$, $u(x_{n_i})$, $v(x_{n_i})$ and $w(x_{n_i})$ as S_i, Q_i, u_i, v_i and w_i. Note that N has exactly one blue edge (its first edge), while both the number of red edges and the number green edges may be either finite or infinite.

Let $e_0, e_1, \ldots$ be the green edges of N, in order. Then for consecutive green edges $e_{k-1} = \{ n_{i-1}, n_i \}$ and $e_k = \{ n_j, n_{j+1} \}$ (thus, $i \leqslant j$), we have

$$w_i <_F v_{i-1} \leqslant_F w_{j+1} <_F v_j \,. \tag{5.15}$$

Indeed, the two strict inequalities here are both consequences of (5.14)(v). For the middle inequality, observe that

$$v_{i-1} \sim_F x_{n_{i-1}+1} \leqslant_F x_{n_i} \leqslant_F x_{n_j} \leqslant_F w_{j+1} \,,$$

where the first relation comes from (5.14)(i), the third from the fact that $i \leqslant j$, and the fourth from the definition of green edges in (5.14)(iv).

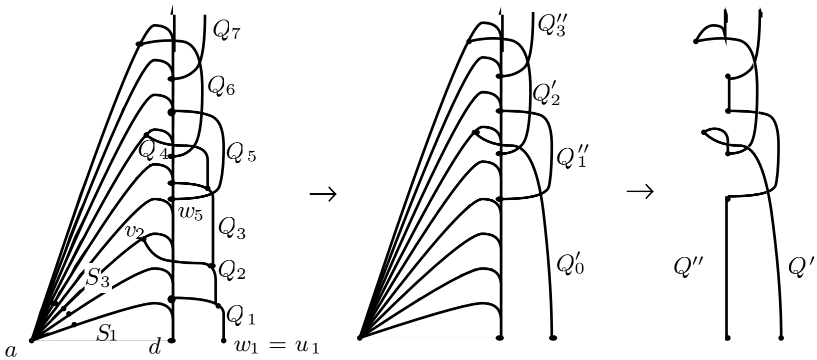

FIGURE 8. Building a sail

We finally turn to the third step in our construction and build a sail from F and the paths Q_i. We first combine some of the Q_i to longer paths Q'_j and Q''_j, which will in turn be combined into two rays Q' and Q'' for the mast and spine of our sail (Fig. 8).

For each $k > 0$ such that e_k is defined, let i and j be such that $e_{k-1} = \{ n_{i-1}, n_i \}$ and $e_k = \{ n_j, n_{j+1} \}$ (so $i \leqslant j$), and consider the unique v_j–w_i path in $\bigcup_{i \leqslant h \leqslant j} Q_h$. If k is even, call this path Q'_k (Fig. 9); if k is odd, call it Q''_k. For $k = 0$, let Q'_0 be the unique v_j–w_1 path in $\bigcup_{h \leqslant j} Q_h$, where $e_0 = \{ n_j, n_{j+1} \}$. Note that this path ends in $(w_1 =)\ u_1 \in C^- \cap \lfloor a \rfloor$; we shall rename this vertex u_1 as u. Finally, if N has a last green edge $e_k = \{ n_{i-1}, n_i \}$, consider the unique ray in $\bigcup_{h \geqslant i} Q_h$ starting at w_i, and call it Q'_∞ (if k is odd) or Q''_∞ (if k is even). Note that, by (5.14)(ii) and (5.15), the paths Q'_k and Q''_k defined in this paragraph are disjoint for different values of $k \in \mathbb{N} \cup \{ \infty \}$, except only that the first vertex of Q'_k might coincide with the last vertex of Q'_{k+2} (and likewise for Q''_k and Q''_{k+2}), where $Q'_{k+2} := Q'_\infty$ or $Q''_{k+2} := Q''_\infty$ if e_{k+1} is the last green edge. With the exception of the one vertex u, the paths Q'_k and Q''_k run entirely inside C and meet F only in their endvertices.

Similarly consider, for each $k > 0$ such that e_k is defined, the unique w_{j+1}–v_{i-1} path in $F - a$, where i and j are again such that $e_{k-1} = \{ n_{i-1}, n_i \}$ and $e_k = \{ n_j, n_{j+1} \}$. By (5.15), we have $v_{i-1} \leqslant_F w_{j+1}$, so every vertex v on this path satisfies

$$v_{i-1} \leqslant_F v \leqslant_F w_{j+1} \,. \tag{5.16}$$

If k is odd, call this path Q'_k; if k is even, call it Q''_k (Fig. 9). For $k = 0$, let Q''_0 be the unique w_{j+1}–d path in $F - a$, where $e_0 = \{ n_j, n_{j+1} \}$. Finally, if N has a last green edge $e_k = \{ n_{i-1}, n_i \}$, consider the unique ray in $F - a$ starting at v_{i-1}, and call it Q'_∞ (if k is even) or Q''_∞ (if k is odd). Note again that, by

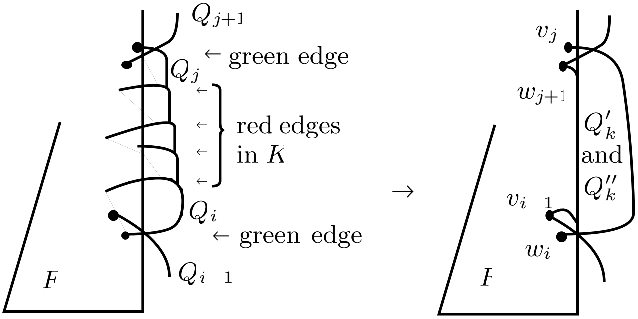

FIGURE 9. Defining the paths Q'_k and Q''_k

(5.15) and (5.16), the paths Q'_k and Q''_k defined in this paragraph are disjoint for different values of $k \in \mathbb{N} \cup \{\infty\}$, and they are all contained in F.

In summary, (5.15) and (5.16) imply that any path Q'_k running through $F - a$ meets only Q'_{k-1} and Q'_{k+1} among all the paths Q'_i and Q''_i. (For paths running outside F, the possible exception is that the first vertex of Q'_k may coincide with the last vertex of Q'_{k+2}, in which case Q'_{k+1} is the trivial path consisting of this vertex.) Similarly, Q''_k meets only Q''_{k-1} and Q''_{k+1}. An easy induction on k now shows that

$$Q' := \bigcup_{i \leqslant \infty} Q'_i \qquad \text{and} \qquad Q'' := \bigcup_{i \leqslant \infty} Q''_i$$

are two disjoint rays, starting at u and d, respectively (Fig. 10).

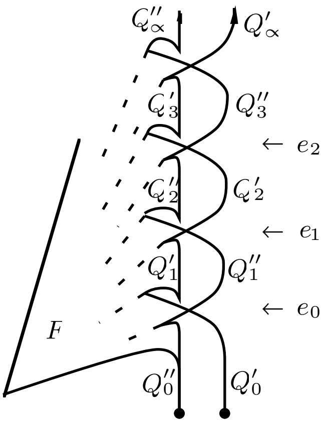

FIGURE 10. Disentangling mast and spine

We shall use Q' and Q'' for the mast and spine of our sail, as follows. If N has infinitely many green edges, we let Q' be the mast and Q'' the spine. Since both Q' and Q'' meet F infinitely often, F contains infinitely many disjoint

Q'–Q'' paths as rungs, and infinitely many a–Q'' paths, disjoint except for a, as spokes. Our sail is thus complete.

If N has only finitely many green edges, either Q' or Q'' shares a tail with Q; let this ray (Q' or Q'') be the spine, the other one the mast. We may now borrow our spokes from F (chosen high enough). As rungs, we use the paths $Q_n w_{n+1}$ (for large enough n) if $v_n \in Q$. If $v_n \in S(v_n) - Q$, we let s_n denote the first vertex of $aS(v_n)$ in Q, and choose $s_n S(v_n) v_n Q_n w_{n+1}$ as a rung. As there are infinitely many such spokes and rungs to choose from, it is easy to choose them disjoint.

We finally have to place our sail (or a fan derived from it) on one of the paths $P(x)$, $x \in X$. To this end, we shall choose $x \in X$ so that $P := P(x)$ satisfies the requirements for placing sails listed in Lemma 3.3. If N has a green edge, this is easy. Let $\{ n_j, n_{j+1} \}$ be its first green edge; then v_j is the starting vertex of Q'_0, and hence on Q'. Now choose x so that $v_j \in S(x)$. Then $aP(x)$ meets both Q' (in v_j) and Q'' (in d), and the requirements in Lemma 3.3 are satisfied.

If N has no green edge, then $Q'' = Q$ corresponds in Lemma 3.3 to the ray Q_2, the sail's spine, while Q' corresponds to Q_1, its mast. Let $x := x_{n_1}$. Since $S(x)$ meets Q, we have $aP(x) \cap Q_2 \neq \emptyset$ as required. As in Lemma 3.3, let q_2^+ be the highest vertex on Q of $aP(x)$. If $P(x)a \cap q_2^+ Q = \emptyset$, the requirements of Lemma 3.3 are again satisfied. So assume that $P(x)a$ hits $q_2^+ Q$, say first in q_2^-. Then q_2^- lies on Q above $S(x)$, and hence cannot precede $v(x)$ on $P(x)$ (by definition of $v(x)$). Therefore $u \in P(x)v(x) \subset P(x)q_2^-$; recall that $u = u_1 = w_1$, which is on $P(x)v(x)$ by the choice of x. Thus $P(x)q_2^- \cap Q_1 \neq \emptyset$ (as $u \in Q' = Q_1$), and the requirements of Lemma 3.3 are again satisfied.

With the placing of our sail on a path $P(x)$ we have thus satisfied the requirement of (5.13), and thereby completed the induction step for (5.8). The proof of Theorem 1 is thus complete.

References

[[1]] B. Bollobás, *Extremal Graph Theory*, Academic Press, London 1978.

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

[[3]] R. Diestel and I. Leader, A proof of the bounded graph conjecture, *Invent. math.* **108** (1992), 131–162.

[[4]] R. Diestel, Dominating functions and topological graph minors, *Contemporary Mathematics* **147** (1993), 461–476.

[[5]] R. Diestel and I. Leader, The growth of infinite graphs: boundedness and finite spreading *Combinatorics, Probability and Computing* **3** (1994) 51–55.

[[6]] R. Diestel and J.M. Brochet, Normal tree orders for infinite graphs, *Trans. Amer. Math. Soc.* **345** (1995), 871–895.

[[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*, Akademische Verlagsgesellschaft, Leipzig 1936 (reprinted: Chelsea, New York 1950).

[[10]] C. Thomassen, communication at the 10th Bielefeld Combinatorial Colloquium, 1990.

Reinhard Diestel, Faculty of Mathematics, TU Chemnitz, 09107 Chemnitz, Germany.