

A Proof of the Bounded Graph Conjecture

Reinhard Diestel and Imre Leader

An infinite graph is called bounded if for every labelling of its vertices with natural numbers there exists a sequence of natural numbers which eventually exceeds the labelling along any ray in the graph. We prove an old conjecture of Halin, which characterizes the bounded graphs in terms of four forbidden topological subgraphs.

1. Introduction

Let $\sigma, \tau: \mathbb{N} \rightarrow \mathbb{N}$ be two sequences of natural numbers. We say that σ *dominates* τ if $\sigma(n) \geq \tau(n)$ for every n greater than some $n_0 \in \mathbb{N}$. Now let G be an infinite graph, and assume the vertices of G are labelled with natural numbers, say by $f: V(G) \rightarrow \mathbb{N}$. Then each *ray* (one-way infinite path) $R = x_0x_1\dots$ in G gives rise to a sequence $\tau_R := (f(x_0), f(x_1), \dots)$; we say that σ *dominates* R if σ dominates this sequence τ_R . We may then ask whether or not there exists a sequence $\sigma: \mathbb{N} \rightarrow \mathbb{N}$ which dominates all the rays in G simultaneously. If this is the case, we say that the labelling f of G is *bounded* by σ ; if not, f is *unbounded*.

An unlabelled graph G is said to be *bounded* if every labelling of its vertices is bounded. Thus, G is bounded if and only if for every labelling $f: V(G) \rightarrow \mathbb{N}$ there exists a sequence $\sigma: \mathbb{N} \rightarrow \mathbb{N}$ which dominates every ray in G with respect to f . If not, G will be called *unbounded*.

Our main aim in this paper is to give a proof of what has become known as the ‘bounded graph conjecture’ [5]. This conjecture, proposed by Halin in 1964 (see [5]) but first published in [4], characterizes the bounded graphs by the exclusion of four simple prototypes of unbounded graphs as topological subgraphs. Not surprisingly, this characterization has some fundamental implications for the concept of boundedness.

Before we state the bounded graph conjecture in detail, let us look at a few examples of bounded and unbounded graphs. K_ω , the complete graph on a countably infinite set of vertices, is clearly unbounded: choose any labelling that uses infinitely many distinct labels, and there will be rays whose labellings grow faster than any fixed $\mathbb{N} \rightarrow \mathbb{N}$ sequence.

Another trivial example of an unbounded graph is the following. For a cardinal κ , let I_κ denote the disjoint union of κ rays. Then I_κ is unbounded for every $\kappa \geq 2^\omega$. Indeed, since there are only 2^ω different $\mathbb{N} \rightarrow \mathbb{N}$ sequences, we can accommodate every possible sequence as the labelling of one of the rays of I_κ . The overall labelling of I_κ then is clearly unbounded, and so I_κ is unbounded.

This trivial uncountable example of an unbounded graph is in fact the only uncountable aspect of the whole bounded graph problem: as we shall see, the task of classifying the bounded or unbounded graphs *without* uncountably many disjoint rays is an entirely ‘countable’ problem. (Indeed, we shall find that such a graph is bounded if and only if all its countable subgraphs are bounded (Corollary 7.3).)

Let us then look at some further countable examples. A graph with only countably many rays is bounded, due to the following well-known observation:

Lemma 1.1. *Every countable family of sequences $\sigma_n: \mathbb{N} \rightarrow \mathbb{N}$ ($n \in \mathbb{N}$) is dominated by a common sequence $\sigma: \mathbb{N} \rightarrow \mathbb{N}$.*

Proof. For each $n \in \mathbb{N}$, set $\sigma(n) := \max \{ \sigma_0(n), \dots, \sigma_n(n) \}$. □

For example, a single ray contains only countably many distinct rays as subgraphs (namely, all its ‘tails’), and is therefore bounded.

A similar argument shows that, more generally, every locally finite connected graph G is bounded. Indeed, given a labelling f and any fixed vertex v of G , it is easy to construct a sequence σ_v which dominates all the rays in G starting at v ; then Lemma 1.1 provides a sequence which dominates all these (countably many) sequences σ_v , and hence bounds f .

The example of K_ω may suggest that countable graphs have to be dense in order to be unbounded. However, this is not the case. T_ω , the tree of countably infinite regular degree,* is also unbounded. To show this, we have to find a labelling $f: V(T_\omega) \rightarrow \mathbb{N}$ such that no $\mathbb{N} \rightarrow \mathbb{N}$ sequence dominates all the rays in T_ω with respect to f . This is easily done. In fact, all we ask is that f be injective, i.e. that all the vertices of T_ω be labelled differently. Given $\sigma: \mathbb{N} \rightarrow \mathbb{N}$, we then choose a ray in T_ω as follows. Let x_0 be any vertex of T_ω whose label is greater than $\sigma(0)$. Since f is injective, x_0 has a neighbour x_1 whose label exceeds $\sigma(1)$. Continuing in this manner, we find a ray $x_0x_1 \dots \subset T_\omega$ such that $f(x_n) > \sigma(n)$ for every n , so this ray is certainly not dominated by σ .

The following example of an unbounded graph is based on the same idea. Let B be the graph obtained from a ray $v_0v_1 \dots$ by adding, for each $n \in \mathbb{N}$, a countably infinite set of independent $v_{3n+1}v_{3n+3}$ paths of length 2 (Fig. 1). Any subdivision of B will be called a *bundle graph*. (See §2 for the definition of subdivision etc.)

To show that B is unbounded, let f be any injective labelling of B , and assume we are given a sequence $\sigma: \mathbb{N} \rightarrow \mathbb{N}$. Let us specify a ray $R = x_0x_1 \dots$ as follows. Put $x_0 := v_0$ and $x_1 := v_1$. For x_2 , we have an infinite choice of neighbours of v_1 , all labelled differently. We may thus choose as x_2 a neighbour of v_1 whose label is greater than $\sigma(2)$. We continue with $x_3 := v_3$ and $x_4 := v_4$, where again we have an infinite choice for x_5 . Proceeding in this manner, we

* This means that every vertex of T_ω has exactly ω neighbours; see §2.

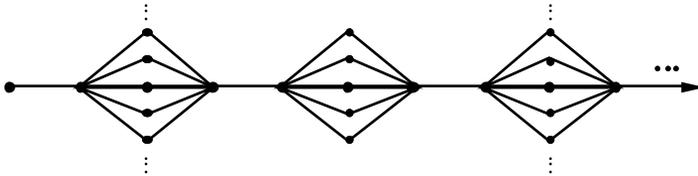


FIGURE 1. The prototype bundle graph B

may choose R in such a way that $f(x_n)$ exceeds $\sigma(n)$ for every third value of n , and thus σ fails to dominate R .

Our last example is similar to B . Let F be the graph obtained from a ray $V = v_0v_1\dots$ by adding disjoint rays P_2, P_4, P_6, \dots with $P_k \cap V = \{v_k\}$, and joining v_{2n+1} to all the new vertices of P_{2n+2} for every $n \in \mathbb{N}$ (Fig. 2). Any subdivision of F will be called a *fan graph*.

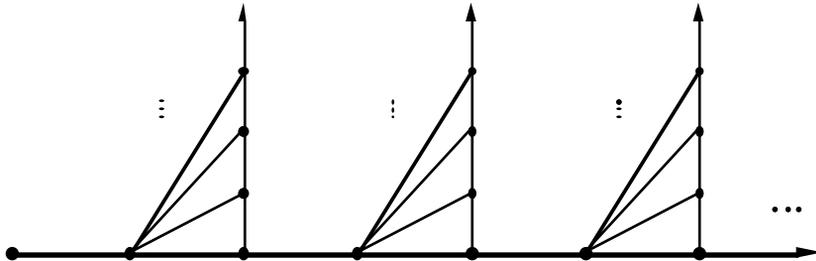


FIGURE 2. The prototype fan graph F

F is again unbounded, the proof being essentially the same as for B . Indeed, given any injective labelling $f: V(F) \rightarrow \mathbb{N}$ and any sequence $\sigma: \mathbb{N} \rightarrow \mathbb{N}$, we can easily find a ray R through F whose labels exceed the corresponding values of σ again and again. All we have to ensure when defining R is that we start at v_0 , and never use an edge of one of the paths P_k in its ‘upward’ direction. (This would force us to trace out the entire tail of P_k , leaving us unable to return to a vertex of type v_{2n+1} with an infinite choice ahead.)

It is perhaps remarkable that F can be made bounded by what would seem to be an inessential change. If the ‘fans’ in F are flipped horizontally, i.e. if the vertices v_{2n+1} are joined to all the vertices of P_{2n} rather than to those of P_{2n+2} (add a ray P_0), the resulting graph is bounded. (This example is due to Halin; he found it so discouraging that it prompted him to doubt the possibility of a proof of his conjecture by elementary methods [5].)

We are now ready to state Halin’s conjecture. If a graph G has a subgraph isomorphic to a subdivision of a graph X , we say that X is a *topological*

subgraph of G .

Bounded Graph Conjecture. (Halin)

Assume the Continuum Hypothesis. A graph is bounded if and only if it contains none of the following graphs as a topological subgraph: I_{ω_1} ; T_ω ; B ; F .

The forward implication of the conjecture should be obvious: we have seen that the graphs I_{2^ω} , T_ω , B and F are all unbounded; these proofs remain essentially the same for subdivisions of the latter three graphs (and identical for $I_{\omega_1} = I_{2^\omega}$), and it is clear that no bounded graph can have an unbounded subgraph.

A word concerning the Continuum Hypothesis (CH). When Halin originally proposed the bounded graph conjecture (in connection with Rado's seminal paper on *Universal graphs and universal functions* [9]), it was understood to refer to countable graphs only, and T_ω , B and F were the only forbidden (topological) subgraphs. There was no mention of CH. Halin later strengthened the conjecture to include graphs of any cardinality, which resulted in the additional exclusion of I_{2^ω} . This in turn suggested the assumption of CH, as it seemed unclear whether or not I_κ is bounded when κ is a cardinal strictly between \aleph_1 and 2^ω .

In fact, it is well-known among set theorists that this latter question is independent of the usual axioms of set theory. Martin's Axiom, for example, implies that every I_κ with $\kappa < 2^\omega$ is bounded, but it is consistent with ZFC that there exist $\kappa < 2^\omega$ for which I_κ is unbounded.

It should be stressed that our proof of the bounded graph conjecture is entirely free from such set theoretic considerations. To make this possible, we take the following (standard) approach. We shall define a cardinal \mathfrak{b} in such a way that it is the least cardinal κ for which I_κ is unbounded (whatever the value of that cardinal may be), and then prove—without assuming CH—that a graph is bounded if and only if it contains none of $I_\mathfrak{b}$, T_ω , B or F as a topological subgraph. Since $\omega < \mathfrak{b} \leq 2^\omega$, and hence $\omega_1 = \mathfrak{b} = 2^\omega$ with CH, this in particular implies the bounded graph conjecture as stated above.

It turns out that we need to show very little about \mathfrak{b} to make this approach work. However, any reader who prefers to work with countable sets and is willing to accept CH will lose nothing by substituting 'countable' whenever he reads ' $< \mathfrak{b}$ '.

Before we get down to the proof of the bounded graph conjecture, let us look at some of its immediate implications. (See also §7.) As a first corollary, the conjecture implies that subdivisions of bounded graphs are again bounded, and that subdivisions of unbounded graphs are unbounded. The first of these two statements may not be unexpected: for countable graphs it follows without much effort from the definition of boundedness. The second assertion is more surprising. Indeed, it would seem that there should be graphs which

are unbounded largely because they are dense (as our example of K_ω), so that subdividing their edges sufficiently often might make them bounded.

Together with the obvious fact that subgraphs of bounded graphs are bounded, the assertion that subdivisions of unbounded graphs are again unbounded is equivalent to saying that boundedness is closed under taking topological subgraphs. We shall also derive a forbidden minor version of the bounded graph conjecture (Theorem 7.7), together with a corollary saying that boundedness is closed under taking minors (with finite branch sets).

Another consequence of the bounded graph conjecture is the fact that a graph is bounded if and only if all its countable subgraphs are bounded—unless, of course, it contains uncountably many disjoint rays, in which case it is trivially unbounded (assuming CH). Thus, boundedness is shown to be an essentially ‘countable’ property, irrespective of the cardinality of the graph considered.

Finally, the bounded graph conjecture implies that if a countable graph admits any unbounded labelling at all, then all its injective labellings are unbounded. (Recall that the latter is true for subdivisions of T_ω , and for bundle and fan graphs.) It is therefore never difficult to find an unbounded labelling for an unbounded graph: its unboundedness lies properly in its structure, not in the complication of any particular labelling.

On the face of it, the bounded graph conjecture seems rather a bold conjecture to make. There appears to be nothing special about the four unbounded graphs which it claims are the only prototypes—indeed it would seem that playing with a few examples should reveal many more similar ‘minimal’ unbounded graphs. This may partly explain the fact that over the last 25 years there seem to have been few serious attempts at proving the conjecture.

Halin himself published his own partial results only very recently. In a paper written in 1986 [4], he proves the bounded graph conjecture for countable trees. Note that of the four forbidden graphs from the conjecture this involves only T_ω ; graphs of the other three types cannot even occur in a countable tree. A few years later, Halin extended this theorem to graphs whose blocks are rayless [5] (see also [2]). The extended result still involves only three of the four forbidden configurations, as fan graphs cannot occur. A very recent result of Halin, also found in [5], confirms the bounded graph conjecture for graphs not containing an infinite set of disjoint rays. (This last theorem builds on Halin’s well-known structural characterization of these graphs by ‘chain-like’ decompositions.)

The following may help to motivate the bounded graph conjecture a little better, taking some of the arbitrariness out of its list of forbidden configurations. (It will also be used later, for a proof of Lemma 2.3.) Let G be a connected unbounded graph. Since locally finite connected graphs are bounded, G has a vertex v of infinite degree. Let us assume that every edge incident with v can be extended to a ray (which is not an enormous assumption), and

let R_0, R_1, \dots be an infinite set of rays all starting at v . Can R_0, R_1, \dots be chosen pairwise independent (i.e., disjoint except for v)? If so, they might be seen as the beginning of a subdivided T_ω in G , and we might try to find similar rays starting at other vertices of the R_n to continue the construction.

If the rays R_0, R_1, \dots cannot be chosen independent, it seems reasonable to expect that we can choose them in such a way that there exists a finite set of vertices, and hence a single vertex w , which each of them meets. We may now apply König's theorem (see §2) to combine suitable initial segments of the rays R_n into what might be the first 'bundle' of a bundle graph, or the first 'fan' of a fan graph in G .

Indeed, let us choose initial segments P_n of the R_n as follows. Let P_0 be the initial segment of R_0 that ends at w . Having defined P_0, \dots, P_n for some $n \in \mathbb{N}$, let P_{n+1} be the initial segment of R_{n+1} which ends at the first vertex of R_{n+1} that is in $P_0 \cup \dots \cup P_n$. Let this vertex be called w_{n+1} , and let $k(n+1)$ be the minimal k for which $w_{n+1} \in P_k$. Call $P_{k(n+1)}$ the *predecessor* of P_{n+1} .

Let T be the tree whose vertices are the paths P_n ($n \in \mathbb{N}$), and in which each vertex P_n (for $n = 1, 2, \dots$) is joined by an edge to its predecessor $P_{k(n)}$. By König's theorem, T contains a ray or has a vertex of infinite degree. If T has a vertex of infinite degree, P_k say, then P_k contains a vertex x such that $w_n = x$ for infinitely many $n > k$. The paths P_n with $w_n = x$ are pairwise independent (i.e., disjoint except for v and x), and they form a 'bundle' (Fig. 3).

On the other hand if T contains a ray, without loss of generality starting at P_0 , then the union of the paths $P_n \subset G$ which are the vertices of this ray are easily seen to form a 'fan' (Fig. 3). (Note that, by the construction of T , if P_m succeeds P_n on the ray in T then $m > n$, so P_n is the predecessor of P_m .)

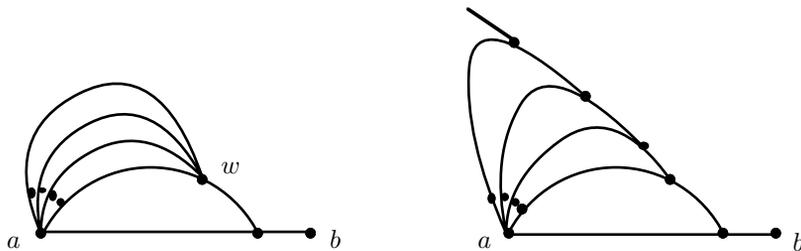


FIGURE 3. A bundle or a fan on a path $P_k = v \dots w$

The main purpose of the above considerations is to persuade the reader that fans and bundles are indeed natural ingredients of unbounded graphs. It must be said, however, that these considerations have not taken us anywhere near a proof of the bounded graph conjecture. Indeed, one might well find various fans and bundles in an unbounded graph: the problem is that, in general, they will be far from disjoint. Our task will thus be to prove that unbounded graphs must contain fans and bundles in such abundance that it is possible to

select an infinite set of them as tidily threaded on a ray as in the graphs B and F .

The paper is organized as follows. We begin with a short section on notation and basic lemmas. This is followed by a section on normal spanning trees, the main structural tool in our proof, and another which contains two simple Ramsey-type lemmas concerning partitions of T_ω . The bounded graph conjecture is then proved in Sections 5 and 6, with a concluding section listing some immediate corollaries.

2. Notation and basic lemmas

All the graphs considered in this paper 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. The first infinite ordinal is denoted by ω , the first uncountable ordinal by ω_1 .

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 . If xy is an edge (for $x, y \in V(G)$), then y is a *neighbour* of x , and the two vertices are *adjacent*. 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.

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*, and a graph is *rayless* if it contains no 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.

If $P = x_1 \dots x_n$ is a path, we write $\overset{\circ}{P}$ for the *interior* of P , the subpath $x_2 \dots x_{n-1}$. The vertices of $\overset{\circ}{P}$ are the *inner* vertices of P . Similarly, if $1 \leq i \leq j \leq n$, we set $Px_i := x_1 \dots x_i$, $P\overset{\circ}{x}_i := x_1 \dots x_{i-1}$, $x_iPx_j := x_i \dots x_j$, $x_jP := x_j \dots x_n$ and $\overset{\circ}{x}_jP := x_{j+1} \dots x_n$ 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 , and P an *extension* of Px_i . Note that a finite path is an initial segment of itself.

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 $G \setminus (X \cup Y)$. 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 $X \subset V(G)$, we write $G[X]$ for the subgraph of G induced by the vertices in X . For $H \subset G$ and $x \in G$, 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 $G \setminus H$ containing x , then $H[x]$ contains precisely those vertices of H that have a neighbour in C .

Two or more paths are *independent* if their interiors are disjoint. If A , B and X are subgraphs of G or subsets of $V(G)$, we say that X *separates* A from B in G if every A - B path in G contains a vertex of X . We use this notation even when A and B are not disjoint; of course, X must then contain $A \cap B$.

When T is a tree (an acyclic connected graph), 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 $x \leq y$ 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 $x \leq z$ is linearly ordered; the subgraph it spans in T will be denoted by $[z]$. The *down-closure* $[U]$ of a set $U \subset V(T)$ is the graph $\bigcup \{[u] \mid u \in U\} \subset T$. For any vertex v of T , the subtree of T spanned by all the vertices $x \geq v$ will be denoted by $T^{\geq v}$. Similarly, we write $T^{>v}$ for $T^{\geq v} \setminus \{v\}$.

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

A *subdivision* of a graph X , denoted by TX , is any graph obtained from X by replacing its edges with independent paths of lengths ≥ 1 . The original vertices of X are then called the *branch vertices* of the TX ; the other vertices of the TX are its *subdividing vertices*. X is a *topological subgraph* of any graph containing a TX . If X is a tree rooted at some vertex r , then any TX will again be assumed to be rooted at r ; the order on $V(X)$ is thus embedded in the tree order on the vertices of the TX .

The following definition and lemma will not be needed until §7. They are included here for completeness, but the reader may wish to skip straight to König's theorem further below.

In analogy to our notation concerning topological subgraphs, we say that a graph H is an HX (where 'H' comes from 'homomorphism') if its vertex set admits a partition $\{V_x \mid x \in V(X)\}$ into subsets V_x spanning connected subgraphs in H so that H contains a V_x - V_y edge if and only if x and y are adjacent in X . The sets V_x are called the *branch sets* of H . A graph X' is said to be a *minor* of H if H has a subgraph H' such that $H' = HX'$ (or, equivalently, if $H = HX$ with $X' \subset X$).

Lemma 2.1. *Let X be a topological subgraph of a graph G . If X is countable or locally finite, then X is a minor of G with all branch sets finite.*

Proof. By assumption, the components of X all have countable vertex sets; choose an enumeration for each of them. Let H be a TX contained in G . For each $x \in V(X)$, let V_x be the subset of $V(H)$ that consists of x and any subdividing vertices on x - y edges of X for which y precedes x in their common component of X . Note that each V_x is finite, and that it spans a connected subgraph in H . Hence, H is an HX as required. \square

The following theorem of König [7], 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 contains either a vertex of infinite degree or a ray.* \square

We conclude with some less standard terminology.

It will often be useful to be able to combine several sequences of natural numbers into one fast growing sequence which dominates them all. For $\sigma: \mathbb{N} \rightarrow \mathbb{N}$, let $\sigma', \sigma'': \mathbb{N} \rightarrow \mathbb{N}$ be defined by setting $\sigma'(n) := \max\{\sigma(0), \dots, \sigma(n)\}$ and $\sigma''(n) := \sigma'(2n)$. Then σ'' dominates σ' , which in turn dominates σ . The sequence σ'' will be called the *majorant* of σ . More generally, let F be a family of $\mathbb{N} \rightarrow \mathbb{N}$ sequences dominated simultaneously by a sequence σ . We shall then also refer to the majorant σ'' of σ as a majorant of the family F . Note that, by Lemma 1.1, every countable family of $\mathbb{N} \rightarrow \mathbb{N}$ sequences has a majorant.

The majorant of a single sequence is thus obtained by first making it monotonically increasing and then doubling its growth rate. It therefore ‘catches up’ with the original sequence even when shifted forward by any constant number of terms. This is the idea behind the following lemma, which is essentially due to Halin [5]:

Lemma 2.2. *Let F be a family of $\mathbb{N} \rightarrow \mathbb{N}$ sequences, and let G be a graph with a labelling $V(G) \rightarrow \mathbb{N}$. If $R, R' \subset G$ are rays sharing a tail, and if some $\sigma \in F$ dominates R , then any majorant of F dominates R' .* \square

Let a and b be distinct vertices of a graph G . The union of an infinite set of independent a - b paths of lengths ≥ 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. The *interior* of a bundle is the union of its fibres. 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 $\{B_1, \dots, B_n\}$ is a (possibly empty) set of disjoint bundles on a common path P , arranged in such a way that the first vertex of P is in none of the B_i and both vertices of $B_i \cap P$ precede both vertices of $B_{i+1} \cap P$ on P for all $i < n$, then the union of P and the B_i is a *partial bundle graph* on P .

Similarly, the union of a ray R with initial vertex b and infinitely many paths from a vertex $a \notin R$ to $R \setminus \{b\}$ is called a *fan from a to b* , or an *a - b*

fan, if those paths are pairwise disjoint except for a and have no inner vertices on R . The ray R in such a fan is called its *spine*; the a - R paths are its *spokes*. The *interior* of a fan F from a to b is the graph $F \setminus \{a, b\}$. F is said to be *on* a path P if $F \cap P = \{a, b\}$ and a precedes b on P . If $\{F_1, \dots, F_n\}$ is a (possibly empty) set of disjoint fans on a common path P , arranged in such a way that the first vertex of P is in none of the F_i and both vertices of $F_i \cap P$ precede both vertices of $F_{i+1} \cap P$ on P for all $i < n$, then the union of P and the F_i is a *partial fan graph* on P .

Using these terms, we can now write down as a lemma a precise summary of the observations we made at the end of the Introduction. The proof of the lemma was sketched there; the details are left to the reader.

Lemma 2.3. *Let a and b be distinct vertices of a graph G , and let \mathcal{P} be an infinite set of a - b paths in G whose second vertices are pairwise distinct. The union $\bigcup \mathcal{P}$ of these paths then contains a bundle B or a fan F . Moreover:*

- (i) *The bundle B can be chosen as an a - x bundle on a path $P \in \mathcal{P}$ (where x is some vertex on P), and so that every a - x path in B is an initial segment of a path in \mathcal{P} .*
- (ii) *The fan F —with spine R , spokes S_0, S_1, \dots , and $S_n \cap R =: \{v_n\}$ for each n —can be chosen so that v_n precedes v_m on R whenever $n < m$, and each of the paths $S_n v_n R v_{n-1}$ is an initial segment of a path in \mathcal{P} (where v_{-1} is the starting vertex of R). \square*

3. Normal spanning trees

In this section we introduce and develop the main structural tool for our proof, 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 \leq_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 [6] characterized the infinite graphs having normal spanning trees; his characterization implies that every countable connected graph contains such a tree. Starting from Jung’s result and using the theory of simplicial decompositions of graphs—see [1]—Halin [3] was able to prove the following:

Theorem 3.1. (Halin)

If G is connected and $G \not\cong \text{TK}_\omega$, then G has a normal spanning tree.

When we come to prove that an unbounded graph G must contain one of the four topological subgraphs listed in the bounded graph conjecture, we may thus assume that G has a normal spanning tree: if not, it would contain a TK_ω , and hence a TT_ω , as desired.

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.

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 \leq_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 3.2. *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 .*

Proof. Let $P = x_1 \dots x_n$ be an x - y path in G , and put $z := \inf \{x, y\}$; we have to show that P meets $\lfloor z \rfloor$. This is trivial if x and y are comparable, so let us assume they are incomparable. Let w be the upper neighbour of z such that $w \leq x$. Then $w \not\leq y$; let x_i be the last vertex of P with $w \leq x_i$. Then $x_{i+1} \in \lfloor z \rfloor$, because T is normal. \square

Note that if $x \in T'$ then $\lfloor x \rfloor \subset T'$. Lemma 3.2 thus implies the following:

Lemma 3.3. *If $P = x \dots y$ is a T' - T' path in G , then x and y are comparable in S .*

Proof. If x and y are incomparable, then neither of them is in $\lfloor x \rfloor \cap \lfloor y \rfloor$, and so $P \cap \lfloor x \rfloor \cap \lfloor y \rfloor \neq \emptyset$ by Lemma 3.2. Since $\lfloor x \rfloor \cap \lfloor y \rfloor \subset T'$, this means that P is not a T' - T' path. \square

Note that Lemma 3.2 does *not* imply that the *interior* of an x - y path meets $\lfloor x \rfloor \cap \lfloor y \rfloor$. The T' - T' path in Lemma 3.3, for example, may have a non-empty interior lying in $G \setminus T'$ somewhere above x and y .

Let S be a subtree of T rooted at r . (In applications, S will be either a normal ray or equal to T' .) Our next lemma says that any vertex of S which is ‘directly accessible’ from some vertex $x \notin S$ (by a path in G whose interior avoids S) must lie below x .

Lemma 3.4. *Let S be a subtree of T rooted at r , and let $x \in V(G)$. Then $S[x] \subset \lfloor x \rfloor$.*

Proof. If $x \in S$, then $S[x] = \{x\} \subset [x]$. If $x \notin S$, let $P = x \dots y$ be an x - S path in G . By Lemma 3.2, P meets $[x] \cap [y] \subset S$, while $\overset{\circ}{P} \cap S = \emptyset$ by the choice of P . Hence $P \cap [x] \cap [y] = \{y\}$, so $y \in [x]$ as required. \square

Lemma 3.4 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 in \leq_T . Let us define the R -height of x , denoted by $\text{ht}(x; R)$, to be the unique maximal vertex of $R[x]$. For example, 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 $\overset{\circ}{h}R$.

The concept of height will provide essential help in organizing our constructions of subgraphs of G such as fan and bundle graphs. For $H \subset G$ or $H \subset V(G)$, we shall say that the R -height of (the vertices in) H is *bounded* if R has a vertex h such that $\text{ht}(x; R) \leq h$ for every vertex $x \in H$. Another subgraph $H' \subset G$ (or $H' \subset V(G)$) is H -*clear* with respect to R if the R -height of each vertex of H' is strictly greater than the R -height of every vertex in H . Then, for example, if we are trying to find a fan graph in G and H is a partial fan graph constructed so far (on the path P , say), we shall aim to extend P into H' and continue our construction there; we can then be sure that fans constructed at later stages will not interfere with earlier fans.

The first assertion in our next lemma is immediate from the definition of R -height. The second is a consequence of the fact that, by definition, the graph $R[x]$ coincides for all vertices x in a common component of $G \setminus R$.

Lemma 3.5. *Let $x, y \in G$, and let R be a normal ray.*

- (i) *If $y \in R[x]$, then $\text{ht}(x; R) \geq y = \text{ht}(y; R)$.*
- (ii) *If x and y are in a common component of $G \setminus R$, then $\text{ht}(x; R) = \text{ht}(y; R)$.* \square

If R and R' are normal rays, and if $x \in R \setminus R'$ and $y \in R' \setminus R$, then $R \cap R'$ separates x from y in G (by Lemma 3.2). This observation extends to any pair of vertices whose R -height and R' -height, respectively, is sufficiently large:

Lemma 3.6. *Let R and R' be distinct normal rays; set $z := \max(R \cap R')$. Suppose that $x, y \in V(G)$ are such that $\text{ht}(x; R) > z$ and $\text{ht}(y; R') > z$. Then $[z] = R \cap R'$ separates x from y in G .*

Proof. Write $u := \text{ht}(x; R)$ and $v := \text{ht}(y; R')$, and let P be any x - y path in G . Moreover, let $x \dots u$ be an x - R path and $y \dots v$ a y - R' path in G . Since the walk $u \dots xPy \dots v$ meets $R \cap R'$ (by Lemma 3.2) but neither $x \dots u$ nor $y \dots v$ meets $R \cap R'$ (by definition), P meets $R \cap R'$ as claimed. \square

Our next lemma asserts that if the R -height of $x \notin R$ is sufficiently large, then its R' -height is the same as that of the vertices on $R \setminus R'$ itself, namely $z = \max(R \cap R')$.

Lemma 3.7. *Let R and R' be distinct normal rays, $z := \max(R \cap R')$, and $x \in V(G)$. If $\text{ht}(x; R) > z$ then $\text{ht}(x; R') = z$.*

Proof. Let P be an x - R path in G whose last vertex y is above z . Then $PyRz$ is an x - $(R \cap R')$ path, so $\text{ht}(x; R') \geq z$ by definition. Now if $\text{ht}(x; R') > z$ then Lemma 3.6, applied with $y = x$, would give that $\lfloor z \rfloor$ separates x from x , a contradiction. Hence $\text{ht}(x; R') = z$, as claimed. \square

If R is a normal ray and H is any subgraph of G , let us say that the R -height of H tends to infinity if for each $h \in R$ only finitely many vertices of H have R -height $\leq h$. Lemma 3.7 then has the following immediate consequence:

Lemma 3.8. *If R and R' are distinct normal rays, and if the R -height of $H \subset G$ tends to infinity, then the R' -height of H is bounded.* \square

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

Lemma 3.9. *Every ray in G follows a unique normal ray.*

Proof. Let $R \subset G$ be a ray. By Lemma 3.6, R can follow at most one normal ray. We show that such a ray exists, by constructing a normal ray $x_0x_1\dots$ which R meets infinitely often. Let us assume that R starts at the root r of T ; since this can be achieved by adding or deleting finitely many vertices of R , there will be no loss of generality.

Let $x_0 := r$. Suppose vertices x_0, \dots, x_n have been chosen so that they induce an increasing path in T and infinitely many vertices of R lie above x_n . Since R can pass through $\lfloor x_n \rfloor$ only finitely often, Lemma 3.2 implies that x_n has a unique upper neighbour y in T such that R has infinitely many vertices above y ; put $x_{n+1} := y$.

In order to show that R meets $x_0x_1\dots$ infinitely often, suppose x_n is the last vertex of R that belongs to $\{x_0, x_1, \dots\}$. (For the existence of such a vertex, recall that $x_0 = r \in R$ by assumption.) Let y be the next vertex on R , and let m be the r - y distance in T . Then y and x_m have the same distance from r in T , and $y \neq x_m$. Hence y and x_m must be incomparable. By definition of x_m , yR contains a vertex x above x_m . By Lemma 3.2, then, yR meets

$$\lfloor x \rfloor \cap \lfloor y \rfloor \subset \lfloor x_m \rfloor = \{x_0, \dots, x_m\},$$

contrary to the choice of x_n . \square

Note that, since normal rays are contained in T' , Lemma 3.9 implies in particular that $G \setminus T'$ is rayless.

Our next lemma says that if a ray R follows a normal ray R' , then, in terms of R' -height, the vertices of R behave like those of R' itself:

Lemma 3.10. *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.*

Proof. Let $x \in R'$ be given; we have to show that at most finitely many vertices of R have R' -height $\leq x$. Let $U := \{v \in V(G) \mid \text{ht}(v; R') \leq x\}$. Then $[x]$ separates $\hat{x}R'$ from U ; recall that $R'[v] \subset [x]$ for $v \in U$, by definition of $\text{ht}(v; R')$. Since R meets xR' infinitely often, this implies that $V(R) \cap U$ is finite. \square

Our last lemma shows how vertices with neighbourhoods of unbounded R -height can be used to find a copy of TK_ω in G .

Lemma 3.11. *Let R be a normal ray.*

- (i) *If the neighbourhood of a vertex $x \in G$ has unbounded R -height, then x is on R .*
- (ii) *If G has infinitely many vertices with neighbourhoods of unbounded R -height, then G contains a TK_ω .*

Proof. (i) If x is not in R , then the R -height of each of its neighbours is at most its own R -height, by Lemma 3.5.

(ii) Starting with $X_0 := \emptyset$, let us define inductively a nested sequence $X_0 \subset X_1 \subset \dots$ of subgraphs of G so that, for each n , $X_n \simeq \text{TK}_n$ and the neighbourhood of any branch vertex of X_n has unbounded R -height. Let $n > 0$, and suppose that $X_{n-1} \simeq \text{TK}_{n-1}$ has already been defined. Let x_1, \dots, x_{n-1} be the branch vertices of X_{n-1} ; the neighbourhoods of these vertices all have unbounded R -heights. Since X_{n-1} is finite, $G \setminus X_{n-1}$ contains a vertex x_n whose neighbourhood has unbounded R -height.

Let us define independent paths P_1, \dots, P_{n-1} joining x_n to x_1, \dots, x_{n-1} , respectively. Let $i \in \{1, \dots, n-1\}$ be given, and assume that P_j has been defined for all $j < i$. Let y_i and y_n be neighbours of x_i and x_n , respectively, whose R -heights z_i and z_n exceed the R -heights of all the vertices in

$$X_n^{i-1} := X_{n-1} \cup \bigcup_{j < i} P_j.$$

Let $y_i \dots z_i$ be a y_i - R path and $y_n \dots z_n$ a y_n - R path in G . Note that the vertices on these two paths have uniform R -heights of z_i and z_n , respectively. Let P_i be an x_i - x_n path whose interior is contained in the walk $y_i \dots z_i R z_n \dots y_n$. Then the vertices of $\overset{\circ}{P}_i$ have R -heights between z_i and z_n , so $\overset{\circ}{P}_i$ avoids X_n^{i-1} .

Putting

$$X_n := X_{n-1} \cup \bigcup_{i < n} P_i,$$

we see that $X_n \simeq \text{TK}_n$ as intended. Moreover, the branch vertices of X_n include those of X_{n-1} , and they all have neighbourhoods of unbounded R -height. This completes the induction step. It is clear that $\bigcup_{n \in \mathbb{N}} X_n$ is isomorphic to TK_ω . \square

4. Partitions of T_ω

In this section we collect for convenience two simple Ramsey-type lemmas, which will be needed repeatedly later.

Both our lemmas are about finding TT_ω subtrees in a TT_ω . If $T \simeq \text{TT}_\omega$, a *colouring* of T will be any map whose domain is the set of branch vertices of T ; the images of such a colouring are called *colours*. T is *monochromatic* (or *red*, *blue* etc.) if all its branch vertices are coloured by (mapped to) the same colour (red, blue etc.), and T is *injectively coloured* if no two of its branch vertices have the same colour.

The proof of our first lemma is straightforward.

Lemma 4.1. *If T is a rooted TT_ω with a red-blue colouring, then T contains either a red TT_ω of the form $T^{\geq v}$ (for some branch vertex v of T) or a blue TT_ω (whose minimal vertex is one of its branch vertices). \square*

As a point of reference, let T^* be a fixed copy of a rooted T_ω , given with an enumeration e_1^*, e_2^*, \dots of its edges such that, for each $n \in \mathbb{N}$, the graph $T_n^* \subset T^*$ consisting of the root of T^* and the edges e_1^*, \dots, e_n^* (together with their endvertices) is connected. Such an enumeration clearly exists; for example, one could first list the edges of T^* using *any* enumeration, and then assemble T^* , starting from the root, by adding at each of ω steps the minimal edge from the list that keeps the covered portion of T^* connected. The n th edge e on the list will then be added after no more than $k(n)$ steps, where $k(n)$ is the sum of the list positions of all the edges on the path P in T^* from the root to the far endvertex of e (induction on the length of P).

We shall use this *standard construction* of T^* in the proof of Lemma 4.2 below, and again in §6. Note that, for each n , exactly one of the two endvertices of e_n^* is in T_{n-1}^* . This vertex will be denoted by x_n^* , the other endvertex of e_n^* by y_n^* .

Lemma 4.2. *Any coloured rooted TT_ω contains either a monochromatic TT_ω or an injectively coloured TT_ω , whose minimal vertex (in either case) is among its branch vertices.*

Proof. Let T be a coloured TT_ω . Assume first that the colouring of T has the following property:

(4.3) *For each vertex $v \in T$, the subtree $T^{\geq v}$ uses infinitely many colours.*

Following the standard construction of T^* , we build up an injectively coloured $\text{TT}_\omega =: T' \subset T$, say with branch vertices x_n and y_n corresponding to the vertices x_n^* and y_n^* of T^* , as follows. We start by picking any branch vertex of T as x_1 . Later, whenever we come to choose a subdivided edge $x_n \dots y_n$ of T' corresponding to an edge e_n^* of T^* , we first select an upper neighbour v of x_n

that does not lie in the (finite) part of T' constructed so far. We then use (4.3) to choose y_n among the branch vertices of $T^{\geq v}$ so that its colour differs from the (finitely many) colours used by branch vertices of T' selected earlier, and let $x_n \dots y_n$ be the unique x_n - y_n path in T .

Suppose now that (4.3) fails, i.e. that T has a vertex v such that $T^{\geq v}$ uses only finitely many colours. Without loss of generality, v is a branch vertex. We may now obtain a monochromatic TT_ω inside $T^{\geq v}$ by repeated application of Lemma 4.1: we just partition the colours used into two classes called ‘red’ and ‘blue’, apply the lemma and iterate. \square

For a thorough study of Ramsey-type properties of T_ω , the reader is referred to Prömel and Voigt [8].

5. Good families

In our introductory discussion of the three countable types of forbidden subgraph in the bounded graph conjecture, the TT_ω , the bundle graph, and the fan graph, we noticed that the proofs of these graphs’ own unboundedness were very similar. In the case of T_ω , we just took any injective labelling and observed that, for any given sequence $\sigma: \mathbb{N} \rightarrow \mathbb{N}$, we could easily find a ray R through T_ω which was not dominated by σ . Indeed, no matter how we had chosen an initial segment of R , we would have an infinite choice of labels for the next vertex of R , which could thus be chosen larger than the corresponding term in σ .

For a bundle graph or a fan graph, finding such a ray R was hardly more difficult. All we had to make sure of was that the paths P we considered for initial segments of R belonged to a certain family \mathcal{P} : a family of paths which could, again and again, be extended to reach unused vertices of infinite degree. In the bundle graph B , \mathcal{P} consisted of the paths from left to right; in the fan graph F , of those towards the right and down the spines of the fans.

The idea of finding such a family of paths in an unbounded graph is central to our proof of the bounded graph conjecture. Our aim in this section, essentially, is to prove that in every unbounded graph such a family of paths exists. In the next section we shall show that one can always reconstruct one of the forbidden subgraphs from the paths in that family.

Let us make precise the kind of family of paths in which we are interested, so that we can state exactly what we shall prove in this section. Let G be an arbitrary connected graph with a normal spanning tree T , and let $T' \subset T$ be the union of all normal rays in G . Call a pair (x, y) of distinct vertices of T' *good* if G contains only finitely many T' - T' paths from x to y , and *bad* otherwise. A non-empty family \mathcal{P} of finite paths in G will be called a *good family* if it satisfies the following four conditions:

- (i) for every path $P \in \mathcal{P}$ there exists an $n \in \mathbb{N}$ such that P has infinitely many extensions in \mathcal{P} of length n ;
- (ii) if $P \in \mathcal{P}$, $x, y \in V(P)$, and xPy is a $T'-T'$ path in G of length ≥ 2 (i.e. with at least 2 edges), then (x, y) is a good pair;
- (iii) every path $P \in \mathcal{P}$ starts at the root of T ;
- (iv) if $P \in \mathcal{P}$ and Q is an initial segment of P , then $Q \in \mathcal{P}$.

Note that the essence of a good family is contained in condition (i): if we choose n minimal, then infinitely many extensions $P' \in \mathcal{P}$ of P agree up to their n th (= penultimate) vertex and differ pairwise at their last vertex. Thus, if P is viewed as the beginning of a ray R being constructed to elude a given sequence σ (as in our discussion above), then the $(n+1)$ th vertex of this ray can be selected from an infinite choice, and therefore in such a way that its label exceeds the corresponding term in σ . (So, by condition (i) alone, the existence of a good family of paths in G implies that G is unbounded.)

Condition (ii) says, less formally, that paths of a good family leave and return to T' only at good pairs of vertices. This assumption is also vital to our proof: it ensures that subgraphs constructable from the paths in a good family have many vertices on T' —for example, we shall see that any vertex of infinite degree has to be on T' —which will give us the amount of control over these subgraphs that we shall need.

Conditions (iii) and (iv) are normalization conditions; they are not essential, but they will make it easier to handle good families.

Recall that, if κ is a cardinal, I_κ denotes the disjoint union of κ rays. Let us define \mathfrak{b} to be the least cardinal κ such that there exists a family of κ distinct $\mathbb{N} \rightarrow \mathbb{N}$ sequences which are not dominated simultaneously by any $\mathbb{N} \rightarrow \mathbb{N}$ sequence. (This cardinal \mathfrak{b} is commonly known as the *bounding number*). Then, in particular, every family of fewer than \mathfrak{b} sequences has a majorant. Note also that $\omega < \mathfrak{b} \leq 2^\omega$ by Lemma 1.1.

As is easy to check, \mathfrak{b} is *regular*, i.e. \mathfrak{b} is not the union (sum) of fewer than \mathfrak{b} cardinals $< \mathfrak{b}$. Moreover, since I_κ contains exactly $\omega \cdot \kappa = \kappa$ rays whenever κ is an infinite cardinal, \mathfrak{b} is also the least cardinal such that $I_{\mathfrak{b}}$ is unbounded.

We can now state the theorem which we shall prove in this section.

Theorem 5.1. *Let G be an unbounded graph with a normal spanning tree, and suppose that G contains neither an $I_{\mathfrak{b}}$ nor a bundle graph. Then there exists a good family of paths in G .*

Proof. As usual, let the normal spanning tree of G be denoted by T , and the union of all normal rays in G by T' .

Let us show straight away that

$$(5.2) \quad |T'| < \mathfrak{b}.$$

Since T' is connected and \mathfrak{b} is regular, it suffices to show that each vertex $v \in T'$ has degree $< \mathfrak{b}$ in T' . Let W be the set of upper neighbours of v in T' and, for each $w \in W$, let R_w be any normal ray containing w . Then the tails wR_w for $w \in W$ are pairwise disjoint. Since $G \not\geq I_{\mathfrak{b}}$ by assumption, this means that v has degree $< \mathfrak{b}$, as claimed.

The proof of the theorem is organized as follows. We begin by constructing a family \mathcal{P} of paths in G which satisfies the four defining conditions of a good family, but which will not, a priori, be known to be non-empty. Such a family is easily obtained: we shall start with the family \mathcal{P}_0 of *all* paths in G which have any chance of eventually being in \mathcal{P} (omitting, for example, any paths which violate condition (ii)), and then recursively discard those paths from this family which fail to meet the extendability condition (i).

The main part of the proof will be to show that \mathcal{P} is indeed non-empty and therefore a good family. The key here will be that the recursion by which \mathcal{P} is obtained takes fewer than \mathfrak{b} steps. At each step α , we shall find a sequence $\sigma_\alpha: \mathbb{N} \rightarrow \mathbb{N}$ which dominates all those rays in G whose initial segments were all in \mathcal{P}_0 , and for which some initial segment was discarded at step α . Then, by the definition of \mathfrak{b} , there will be a sequence σ which dominates *all* such rays in G simultaneously. Since G is by assumption unbounded, and therefore contains a ray R which is not dominated by σ , this means that the initial segments of such a ray R (after some normalization) are either still in \mathcal{P} when the recursion ends, or were not all in \mathcal{P}_0 at the start of the recursion. In the former case \mathcal{P} will be non-empty, as desired. In the latter case, it will turn out that R leaves and returns to T' infinitely often at bad pairs (x, y) ; we shall then be able to construct a bundle graph in G along R (typically with the bundles occurring between these x and y), in contradiction to our assumptions.

Let us define \mathcal{P} precisely. Let \mathcal{Q}_0 be the set of all paths $P \subset G$ starting at the root of T and ending on T' which only leave T' at good pairs of vertices, i.e., which have the property that if xPy is a T' - T' path in G of length ≥ 2 then (x, y) is a good pair. Let \mathcal{P}_0 be the set of all initial segments of paths in \mathcal{Q}_0 ; note that \mathcal{P}_0 includes \mathcal{Q}_0 . Then \mathcal{P}_0 satisfies conditions (ii), (iii) and (iv) from the definition of a good family.

We claim that $|\mathcal{P}_0| < \mathfrak{b}$. As clearly $|\mathcal{P}_0| = |\mathcal{Q}_0|$, it suffices to show that $|\mathcal{Q}_0| < \mathfrak{b}$. Now, each path $P \in \mathcal{Q}_0$ is determined by its (finite) sequence of vertices on T' and the choice of its subpaths xPy for successive $x, y \in V(P) \cap V(T')$. As $|T'| < \mathfrak{b}$, there are only $< \mathfrak{b}$ finite sequences of vertices of T' , and since each xPy has length ≥ 2 only if (x, y) is good, the choice of subpaths xPy is finite for each pair (x, y) . Hence $|\mathcal{Q}_0| < \mathfrak{b}$, as required.

For all ordinals $\alpha > 0$, let us inductively define subsets \mathcal{P}_α of \mathcal{P}_0 , as follows. Suppose first that α is a successor, say $\alpha = \beta + 1$. If \mathcal{P}_β contains a path $P =: P_\beta$ which violates (i) in the definition of a good family (i.e., a path P such that, for each $n \in \mathbb{N}$, \mathcal{P}_β contains only finitely many extensions of P of length n), let \mathcal{P}_α be obtained from \mathcal{P}_β by deleting P_β and all its extensions $P' \in \mathcal{P}_\beta$. If

\mathcal{P}_β contains no such path P , let $\mathcal{P}_\alpha := \mathcal{P}_\beta$. If α is a limit, let $\mathcal{P}_\alpha := \bigcap_{\beta < \alpha} \mathcal{P}_\beta$.

Since $|\mathcal{P}_0| < \mathfrak{b}$, there exists an $\alpha < \mathfrak{b}$ such that $\mathcal{P}_{\alpha+1} = \mathcal{P}_\alpha$; let α^* be the least such α , and set $\mathcal{P}_{\alpha^*} =: \mathcal{P}$. Since each \mathcal{P}_α satisfies conditions (ii), (iii) and (iv) from the definition of a good family, and since $\mathcal{P} = \mathcal{P}_{\alpha^*}$ in addition satisfies condition (i), we see that \mathcal{P} is a good family if and only if it is non-empty. To complete the proof, then, it suffices to show that if \mathcal{P} is empty then G contains a bundle graph, contrary to our assumption.

Let us suppose that \mathcal{P} is indeed empty. We shall first prove that G contains a ray R which leaves T' infinitely often at bad pairs (x, y) , and then use R to find a bundle graph in G .

(5.3) *There exists a ray $R \subset G$ which contains infinitely many T' - T' paths xRy such that (x, y) is a bad pair of vertices.**

To prove (5.3), let us suppose for a contradiction that any ray in G leaves T' only finitely often at bad pairs of vertices. We show that G is bounded, contrary to our assumptions.

Let $f: V(G) \rightarrow \mathbb{N}$ be an arbitrary labelling of G . We shall define sequences $\sigma_\alpha: \mathbb{N} \rightarrow \mathbb{N}$, one for each $\alpha < \alpha^*$, whose majorant σ^* will bound G . For each $\alpha < \alpha^*$, let the k th term of σ_α be defined as

$$\sigma_\alpha(k) := \begin{cases} f(x_k) & \text{if } k \leq |P_\alpha| \\ \max \{ f(v(P)) \mid P \in \mathcal{P}_\alpha \setminus \mathcal{P}_{\alpha+1} \text{ and } |P| = k \} & \text{if } k > |P_\alpha|, \end{cases}$$

where x_k is the k th vertex of P_α and $v(P)$ is the last vertex of P . Recall that, by definition, P_α has only finitely many extensions $P \in \mathcal{P}_\alpha$ of any given length; since $\mathcal{P}_\alpha \setminus \mathcal{P}_{\alpha+1}$ is precisely the set of P_α and all its extensions in \mathcal{P}_α , the maximum used above is therefore just the maximum of a finite set. Finally, let σ^* be the majorant of $\{\sigma_\alpha \mid \alpha < \alpha^*\}$; it exists, because $\alpha^* < \mathfrak{b}$.

To show that σ^* bounds G , let $R \subset G$ be any ray. By assumption, R leaves T' only finitely often at bad pairs of vertices; let R' be a tail of R which leaves T' only at good pairs. (Recall that since R follows a normal ray (Lemma 3.9), it never leaves T' altogether.) Let Q be a path in T' from the root of T to R' , and let v be the last vertex of Q . Then the ray $R'' := QvR' (= QvR)$ still leaves T' at good pairs only, so every initial segment of R'' is in \mathcal{P}_0 .

Now since $\mathcal{P} = \emptyset$ by assumption, each initial segment of R'' is discarded at some (non-limit) step in the recursive definition of \mathcal{P} ; let $\alpha < \alpha^*$ be minimal such that $\mathcal{P}_\alpha \setminus \mathcal{P}_{\alpha+1}$ contains an initial segment of R'' . Then P_α is itself an initial segment of R'' . Moreover, all the extensions of P_α in R'' are also in \mathcal{P}_α (by the minimality of α), and were hence discarded together with P_α . Thus,

* As the reader may expect (and the proof of (5.3) will indeed show), the paths xRy can be chosen so as to have lengths ≥ 2 . However, this will not be needed when (5.3) is used below.

$\mathcal{P}_\alpha \setminus \mathcal{P}_{\alpha+1}$ contains every initial segment of R'' of length $k \geq |P_\alpha|$. Therefore σ_α dominates R'' , by definition of σ_α .

By Lemma 2.2, this implies that σ^* dominates R , so σ^* bounds G as claimed. As this contradicts our assumption that G is unbounded, we have thus completed the proof of (5.3).

Let $R \subset G$ be a ray as in (5.3). Let $(x_1, y_1), (x_2, y_2), \dots$ be an infinite sequence of disjoint bad pairs of vertices, such that x_i precedes y_i on R for each i and y_i precedes x_j on R for $i < j$. For each i , let H_i be the union of all $T'-T'$ paths in G from x_i to y_i ; note that, since (x_i, y_i) is bad, H_i is infinite. We remark that the graphs H_i are not necessarily disjoint, and that H_i may well meet R even outside the segment $x_i R y_i$.

For each i , let us find a bundle $B_i \subset H_i$, say from a vertex x'_i to a vertex y'_i , together with (possibly trivial) disjoint paths $C_i = x_i \dots x'_i$ and $C'_i = y'_i \dots y_i$ in H_i which avoid the interior of B_i . (These paths will be called the *connectors* of B_i ; Fig. 4.)

FIGURE 4. The bundles B_i and their connectors

To define B_i for a given i , recall that since $G \setminus T'$ is rayless (Lemma 3.9), H_i too is rayless. Since H_i is infinite and connected, it has a vertex v of infinite degree, by König's theorem. As H_i is the union of x_i - y_i paths, each of the edges $e = uv$ incident with v is contained in such a path $P(e)$. Now if v precedes u on $P(e)$ for infinitely many of these edges e then, by Lemma 2.3, the union of their paths $P(e)$ contains a v - w bundle B (for some w) on the segment vP of one of these paths P . (Lemma 2.3 does not give a fan, because H_i is rayless.) Put $C_i := Pv$ and $C'_i := wP$ in this case, and let B_i be obtained from B by deleting those fibres which meet Pv . On the other hand, if u precedes v on $P(e)$ for infinitely many of the edges e , then the union of the corresponding paths $P(e)$ contains a w - v bundle B (for some w) on the segment Pv of one of these paths P (apply Lemma 2.3 'backwards'). In that case, we let $C_i := Pw$ and $C'_i := vP$, and let B_i be obtained from B by deleting those fibres which meet vP . (This construction of a bundle in a graph such as H_i is due to Halin [5].)

Now that we have defined an infinite sequence of bundles along our ray R , it may seem that obtaining a bundle graph from them would be only a small step away. As remarked earlier, however, the subgraphs H_i containing our bundles need not be disjoint. In fact, it might even happen that two of the

bundles, as well as their connectors, coincide (except for the two vertices x_i and y_i).

We shall deal with this problem as follows. By Lemma 3.9, R follows some normal ray R' . A natural way to ensure that different bundles B_i are disjoint, therefore, is to separate them by their R' -height: if we can find an infinite subsequence i_1, i_2, \dots spaced out sufficiently that each B_{i_k} is B_{i_j} -clear with respect to R' for every $j < k$, then these bundles B_{i_k} will indeed define a bundle graph along R . For this approach to work it is essential, of course, that the R' -height of each bundle B_{i_k} be bounded. Unfortunately, this is not necessarily the case—so let us deal with this problem first.

Suppose there are infinitely many i such that the R' -height of B_i is unbounded. Since $(H_i \setminus \{x_i, y_i\}) \cap T' = \emptyset$, every fibre of B_i is disjoint from R' , and so the vertices of each individual fibre have constant R' -height. If the R' -height of B_i is unbounded, this implies that x'_i has a neighbourhood of unbounded R' -height, and hence that $x'_i = x_i$ (Lemma 3.11 (i)). If this happens for infinitely many i , the corresponding vertices x_i —which are all distinct—form an infinite set of vertices each with a neighbourhood of unbounded R' -height. Then G contains a TK_ω (and hence a bundle graph) by Lemma 3.11 (ii).

We may therefore assume from now on:

(5.4) *The R' -height of B_i is bounded for all but finitely many i .*

Using R , the bundles B_i and their connectors C_i and C'_i , let us construct a bundle graph in G inductively, adding one bundle at a time. More formally, let us construct by induction on n a nested sequence $(G_n)_{n \in \mathbb{N}}$ of graphs with the following properties:

- (i) G_n is a partial bundle graph (with n bundles), on a path P_n which ends in a vertex $t_n \in R$;
- (ii) P_n has a vertex $s_n \neq t_n$ such that all the bundles of G_n are between vertices preceding s_n on P_n ;
- (iii) $s_n P_n \subset R$, and s_n precedes t_n on R ;
- (iv) $t_n R$ is S_n -clear with respect to R' , where $S_n := G_n \setminus \overset{\circ}{s}_n P_n$.

(See Figure 5.)

FIGURE 5. The partial bundle graph G_n

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

(5.5) G_n has bounded R' -height.

To define G_0 , let s_0 be the first vertex of R , and put $G_0 := P_0 := R t_0$ where $t_0 \in R$ is chosen so that $t_0 R$ is $\{s_0\}$ -clear with respect to R' ; this can be done by Lemma 3.10 (which we shall use freely in our further construction). Suppose now we have defined G_k , for $k = 0, \dots, n$, in accordance with conditions (i)–(iv). We shall now define G_{n+1} .

By (5.4) and (5.5), choose $i \in \mathbb{N}$ so that B_i has bounded R' -height and $x_i R$ is G_n -clear with respect to R' . Note that $t_n R x_i$ meets G_n only in t_n : it avoids S_n by (iv), and it meets $s_n R t_n$ only in t_n , because s_n precedes t_n on R by (iii) and t_n precedes x_i on R by the choice of x_i . Note that also

$$H'_i := (C_i \cup B_i \cup C'_i) \setminus \{x_i, y_i\}$$

is G_n -clear with respect to R' : since each $v \in H'_i$ can be connected to x_i by a v - T' path, the R' -height of v is at least that of x_i (Lemma 3.5).

Using our assumption that B_i (and hence H'_i) has bounded R' -height, choose $u \in y_i R$ so that $u R$ is X -clear with respect to R' , where

$$X := G_n \cup H'_i.$$

Let B'_i be obtained from B_i by removing those fibres that have a vertex on $t_n R u$; as there are at most finitely many such fibres, B'_i is still a bundle. The interior of B'_i then avoids $t_n R u$, and it also avoids G_n (because H'_i does, and H'_i contains the interior of B'_i).

The connectors C_i and C'_i , however, may still meet $t_n R x_i$ or $y_i R u$. Let x''_i be the last vertex on C_i that is in $t_n R u$, and let y''_i be the first vertex on C'_i that is in $t_n R u$ (Fig. 6). Rename x''_i and y''_i as a and b , where a precedes b on R , and set

$$U := G_n \cup t_n R a \cup (x''_i C_i x'_i \cup B'_i \cup y'_i C'_i y''_i) \cup b R u.$$

Then U is a partial bundle graph, with $n + 1$ bundles, whose R' -height is bounded.

In order to complete the induction step formally, set $s_{n+1} := u$, and choose $t_{n+1} \in u R$ so that $t_{n+1} R$ is U -clear with respect to R' . Note that $u R t_{n+1} \cap U = \{u\}$, because $U \subset X \cup R u$ and $u R t_{n+1}$ is X -clear by the choice of u . Set $G_{n+1} := U \cup u R t_{n+1}$, and let

$$P_{n+1} := P_n \cup t_n R a \cup (x''_i C_i x'_i \cup F \cup y'_i C'_i y''_i) \cup b R t_{n+1},$$

FIGURE 6. Disentangling R and H'_i

where F is some x'_i - y'_i path in B'_i . Then G_{n+1} is a partial bundle graph on P_{n+1} , satisfying conditions (i)–(iv) for $n+1$.

The graph $\bigcup_{n \in \mathbb{N}} G_n$, finally, is a bundle graph as desired. This completes the proof of Theorem 5.1. \square

6. Fans and bundles from a good family tree

This section is devoted to the proof of the following theorem.

Theorem 6.1. *Let G be a graph with a normal spanning tree, and suppose there exists a good family of paths in G . Then G contains a TT_ω , a bundle graph, or a fan graph.*

Proof. As before, let the normal spanning tree of G be denoted by T , and the union of all normal rays by T' . Let \mathcal{P} be a good family of paths in G . The method of our proof is as follows.

We shall represent the structure of \mathcal{P} by a certain tree \mathcal{T} , which will act as something like a covering space for the paths $P \in \mathcal{P}$ in G . Thus, the paths of \mathcal{P} will ‘lift’ to paths in \mathcal{T} , which will be arranged in such a way as to reflect the relation of extension between their projections in \mathcal{P} .

Our main task will be to study the relationship between \mathcal{T} and G . If G (more precisely, the subgraph $\bigcup \mathcal{P}$ of G) resembles \mathcal{T} closely, we shall try to find a tree of type TT_ω in G . On the other hand, if the paths of \mathcal{P} intersect a lot more (in G) than do their lifts in \mathcal{T} , we shall seek to construct a bundle graph or a fan graph in G . The distinction between the latter two cases will depend on the relationship between \mathcal{T} and our normal spanning tree T of G .

Recall that the essence of the fact that \mathcal{P} is a good family of paths lies in its extension property, condition (i) of the definition of a good family: for each of its paths, \mathcal{P} contains infinitely many extensions of some common length. We start by ‘pruning’ \mathcal{P} to a subfamily containing no more paths than the extension property requires.

To be precise, let us define subsets $\mathcal{Q}_0, \mathcal{Q}_1, \dots$ of \mathcal{P} as follows. Pick any non-trivial path $P_0 \in \mathcal{P}$, and set $\mathcal{Q}_0 := \{P_0\}$. Having defined $\mathcal{Q}_0, \dots, \mathcal{Q}_n$, we want \mathcal{Q}_{n+1} to be a set of extensions of the paths in \mathcal{Q}_n just large enough to witness the fact that $\mathcal{Q}_n \subset \mathcal{P}$. For each $P \in \mathcal{Q}_n$, take the minimal k such that P has infinitely many extensions of length k in \mathcal{P} . Then infinitely many of these extensions agree up to their penultimate vertices: recall that \mathcal{P} is closed under taking initial segments. We may thus choose infinitely many extensions $P^{(0)}, P^{(1)}, \dots$ of P in \mathcal{P} , each of length k , which differ only in their last vertices. Set

$$\mathcal{Q}_{n+1} := \{P^{(i)} \mid P \in \mathcal{Q}_n, i \in \mathbb{N}\},$$

and continue inductively. If we now let \mathcal{Q} be the set of all initial segments of paths in $\bigcup_{n \in \mathbb{N}} \mathcal{Q}_n$, we see that \mathcal{Q} is a good family.

In structure, \mathcal{Q} is rather similar to a TT_ω . Let us make this precise. Call a graph H a *long-rooted* TT_ω if H is a subdivision of a rooted tree whose root r has degree 1 and all of whose other vertices have countably infinite degree. The edge incident with r in this tree then corresponds to a path P in H , which will be called the *long root* of H . (Thus, P is the unique path $r \dots x \subset H$ for which $H \setminus P \hat{x} \simeq \text{TT}_\omega$.) To avoid confusion, we may refer to r itself as the *root vertex* of H .

Now define a tree \mathcal{T} as follows. The vertex set of \mathcal{T} is \mathcal{Q} itself. For $P, P' \in \mathcal{Q}$, we join P to P' by an edge whenever P' is an extension of P with $|P'| = |P| + 1$, i.e. if P may be obtained from P' by removing its last vertex. Then \mathcal{T} is a long-rooted TT_ω , whose root vertex ϕ is the trivial path consisting of the first vertex of P_0 (which in turn is the root of T). Note that, for $\alpha, \beta \in V(\mathcal{T})$, we have $\alpha \leq_{\mathcal{T}} \beta$ if and only if α is an initial segment of β .

Any long-rooted TT_ω contained in \mathcal{T} and also rooted at ϕ will be called a *good family tree*. Clearly, the vertex set of such a good family tree is again a good family of paths in G . In fact, the only property of \mathcal{T} we shall ever use is that it is a good family tree: we shall repeatedly prune \mathcal{T} to a smaller good family tree \mathcal{T}' , and then rename \mathcal{T}' as \mathcal{T} .

For $\alpha \in V(\mathcal{T})$, let $v(\alpha)$ denote the last vertex of the path α . Note that if β is another vertex of \mathcal{T} and $\beta <_{\mathcal{T}} \alpha$, then the path $\beta \subset G$ is a proper initial segment of the path α , and so $v(\beta) \neq v(\alpha)$. Likewise, $v(\alpha)$ and $v(\beta)$ will be distinct if the paths α and β were picked for \mathcal{Q} at the same step n , and as extensions $P^{(i)}$ of the same path $P \in \mathcal{Q}_{n-1}$. We may therefore note the following:

(6.2) *If \mathcal{T}' is a good family tree, and if $\alpha, \beta \in V(\mathcal{T}')$ are either comparable or upper neighbours of the same vertex $\gamma \in \mathcal{T}'$, then $v(\alpha) \neq v(\beta)$.*

By virtue of (6.2), v can be interpreted as mapping increasing paths $\pi \subset \mathcal{T}$ to ‘real’ paths $P \subset G$. For example, if $\pi = \xi_0 \dots \xi_n$ is a ϕ - α path in \mathcal{T} , then $v(\pi)$ is the path $v(\xi_0) \dots v(\xi_n)$ in G , which is none other than the path $\alpha \subset G$

itself. More generally, if $\pi \subset \mathcal{T}$ is an increasing path from α to β , and if $v(\alpha) =: a$, then $v(\pi)$ is precisely the subpath $a\beta$ of the path $\beta \subset G$. When \mathcal{T}' is any subgraph of \mathcal{T} other than a path, we may use $v(\mathcal{T}')$ informally to denote the set $\{v(\alpha) \mid \alpha \in \mathcal{T}'\} \subset V(G)$.

Let \mathcal{T}' be a good family tree. We shall call $\alpha \in V(\mathcal{T}')$ a *burst vertex* of \mathcal{T}' if α has infinite degree. (Thus, the burst vertices of \mathcal{T}' are precisely its branch vertices other than ϕ when \mathcal{T}' is viewed as a long-rooted TT_ω .) When $\alpha \in \mathcal{T}'$ is a burst vertex, the *burst* at α in \mathcal{T}' is the union of all the paths $\alpha \dots \gamma \subset \mathcal{T}'$ such that γ is minimal among the burst vertices above α . These paths will be called the *segments* of the burst; the burst itself will be denoted by $B_{\mathcal{T}'}(\alpha)$. When $\mathcal{T}' = \mathcal{T}$ or the precise choice of the good family tree considered is irrelevant, we may abbreviate $B_{\mathcal{T}'}(\alpha)$ to $B(\alpha)$. Thus, a burst $B(\alpha)$ is a subdivided infinite star, contained in \mathcal{T} and centred at α , its segments being the star's subdivided edges.

Let us call a burst $B(\alpha)$ *blocked* by a vertex $b \in G$ if infinitely many of its segments contain a vertex $\beta \neq \alpha$ with $v(\beta) = b$. If $B(\alpha)$ is not blocked by any vertex of G , it will be called *free*.

Passing to a good family tree $\mathcal{T}' \subset \mathcal{T}$ if necessary, we may assume that the following is true for \mathcal{T} :

(6.3) *Whenever a burst $B(\alpha)$ is blocked, there is a vertex $b_\alpha \in G$ such that every segment π of $B(\alpha)$ contains a (unique) vertex $\beta_\pi \neq \alpha$ with $v(\beta_\pi) = b_\alpha$.*

Moreover, since any segment of a burst $B_{\mathcal{T}'}(\alpha)$ in a good family tree \mathcal{T}' contains a segment of $B_{\mathcal{T}}(\alpha)$, we see that (6.3) is in fact true for all good family trees. (Recall that any good family tree is a subtree of \mathcal{T} .)

Let α_0 be the lowest burst vertex in \mathcal{T} . Then $[\alpha_0]$ is the long root of \mathcal{T} , and $\mathcal{T}^{\geq \alpha_0} \simeq \text{TT}_\omega$. Applying Lemma 4.1 to $\mathcal{T}^{\geq \alpha_0}$, we see that either there is a burst vertex $\alpha \in \mathcal{T}$ such that all bursts at vertices in $\mathcal{T}^{\geq \alpha} =: \tilde{\mathcal{T}}$ are free, or $\mathcal{T}^{\geq \alpha_0}$ has a subtree $\tilde{\mathcal{T}} \simeq \text{TT}_\omega$ (whose minimal vertex is a branch vertex) such that for every branch vertex α of $\tilde{\mathcal{T}}$ the burst $B_{\mathcal{T}}(\alpha)$ is blocked. In either case, let \mathcal{T}' be obtained from $\tilde{\mathcal{T}}$ by adding the $\phi\text{-}\tilde{\mathcal{T}}$ path in \mathcal{T} as a long root. Then \mathcal{T}' is a good family tree. Moreover, in the first case every burst in \mathcal{T}' is free, while in the second case (6.3) implies that every burst in \mathcal{T}' is blocked. Hence, passing to this good family tree \mathcal{T}' , we may assume that either all bursts in \mathcal{T} are blocked or all bursts in \mathcal{T} are free.

Let us quickly dispose of the case when all bursts are free, by showing how to find a TT_ω in G . If all the bursts in \mathcal{T} are free then, for every finite set $F \subset V(G)$ and every burst vertex α , there is a segment of $B(\alpha)$ such that for every vertex $\beta \neq \alpha$ on this segment we have $v(\beta) \notin F$. Using (6.2), and following the method of the proof of Lemma 4.2, we may thus construct a good family tree $\mathcal{T}' \subset \mathcal{T}$ on which the function v is injective. (Choose segments $\pi = \alpha \dots \gamma$ inductively as in the standard construction of $T^* \simeq T_\omega$, in such

a way that $v(\hat{\alpha}\pi)$ is disjoint from $v(\pi')$ for the—finitely many—segments π' chosen earlier.) The image of \mathcal{T}' under v then spans a long-rooted TT_ω in G .

We may therefore assume from now on that all bursts in \mathcal{T} are blocked. By (6.3), this assumption carries over to all good family trees:

(6.4) *If \mathcal{T}' is a good family tree and α is a burst vertex of \mathcal{T}' , then $B_{\mathcal{T}'}(\alpha)$ is blocked by the vertex $b_\alpha \in G$.*

Let us continue our investigation of the bursts in good family trees. We claim the following:

(6.5) *If α is a burst vertex in a good family tree \mathcal{T}' , then $v(\alpha) \in T'$.*

Indeed, suppose that $v(\alpha) \notin T'$. Let β be the highest vertex of \mathcal{T}' below α such that $v(\beta) \in T'$; β exists, because $v(\phi) \in T'$. Let π_0, π_1, \dots be the segments of $B_{\mathcal{T}'}(\alpha)$, and for each $n \in \mathbb{N}$ let ρ_n be a ray in \mathcal{T}' which starts at β and contains π_n . Since these rays are increasing in \mathcal{T}' , they correspond to rays $v(\rho_n) =: R_n$ in G (recall the remark following (6.2)). Each of the rays R_n starts at $v(\beta)$; let $t_n = v(\tau_n)$ be its first inner vertex on T' . This vertex t_n exists by Lemma 3.9, and it comes after $v(\alpha)$ on R_n by the choice of β . Moreover, the vertices t_n are all in $T'[v(\alpha)]$, which is finite by Lemma 3.4. Therefore infinitely many of the t_n are the same (t , say), and the corresponding paths $P_n := R_n t$ form an infinite family of $T'-T'$ paths in G from $v(\beta)$ to t . The paths P_n are distinct, because v is injective on the set of upper neighbours of α in \mathcal{T}' . Therefore $(v(\beta), t)$ is a bad pair. Since P_n is a subpath of the path $\tau_n \in V(\mathcal{T}')$ (for any n), this contradicts the fact that $V(\mathcal{T}')$ is a good family of paths in G . This completes the proof of (6.5).

Let $\alpha \in \mathcal{T}$ be a burst vertex. Let us show that $v(\hat{\alpha}\pi\hat{\beta}_\pi) \cap T' \neq \emptyset$ for all but finitely many exceptional segments π of $B(\alpha)$. In other words: for almost every segment of $B(\alpha)$ the corresponding path in G meets T' before it hits the blocking vertex b_α .

Suppose, for a contradiction, that $B(\alpha)$ has infinitely many segments π_0, π_1, \dots such that $v(\hat{\alpha}\pi_n\hat{\beta}_{\pi_n}) \cap T' = \emptyset$ for every n . We argue as in the proof of (6.5). For each n , let $P_n := v(\pi_n)$, and let t_n be the first vertex on $b_\alpha P_n$ that is in T' ; such a vertex exists, because the last vertex of P_n corresponds to another burst vertex of \mathcal{T} , and is therefore in T' by (6.5). Since each of the vertices t_n is in $T'[b_\alpha]$, which is finite, infinitely many of them are the same: t , say. The corresponding paths $P_n t$ are therefore $T'-T'$ paths in G from $v(\alpha)$ to t ; they are distinct, because the upper neighbours of α in \mathcal{T} have distinct images under v . Hence $(v(\alpha), t)$ is a bad pair, a contradiction.

For every burst vertex α , let us remove from \mathcal{T} the finitely many exceptional segments of $B(\alpha)$ (and all the vertices above these segments). So, by passing to a good family tree $\mathcal{T}' \subset \mathcal{T}$ if necessary, we may assume that, for every burst vertex $\alpha \in \mathcal{T}$ and every segment π of $B(\alpha)$, the corresponding path

$v(\pi) \subset G$ meets T' before it hits the blocking vertex b_α . Again, this carries over to all good family trees:

(6.6) *If α is a burst vertex in a good family tree T' , and if π is a segment of $B_{T'}(\alpha)$, then $v(\hat{\alpha}\pi\hat{\beta}_\pi) \cap T' \neq \emptyset$. The first vertex τ of $\hat{\alpha}\pi\hat{\beta}_\pi$ with $v(\tau) \in T'$ will be denoted by τ_π , and we write t_π for $v(\tau_\pi)$.*

When R is a normal ray, let us extend the definition of R -height—and the notion of ‘ (\cdot) -clear’ associated with it—from $V(G)$ to $V(\mathcal{T})$ in the natural way, setting $\text{ht}(\alpha; R) := \text{ht}(v(\alpha); R)$.

Let α be a burst vertex in some good family tree. We shall call its burst $B(\alpha)$ in this tree a *long burst* if there exists a normal ray $R \subset G$ such that the neighbourhood of α in $B(\alpha)$ has unbounded R -height. Otherwise, $B(\alpha)$ is a *short burst*.

For each burst vertex $\alpha \in \mathcal{T}$ where $B(\alpha)$ is a long burst, let us pick a normal ray $R_\alpha \subset G$ such that the neighbourhood of α in $B(\alpha)$ has unbounded R_α -height. We may then select infinitely many segments of $B(\alpha)$ so that the neighbours of α on these segments have distinct R_α -heights. The R_α -height on the set of these neighbours then tends to infinity. Passing to a good family tree $T' \subset \mathcal{T}$ if necessary, we may therefore assume that, for each burst vertex $\alpha \in T'$ with $B(\alpha)$ long, the R_α -height of its neighbours in $B(\alpha)$ tends to infinity.

The same is true then for every infinite subset of the set of upper neighbours of α in \mathcal{T} . Our assumption therefore extends to all good family trees:

(6.7) *Whenever $B(\alpha)$ is a long burst in some good family tree, the R_α -height of the neighbours of α in $B(\alpha)$ tends to infinity.*

As earlier when we considered free versus blocked bursts, let us apply Lemma 4.1 to the tree $\mathcal{T}^{\geq \alpha_0} \simeq \text{TT}_\omega$, where α_0 is the lowest burst vertex in \mathcal{T} . The lemma now implies that either there is a burst vertex $\alpha \in \mathcal{T}$ such that all bursts at vertices in $\mathcal{T}^{\geq \alpha} =: \tilde{\mathcal{T}}$ are short, or else $\mathcal{T}^{\geq \alpha_0}$ has a subtree $\tilde{\mathcal{T}} \simeq \text{TT}_\omega$ (whose minimal vertex is a branch vertex) such that for every branch vertex α of $\tilde{\mathcal{T}}$ the burst $B_{\mathcal{T}}(\alpha)$ is long. Let T' be obtained from $\tilde{\mathcal{T}}$ by adding the $\phi\text{-}\tilde{\mathcal{T}}$ path in \mathcal{T} as a long root. Then T' is a good family tree. Moreover, in the first case every burst in T' is short, while in the second case (6.7) (applied to \mathcal{T}) implies that every burst in T' is long. Passing to this good family tree T' , we may thus assume that either all bursts of \mathcal{T} are short or all bursts of \mathcal{T} are long.

We shall consider these two cases separately.

Case 1: *All bursts in \mathcal{T} are short.*

We shall construct a bundle graph in G by induction, adding one bundle at a time. The bundle graph will consist of the image of a ray in \mathcal{T} , together with bundles made up from some of the bursts along the ray.

To be precise, let us construct a sequence $G_0 \subset G_1 \subset \dots$ of subgraphs of G , and a sequence $\alpha_0 < \alpha_1 < \dots$ of vertices of \mathcal{T} , so that the following four conditions are satisfied for every $n \in \mathbb{N}$:

- (i) G_n is a partial bundle graph (with n bundles) on the path $\alpha_n \subset G$;
- (ii) α_n is a burst vertex of \mathcal{T} ;
- (iii) the vertex $v(\alpha_n)$ has degree 1 in G_n ;
- (iv) the image of $V(\mathcal{T}^{>\alpha_n})$ under the map v is disjoint from $V(G_n)$.

Roughly speaking, condition (i) asserts that our construction has been successful so far, while conditions (ii)–(iv) will enable us to continue the construction without running into obstructions.

Note that the induction starts, with α_0 as any burst vertex of \mathcal{T} and G_0 as the path α_0 in G . Note also that once we have constructed the graphs G_n , their union is a bundle graph in G , as required.

Let us then turn to the induction step. Suppose we have constructed $G_0 \subset \dots \subset G_n$ and $\alpha_0 < \dots < \alpha_n$ in accordance with conditions (i)–(iv). Put $\alpha := \alpha_n$ and $a := v(\alpha)$, and let Π be an infinite set of segments of $B(\alpha)$. Recall that, for each $\pi \in \Pi$, $t_\pi = v(\tau_\pi)$ is the first vertex of $v(\hat{\alpha}\pi\hat{\beta}_\pi)$ in T' .

Note that the vertices t_π cannot be the same (say t) for infinitely many $\pi \in \Pi$: the pair (a, t) would then be bad, contradicting the fact that $av(\pi)t$ is a $T'-T'$ path contained in the path $\tau_\pi \in V(\mathcal{T})$ (for any π with $t_\pi = t$). Moreover, by Lemma 3.3 each t_π is comparable with a . As G has only finitely many vertices below a , it follows that $a < t_\pi$ for all but finitely many $\pi \in \Pi$. Replacing Π with a suitable infinite subset if necessary, we may therefore assume the following:

(6.8) *The vertices t_π ($\pi \in \Pi$) are all distinct and lie strictly above a in T' .*

Let

$$H := \bigcup_{\pi \in \Pi} aT't_\pi.$$

H is an infinite subtree of T' . Hence, by König's theorem, H contains either a ray or a vertex of infinite degree. If H contains a ray, then a tail of this ray is increasing in T , and is therefore a tail of some normal ray R . By construction of H and Lemma 3.5, the vertices t_π have unbounded R -height. The neighbours of α in $B(\alpha)$ then also have unbounded R -height (again by Lemma 3.5), which contradicts the fact that $B(\alpha)$ is short. Hence H has a vertex z of infinite degree.

Replacing Π with a suitable infinite subset, we may assume that every t_π lies above z , and that $[t_\pi] \cap [t_{\pi'}] = [z]$ whenever $\pi \neq \pi'$ (Fig. 7).

By Lemma 3.2, this implies that

(6.9) *If $\pi, \pi' \in \Pi$ are distinct, then any t_π – $t_{\pi'}$ path in G meets $[z]$.*

FIGURE 7. The graph $H \subset T'$

For each $\pi \in \Pi$, let

$$P_\pi := v(\pi\beta_\pi);$$

recall that $v(\beta_\pi)$ is the blocking vertex b_α for $B(\alpha)$. Then

$$\mathcal{P}_\Pi := \{P_\pi \mid \pi \in \Pi\}$$

is an infinite set of a - b_α paths in G whose second vertices are distinct. By Lemma 2.3, the union $\bigcup \mathcal{P}_\Pi$ of these paths contains a bundle or a fan; we claim it must contain a bundle:

(6.10) $\bigcup \mathcal{P}_\Pi$ contains a bundle B as specified in Lemma 2.3 (i).

To prove (6.10), suppose $\bigcup \mathcal{P}_\Pi$ contains no such bundle. Then, by Lemma 2.3, $\bigcup \mathcal{P}_\Pi$ contains a fan F as specified in Lemma 2.3 (ii). We shall construct, inductively, an infinite sequence Q_0, Q_1, \dots of disjoint paths in F , each joining two vertices of $\{t_\pi \mid \pi \in \Pi\}$. (Fig. 8). This will contradict (6.9), which asserts that these vertices are pairwise separated in G by the finite subgraph $[z]$.

FIGURE 8. Finding the paths Q_0, Q_1, \dots in F

Let $k \in \mathbb{N}$ and assume that Q_i has been chosen, for every $i < k$, so that $a \notin Q_i$. Let $\ell \in \mathbb{N}$ be large enough that none of the paths Q_i meets $v_\ell R$ or any of the spokes S_n with $n > \ell$. By Lemma 3.9, $v_\ell R$ has a vertex t on T' . Let m be such that $t \in v_{m-1} R v_m$. By the choice of F (see Lemma 2.3), the path $S_m v_m R v_{m-1}$ is an initial segment of a path $P_\pi \in \mathcal{P}_\Pi$. Since t_π is the first vertex of $\mathring{a}P_\pi$ in T' , it precedes t on P_π or coincides with it, and is therefore itself on the path $S_m v_m R v_{m-1}$. Now pick $t' \in \mathring{v}_m R \cap T'$, say between v_{n-1} and v_n on R , find $\pi' \in \Pi$ such that $S_n v_n R v_{n-1}$ is an initial segment of $P_{\pi'}$, and note that $t_{\pi'} \in S_n v_n R v_{n-1}$. Let Q_k be the t_π - $t_{\pi'}$ path in $S_m \cup v_{m-1} R v_n \cup S_n$. Then $a \notin Q_k$, and $Q_k \cap Q_i = \emptyset$ for all $i < k$. This completes the induction step, and hence the proof of (6.10).

Note that, according to its specification by Lemma 2.3 (i), the bundle B in (6.10) is an a - x bundle on a path $P \in \mathcal{P}_\Pi$, where x is some vertex of P . Moreover, each of the a - x paths in B is an initial segment of some path $P_\pi \in \mathcal{P}_\Pi$; let $F_\pi \subset P_\pi$ be the corresponding fibre of B .

Our aim is to obtain our new partial bundle graph G_{n+1} by adding the bundle B on to G_n . To ensure that G_{n+1} will be compatible with the disjointness condition (iv), however (for some suitable $\alpha_{n+1} \in \mathcal{T}$), we first have to delete some fibres from B .

We first wish to delete those fibres F_π from B which do not contain t_π . Before we do this, we have to check that only finitely many fibres of B will be affected, i.e. that $t_\pi \in F_\pi$ for all but finitely many of the fibres F_π . If this is not the case, then $F_\pi \subset G \setminus T'$ for infinitely many $\pi \in \Pi$ (by definition of the vertices t_π ; recall that F_π is an initial segment of $\mathring{a}P_\pi$). Then $t_\pi \in T' [x]$ for all these π , and hence $t_\pi \in [x]$ by Lemma 3.4. Since $[x]$ is finite, this means that the vertices t_π coincide for infinitely many of the fibres F_π , contradicting (6.8).

So, removing finitely many fibres from B if necessary, we may assume that $t_\pi \in F_\pi$ for each fibre F_π of B . Again removing finitely many fibres, we may assume that no fibre of B meets $[z]$. Hence, (6.9) implies the following:

(6.11) *The fibres of B are pairwise separated by $[z]$ in G .*

(In particular, $x \in [z]$.)

We are now almost ready to attach the bundle B to our partial bundle graph G_n . Choose $\beta \in B(\alpha)$ so that $v(\beta) = x$ and the path $\alpha T \beta$ has length at least 2. Among all $\gamma \geq \beta$, choose one with $|\gamma \cap [z]|$ maximal, and pick a burst vertex $\alpha_{n+1} > \gamma$. By the choice of γ , we have

$$(6.12) \quad v(\mathcal{T}^{\geq \alpha_{n+1}}) \cap [z] = \emptyset.$$

Indeed, if $\delta \in \mathcal{T}^{\geq \alpha_{n+1}}$ (and thus $\delta > \gamma$), then the path $\delta \subset G$ is a proper extension of the path γ . The choice of γ implies that $\delta \cap [z] \subset \gamma$; in particular, $v(\delta)$ (which is the last vertex of δ) is not in $[z]$.

As $G[v(\mathcal{T}^{\geq \alpha_{n+1}})]$ is connected, (6.11) and (6.12) imply that at most one fibre of B contains a vertex in $v(\mathcal{T}^{\geq \alpha_{n+1}})$; delete any such fibre from B . Moreover, only finitely many fibres of B can have a vertex on the (finite) path $\alpha_{n+1} \setminus \alpha_n (= v(\mathring{\alpha}_n \mathcal{T} \alpha_{n+1}))$; delete those fibres as well. We may now set

$$G_{n+1} := G_n \cup \alpha_{n+1} \cup B.$$

Let us quickly verify the induction hypothesis for $n+1$. We certainly have $G_n \subset G_{n+1}$, and α_{n+1} is a burst vertex of \mathcal{T} (assertion (ii)) strictly above α_n .

For assertion (i), recall first that α_{n+1} is an extension of the path α_n , and that $\alpha_{n+1} \setminus \alpha_n \cap G_n = \emptyset$ by condition (iv) for n , because $V(\alpha_{n+1} \setminus \alpha_n) \subset v(\mathcal{T}^{> \alpha_n})$. Hence, G_n is a partial bundle graph on α_{n+1} . The interior of the new bundle B is contained in the union of the paths $v(\mathring{\pi})$, where π ranges over the segments of the burst $B(\alpha) = B(\alpha_n)$. The interior of B is therefore again disjoint from G_n by condition (iv) for n . As the interior of B also avoids the path $\alpha_{n+1} \setminus \alpha_n$ (by construction; recall the deletion of fibres just before the definition of G_{n+1}), B is also a bundle on the path α_{n+1} . Finally, if B' is any of the bundles of G_n , then both vertices of $B' \cap \alpha_{n+1}$ precede $v(\alpha_n) (= a)$ on α_{n+1} by condition (iii) for n , and hence they also precede the two vertices of $B \cap \alpha_{n+1}$. Hence G_{n+1} is a partial bundle graph, with $n+1$ bundles, on the path $\alpha_{n+1} \subset G$.

By the choice of α_{n+1} as strictly above β , we know that $v(\alpha_{n+1})$ has degree 1 in G_{n+1} (assertion (iii)).

For assertion (iv), note that $v(\mathcal{T}^{> \alpha_{n+1}})$ is disjoint from $V(G_n)$ by condition (iv) for n , since $\alpha_{n+1} > \alpha_n$ and hence $\mathcal{T}^{> \alpha_{n+1}} \subset \mathcal{T}^{> \alpha_n}$. It is certainly disjoint from $V(\alpha_{n+1}) = v(\phi \mathcal{T} \alpha_{n+1})$, and it is finally disjoint from the interior of B , since we deleted the one fibre of B that could possibly meet $v(\mathcal{T}^{\geq \alpha_{n+1}})$. Hence, $v(\mathcal{T}^{> \alpha_{n+1}})$ is disjoint from $V(G_{n+1})$.

This completes the induction step in our construction of a bundle graph in G , and so completes Case 1.

Case 2: *All bursts of \mathcal{T} are long.*

Recall that, by (6.7), each burst vertex $\alpha \in \mathcal{T}$ is associated with a normal ray R_α , so that the R_α -height of the neighbours of α in $B(\alpha)$ tends to infinity. Let again α_0 be the lowest burst vertex in \mathcal{T} . Applying Lemma 4.2 to $\mathcal{T}^{\geq \alpha_0}$, we see that \mathcal{T} contains a good family tree \mathcal{T}' with the property that the rays R_α are either all the same or all distinct for the burst vertices $\alpha \in \mathcal{T}'$. Let us rename \mathcal{T}' as \mathcal{T} , and address the two cases in turn.

If all the rays R_α are the same, say R , choose a sequence of burst vertices $\alpha_0 < \alpha_1 < \dots$ in \mathcal{T} . Then $v(\alpha_0), v(\alpha_1), \dots$ are distinct vertices of G , whose neighbourhoods have unbounded R -height. By Lemma 3.11 (ii), this implies that G contains a TK_ω , and hence a TT_ω .

We may therefore assume from now on that the rays R_α are all distinct:

(6.12) *If α and α' are distinct burst vertices of \mathcal{T} , then R_α and $R_{\alpha'}$ are distinct normal rays in G .*

Similarly to our construction of a bundle graph in the short burst case, we shall now construct a fan graph in G along the image of a ray from \mathcal{T} . The construction will be more complicated, however: the individual fans will be made not just from bursts $B(\alpha)$ along that ray, but partly also from the rays R_α corresponding to these bursts.

To be precise, let us construct a sequence $G_0 \subset G_1 \subset \dots$ of subgraphs of G , and a sequence $\beta_0 < \beta_1 < \dots$ of vertices of \mathcal{T} , so that the following three conditions are satisfied for every $n \in \mathbb{N}$:

- (i) G_n is a partial fan graph (with n fans) on the path $\beta_n \subset G$;
- (ii) for every burst vertex $\gamma > \beta_n$ of \mathcal{T} , the R_γ -height of G_n is bounded;
- (iii) there is a burst vertex $\nu_n > \beta_n$ such that the image of $V(\overset{\circ}{\beta}_n \mathcal{T} \nu_n \cup \mathcal{T}^{>\nu_n})$ under the map v is disjoint from $V(G_n)$.

Again, roughly speaking, condition (i) asserts that our construction has been successful so far, while conditions (ii) and (iii) will enable us to continue the construction.

We apply induction on n . Taking $G_0 := \beta_0 := \phi$, and letting ν_0 be any burst vertex of \mathcal{T} , we see that the induction starts. Note also that once we have constructed the graphs G_n for all $n \in \mathbb{N}$, their union is a fan graph in G , as desired.

Let us then turn to the induction step. Suppose we have constructed subgraphs $G_0 \subset \dots \subset G_n$ of G and found vertices $\beta_0 < \dots < \beta_n$ and ν_n of \mathcal{T} in accordance with conditions (i)–(iii). In order to define G_{n+1} , we wish to construct a new fan to add to G_n . However, it turns out to be difficult to pinpoint directly a fan that misses the whole of G_n , and which can be added to G_n in a manner that allows the induction to continue. We shall instead define two fans (to be called F and F'), and our aim will be to show that at least one of these fans is suitable for attachment to G_n .

Put $\alpha := \nu_n$ and $R := R_\alpha$, and let $a := v(\alpha)$. As the path $\alpha \subset G$ is finite, condition (ii) implies that the R -height of $G_n \cup \alpha$ is bounded. We may therefore choose a neighbour $\delta = \delta_0$ of α in $B(\alpha)$ that is $(G_n \cup \alpha)$ -clear with respect to R (cf. (6.7)). Let $r_0 \in R$ be the R -height of δ . Let us further choose segments π_1, π_2, \dots of $B(\alpha)$ such that, writing δ_i for the neighbour of α on π_i and r_i for the R -height of δ_i , we have $r_0 < r_1 < r_2 < \dots$. For each $i \in \mathbb{N}$, denote $v(\delta_i)$ by d_i , and let P_i be a d_i - R path in G that ends in r_i (using the definition of R -height). See Figure 9.

Since $P_i \overset{\circ}{r}_i \cap R = \emptyset$, Lemma 3.5 implies that the R -height is constant on $P_i \overset{\circ}{r}_i$, and thus equal to the R -height of d_i . Hence we have:

(6.13) *For each $i \in \mathbb{N}$, every vertex of P_i has R -height r_i .*

FIGURE 9. Constructing a fan from $B(\alpha)$ and R_α

Define a fan F from a to $d := d_0$, as follows. The spine of F is the ray $P_0 r_0 R$. Its spokes are the paths $ad_i P_i$ for $i > 0$. (In other words, the i th spoke consists of the edge ad_i followed by the path P_i .) By (6.13), the spokes of F are pairwise disjoint (except for a) and avoid P_0 , so F is indeed a fan. As further consequences of (6.13), we have the following:

(6.14) *The interior of F is $(G_n \cup \alpha)$ -clear with respect to R .*

(6.15) *The R -height of F tends to infinity.*

In order to define our second fan F' , choose a burst vertex $\alpha' > \delta$, let $R' := R_{\alpha'}$, and put $a' := v(\alpha')$. By condition (ii), the R' -height of $G_n \cup \alpha'$ is bounded. Moreover, since $R' \neq R$, Lemma 3.8 and (6.15) tell us that the R' -height of F is bounded. Choose a neighbour $\delta' = \delta'_0$ of α' in $B(\alpha')$ that is $(G_n \cup F \cup \alpha')$ -clear with respect to R' . Put $d' := v(\delta')$, and define $\pi'_i, \delta'_i, r'_i, d'_i, P'_i$ and a fan F' from a' to d' exactly as above (Fig. 10).

Statements 6.13–6.15 then have the following equivalents for F' :

(6.13') *For each $i \in \mathbb{N}$, every vertex of P'_i has R' -height r'_i .*

(6.14') *The interior of F' is $(G_n \cup F \cup \alpha')$ -clear with respect to R' .*

(6.15') *The R' -height of F' tends to infinity.*

Setting $z := \max(R \cap R')$, we see by (6.15) and Lemma 3.7 that there are vertices in F with R' -height z . The R' -height of δ' , and hence of every vertex in F' , is therefore greater than z :

(6.16) *F' is $\lfloor z \rfloor$ -clear with respect to R' .*

FIGURE 10. The fans F and F'

Now let $U := \{x \in F \mid \text{ht}(x; R) \leq z\}$, and put $W := U \cup [z]$. Note that U , and hence W , is finite by (6.15). We claim that

(6.17) W separates F from F' in G .

Indeed, consider F' and $F \setminus W$. Each vertex of F' has R' -height greater than z , while each vertex of $F \setminus W$ has R -height greater than z . By Lemma 3.6, this implies that $[z]$ separates F' from $F \setminus W$ in G , so W separates F from F' as claimed.

Among all $\gamma \geq \delta'$, choose one with $|V(\gamma) \cap W|$ maximal, and pick a burst vertex $\nu > \gamma$. By the choice of γ , we have $v(\mathcal{T}^{\geq \nu}) \cap W = \emptyset$. Since $G[v(\mathcal{T}^{\geq \nu})]$ is connected, this implies that $v(\mathcal{T}^{\geq \nu})$ meets at most one of F and F' . Let $F'' \in \{F, F'\}$ be such that $v(\mathcal{T}^{\geq \nu}) \cap V(F'') = \emptyset$, and define α'' , R'' , a'' , δ'' , d'' , r''_i and P''_i accordingly. Thus,

(6.18) F'' is a fan from a'' to d'' which avoids $v(\mathcal{T}^{\geq \nu})$.

Moreover, statements 6.13–6.15 and their analogues for F' imply the following:

(6.13'') For each $i \in \mathbb{N}$, every vertex of P''_i has R'' -height r''_i .

(6.14'') The interior of F'' is $(G_n \cup \alpha'')$ -clear with respect to R'' .

(6.15'') The R'' -height of F'' tends to infinity.

Our final task now is to attach this new fan to G_n . Set

$$X := \{v(\beta) \mid \delta'' \leq \beta \leq \nu\} \cap V(F'').$$

Note that $X \neq \emptyset$, as $d'' \in X$. Let $Y \subset X$ consist of the vertices in X of maximal R'' -height. Pick a vertex $b \in Y$ as follows. If Y contains a vertex on R'' , let b be this (unique) vertex. If $Y \cap V(R'') = \emptyset$, then all the vertices of Y lie on a common spoke of F'' or on P_0'' (by (6.13'')); let b be the vertex of Y closest to R'' along this spoke (respectively, along P_0'').

Define a fan E from a'' to b as follows. Let r be the R'' -height of b . As the spokes of E , take those spokes of F'' whose last vertex is above r . If $b \in R''$ (and $r = b$), take bR'' as the spine of E . If $b \notin R''$, say $b \in P_k''$, let the spine of E be the ray $bP_k''rR''$. In either case, E is a fan from a'' to b contained in F'' (Fig 11).

FIGURE 11. The fan $E \subset F''$

Set $\nu_{n+1} := \nu$, and let β_{n+1} be the unique vertex $\beta \leq \nu$ with $v(\beta) = b$. Note that, in fact, $\delta'' \leq \beta < \nu$ by (6.18) and the definition of X , so

$$\alpha'' < \delta'' \leq \beta_{n+1} < \nu_{n+1}.$$

Finally, let

$$G_{n+1} := G_n \cup \beta_{n+1} \cup E.$$

Let us verify the induction hypothesis for $n + 1$. We certainly have $G_n \subset G_{n+1}$ and $\beta_n < \beta_{n+1}$.

For assertion (i), note first that the path $\beta_{n+1} \setminus \beta_n = v(\overset{\circ}{\beta}_n \mathcal{T} \beta_{n+1})$ avoids G_n by condition (iii) for n . The interior of E avoids $G_n \cup \alpha''$ by (6.14'') (since $E \subset F''$), and it avoids the path $\nu_{n+1} \setminus \alpha''$ by the choice of b . (Recall that δ'' is

an upper neighbour of α'' in \mathcal{T} .) In particular, E is a fan on $\beta_{n+1} \setminus \beta_n$ (recall that $\beta_n < \alpha \leq \alpha''$), and (i) follows.

For a proof of assertion (ii), let γ be any burst vertex above β_{n+1} . Then $\gamma \neq \alpha''$, so the R_γ -height of E is bounded by (6.15'') and Lemma 3.8. As the path β_{n+1} is finite, this implies that (ii) carries over from n to $n+1$.

To verify (iii), note that G_n avoids the set

$$V := v(\mathring{\beta}_{n+1} \mathcal{T} \nu_{n+1} \cup \mathcal{T}^{>\nu_{n+1}})$$

by condition (iii) for n , because $\nu_{n+1} > \nu_n$. We saw in the proof of (i) that the interior of E avoids the path $\nu_{n+1} \setminus \alpha''$, so E avoids V by (6.18). The path β_{n+1} avoids V by (6.2). Hence G_{n+1} avoids V .

This completes the induction step in our construction of a fan graph in G , and so completes Case 2. The proof of Theorem 6.1 is thus complete. \square

7. The bounded graph theorem

To conclude the paper, let us restate our principal result and complete its formal proof. Recall that \mathfrak{b} denotes the least cardinal such that there exists a family of \mathfrak{b} distinct $\mathbb{N} \rightarrow \mathbb{N}$ sequences which are not dominated by a common $\mathbb{N} \rightarrow \mathbb{N}$ sequence. (So, $\omega < \mathfrak{b} \leq 2^\omega$, and $\omega_1 = \mathfrak{b} = 2^\omega$ with CH.) Recall further that $I_{\mathfrak{b}}$ denotes the disjoint union of \mathfrak{b} rays, and that T_ω is the tree of countably infinite regular degree. As in the Introduction, let B and F be the prototype bundle graph and the prototype fan graph, respectively (see Figs. 1 and 2).

Theorem 7.1. (Bounded graph theorem)

A graph is bounded if and only if it contains none of the following graphs as a topological subgraph: $I_{\mathfrak{b}}$; T_ω ; B ; F .

Proof. Since boundedness is closed under taking subgraphs, the ‘only if’ part follows from the fact that the four forbidden graphs and their subdivisions are themselves unbounded.

For the ‘if’ part, let us suppose that a graph G contains no subdivision of any of $I_{\mathfrak{b}}$, T_ω , B or F ; we show that G is bounded. Since $G \not\supseteq I_{\mathfrak{b}}$, fewer than \mathfrak{b} of its components contain a ray. By definition of \mathfrak{b} , we may therefore assume that G is connected. Since G does not contain a TK_ω (which would contain a TT_ω), G has a normal spanning tree by Halin’s theorem (3.1). By Theorem 6.1, then, there is no good family of paths in G , so G is bounded by Theorem 5.1. \square

Remark. *If the Continuum Hypothesis is assumed then $\omega_1 = \mathfrak{b} = 2^\omega$, so Theorem 7.1 implies the bounded graph conjecture as stated in the Introduction.*

In the remainder of this section we list a few immediate consequences of the bounded graph theorem. The following three corollaries were already discussed in the Introduction.

Corollary 7.2. *Subdivisions of bounded graphs are bounded, and subdivisions of unbounded graphs are unbounded.* \square

Corollary 7.3. *Let G be a graph. Unless G contains \mathfrak{b} disjoint rays (in which case it is trivially unbounded), G is bounded if and only if all its countable subgraphs are bounded.* \square

Corollary 7.4. *If a countable graph admits some unbounded labelling then all its injective labellings are unbounded.* \square

Our next corollary is a strengthening of Theorem 5.1. Let us call a non-empty family \mathcal{P} of finite paths *propagating* if it satisfies condition (i) of the definition of a good family, i.e., if for every path $P \in \mathcal{P}$ there exists an $n \in \mathbb{N}$ such that P has infinitely many extensions in \mathcal{P} of length n . Clearly, any graph with a propagating family of paths must be unbounded, and any graph containing a subdivision of T_ω , B or F contains a propagating family of paths.

Corollary 7.5. *A graph G not containing $I_{\mathfrak{b}}$ is bounded if and only if there is no propagating family of paths in G .* \square

One may ask whether or not boundedness can also be characterized by the exclusion of minors rather than topological subgraphs. Since the minor relation is transitive, this would require that boundedness be closed under taking minors, i.e. that minors of bounded graphs should again be bounded. With the most general definition of an infinite minor, as given in §2, this is easily seen to be false: the infinite grid, for example, is locally finite and hence bounded, but contracting every second vertical ray yields a graph which contains both a copy of F and a bundle graph (and is therefore unbounded).

However, a slightly less general definition of an infinite minor (which still generalizes the usual finite definition) will work: if we require that all the branch sets be finite—i.e., if all our minors are obtained by first taking a subgraph and then contracting *finite* connected parts of that subgraph—then boundedness does turn out to be closed under taking minors. Indeed, if H is a bounded graph and $H \supseteq H' = HG$ with all branch sets finite, it is not difficult to see that any TB , TF , TT_ω or $I_{\mathfrak{b}}$ subgraph of G could be lifted to a corresponding TB , TF , TT_ω or $I_{\mathfrak{b}}$ subgraph of H' , and hence of H , contradicting the boundedness of H . Therefore G , too, must be bounded:

Corollary 7.6. *Boundedness is closed under taking minors with finite branch sets.* \square

By Lemma 2.1, we thus have the following forbidden minor version of Theorem 7.1:

Theorem 7.7. *A graph is bounded if and only if it contains none of the following graphs as a minor (with all branch sets finite): I_b ; T_ω ; B ; F . \square*

References

- [1] R. Diestel, *Graph Decompositions—a study in infinite graph theory*, Oxford University Press, Oxford 1990.
- [2] R. Diestel, On spanning trees and k -connectedness in infinite graphs, *J. Combin. Theory B* (to appear).
- [3] 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.
- [4] R. Halin, Some problems and results in infinite graphs, in: (L.D. Andersen et al., Eds.) *Graph Theory in Memory of G. A. Dirac* (Annals of Discrete Mathematics **41**), North-Holland Publ. Co., Amsterdam/London 1989.
- [5] R. Halin, Bounded graphs, in: (R. Diestel, Ed.) *Directions in infinite graph theory and combinatorics* (Annals of Discrete Mathematics), to appear.
- [6] H. A. Jung, Wurzelbäume und unendliche Wege in Graphen, *Math. Nachr.* **41** (1969), 1–22.
- [7] D. König, *Theorie der endlichen und unendlichen Graphen*, Akademische Verlagsgesellschaft, Leipzig 1936 (reprinted: Chelsea, New York 1950).
- [8] H. J. Prömel and B. Voigt, *Aspects of Ramsey Theory*, Springer Verlag, in preparation.
- [9] R. Rado, Universal graphs and universal functions, *Acta Arith.* **9** (1964), 331–340.

Reinhard Diestel
Faculty of Mathematics (SFB 343)
Bielefeld University
D-4800 Bielefeld
Germany

Imre Leader
Department of Pure Mathematics
16 Mill Lane
Cambridge CB2 1SB
England