

Dominating functions and topological graph minors

Reinhard Diestel

An infinite graph G is called dominating if its vertices can be labelled with natural numbers in such a way that for every function $f: \omega \rightarrow \omega$ there is a ray in G whose sequence of labels eventually exceeds f . Conversely, G is called bounded if for every labelling of its vertices with natural numbers there exists a function $f: \omega \rightarrow \omega$ which eventually exceeds the labelling along any ray in G . This expository paper describes recent classifications of the dominating and the bounded graphs by forbidden topological minors, and indicates some connections of these results to infinite games.

Introduction

If f and g are functions from ω to ω , we say that f *dominates* g if $f(n) \geq g(n)$ for all but finitely many $n \in \omega$. A family \mathcal{F} of $\omega \rightarrow \omega$ functions is called a *dominating family* if every function $g: \omega \rightarrow \omega$ is dominated by some $f \in \mathcal{F}$. The least cardinality of a dominating family is denoted by \mathfrak{d} .

Similarly, a family \mathcal{F} of functions from ω to ω is said to be *bounded* by a function $g: \omega \rightarrow \omega$ if g dominates every $f \in \mathcal{F}$; if no such g exists, \mathcal{F} is called *unbounded*. The least cardinality of an unbounded family is denoted by \mathfrak{b} .

It is not difficult to see that any unbounded family of functions must be uncountable. Indeed, if $\mathcal{F} = \{f_n \mid n \in \omega\}$ is a countable family of $\omega \rightarrow \omega$ functions, then $g: \omega \rightarrow \omega$ defined by $g(n) := \max\{f_0(n), \dots, f_n(n)\}$ is a bounding function for \mathcal{F} . Thus, $\mathfrak{b} > \omega$. Since any dominating family is clearly unbounded, we further have $\mathfrak{b} \leq \mathfrak{d}$. Finally, the set of *all* $\omega \rightarrow \omega$ functions clearly defines a dominating family; since there are exactly continuum many $\omega \rightarrow \omega$ functions, we have

$$\omega < \mathfrak{b} \leq \mathfrak{d} \leq 2^\omega.$$

Set theorists have traditionally been interested in the question of when the above inequalities may be strict, and how \mathfrak{b} and \mathfrak{d} compare with other cardinals between ω and 2^ω . In other words, it has been asked just how much the cardinality of a family of functions has to be constrained in order to force it to become bounded or to cease to be dominating. This question depends on the axioms of set theory assumed, and we shall not pursue it further here. Instead, we shall ask how we can force a family of functions to become bounded (or to cease to be dominating) by restricting the independence of its members.

This can be done naturally as follows. We shall label the vertices of an infinite graph with integers, and consider as our family of functions the labellings

along one-way infinite paths, or *rays*, in this graph. Depending on how much different rays in the graph intersect, it may turn out that their labellings can never form a dominating or an unbounded family, even when the graph contains continuum many distinct rays.

To be precise, let us say that a graph G is *dominating* if there exists a labelling $\ell: V(G) \rightarrow \omega$ of its vertices such that for every $f: \omega \rightarrow \omega$ there is a ray $v_0v_1 \dots$ in G with $\ell(v_n) \geq f(n)$ for all but finitely many n . Similarly, G is called *bounded* if for every labelling $\ell: V(G) \rightarrow \omega$ there exists a function $f: \omega \rightarrow \omega$ such that, whenever $v_0v_1 \dots$ is a ray in G , we have $\ell(v_n) \leq f(n)$ for all but finitely many n . If G is a graph with a fixed labelling, we shall not always distinguish between a ray in G and its sequence of labels, so that we may speak of functions dominating rays and vice versa.

Note that, by definition, supergraphs of dominating graphs are again dominating, and subgraphs of bounded graphs are again bounded. We may therefore hope to classify the bounded and the dominating graphs by identifying some particular ‘minimal’ prototypes of unbounded or dominating graphs, and showing that a graph is unbounded or dominating if and only if it contains one of the respective prototypes. In fact, we shall see that a graph is bounded if and only if it contains none of four specified graphs as a topological minor. The dominating graphs will be characterized similarly.

The prototype unbounded or dominating graphs appearing in these characterizations will be discussed in Section 1, together with some other examples. In Section 2 we present the two classification theorems, and take a glance at how they are proved. Section 3 introduces a framework for the definition of infinite games related to the domination of functions arising from rays in labelled graphs; two specific games, the dominating game and the bounding game, are analysed and winning strategies given.

The notation used will be standard; see e.g. [1]. When a graph G contains a subdivision of a graph H as a subgraph, we also say that H is a *topological minor* of G . If P is a path containing vertices u and v , we write uP , Pv and uPv for the obvious subpaths of P starting in u and/or ending in v .

1. Examples

In this section we collect together a few typical examples of graphs that are, or fail to be, bounded or dominating. Note that any dominating graph is automatically unbounded.

We have already seen that any unbounded family of functions must be uncountable. If a graph contains only countably many rays, it is therefore trivially bounded.

For a more interesting example of bounded graphs, consider any connected locally finite graph G , and pick a labelling. For each vertex $v \in G$, we may

define a function $f_v: \omega \rightarrow \omega$ by setting $f_v(n)$ to be the maximum of all the labels of vertices at distance at most n from v . Clearly, f_v dominates (the labelling along) every ray in G that starts at v . Since G has only countably many vertices, there exists a function f which dominates every f_v with $v \in V(G)$, and hence every ray in G . Therefore G is bounded.

On the other hand, it is not difficult to find graphs that are unbounded or even dominating. For example, the union of \mathfrak{d} disjoint rays is dominating, and the union of \mathfrak{b} disjoint rays is unbounded: just label each ray by a different member of some dominating or unbounded family of functions.

So how about countable unbounded graphs? The complete graph K_ω on a countably infinite set of vertices is clearly dominating. Indeed, consider any labelling ℓ that uses arbitrarily large labels: then, for any $f: \omega \rightarrow \omega$, we may find a ray $v_0v_1\dots$ with $\ell(v_n) > f(n)$ for every n . Similarly, the ω -regular tree T_ω is dominating, and hence unbounded: just label its vertices injectively, i.e. so that any two labels are different.

In fact, any subdivision T of T_ω is unbounded. To see this, let again $\ell: V(T) \rightarrow \omega$ be any injective labelling, and let $f: \omega \rightarrow \omega$ be any given function. We may then choose a ray $v_0v_1\dots$ in T as follows. If $n = 0$ or v_{n-1} is a branch vertex of T , choose v_n so that $\ell(v_n) > f(n)$; this can be done, because ℓ is injective. Otherwise let v_n be any hitherto unused neighbour of v_{n-1} . (This vertex is unique unless $n = 1$.) Since any ray in T contains infinitely many branch vertices, we see that $\ell(v_n) > f(n)$ for infinitely many n , and hence that T is unbounded.

It is an interesting fact that, by contrast, subdivisions of T_ω need not be dominating. Indeed, consider any enumeration $e: \omega \rightarrow E(T_\omega)$ of the edges of T_ω , and let T be the tree obtained from T_ω by subdividing $e(n)$ exactly n times for each n .

Proposition 1.1. *T is not dominating.*

Proof. Let $\ell: V(T) \rightarrow \omega$ be any labelling of the vertices of T . Let $f: \omega \rightarrow \omega$ be an increasing function satisfying $f(n) > \max \{ \ell(v) \mid v \in e(n) \}$ for all $n \in \omega$. We show that, for any ray $R = v_0v_1\dots$ in T and any $i \in \omega$, there exists a $k > i$ with $f(k) > \ell(v_k)$.

Given such R and i , choose $j, k \in \omega$ with $i < j < k$ so that $v_jv_k = e(n)$ for some n . Then R traces out the subdivided edge $e(n)$, so in particular we have $k \geq |v_jRv_k| = n + 2$. Since f is increasing and $f(n) > \ell(v_k)$ by definition of f , this gives $f(k) \geq f(n) > \ell(v_k)$ as desired. \square

If, as we have seen, T_ω is dominating but subdivisions of it need not be, can we say exactly which kinds of subdivision of T_ω yield a dominating graph? Indeed we can: as Theorem 2.3 of Section 2 will show, a subdivision of T_ω is dominating if and only if it contains a particularly simple subdivision of T_ω , to

be called ‘uniform’, in which for every branch vertex there exists a bound on how often the edges incident with this vertex in T_ω have been subdivided.

More precisely, let us call a subdivision T of T_ω *uniform* if it has a branch vertex r , called its *root*, such that whenever v is a branch vertex, all the subdivided edges at v that are *not* contained in the unique path from v to r have the same length.

Proposition 1.2. *Uniform subdivisions of T_ω are dominating.*

Proof. Let T be a uniform subdivision of T_ω , with root r say. Let $\ell: V(T) \rightarrow \omega$ be any injective labelling; we show that for every function $f: \omega \rightarrow \omega$ there is a ray R in T which dominates f .

Since any ray in T is a concatenation of paths that are subdivided edges of T_ω , we may define $R = v_0v_1\dots$ inductively by choosing these subdivided edges one at a time. Let $v_0 = r$. Suppose that v_n has been defined for every $n \leq m$, and that v_m is a branch vertex of T . Then all the (infinitely many) subdivided edges at v_m that are not contained in the portion of R defined so far have the same length k , and so we can find one of them, say $s_0\dots s_k$ where $s_0 = v_m$, such that $\ell(s_i) \geq f(m+i)$ whenever $0 < i \leq k$. Setting $v_{m+i} = s_i$, we obtain $\ell(v_{m+i}) \geq f(m+i)$ for all these i ; moreover, v_{m+k} is again a branch vertex of T . This completes the induction step, and hence the construction of R . Since $\ell(v_n) \geq f(n)$ for every $n > 0$, it is clear that R dominates f . \square

Surprisingly, there is yet another way to obtain a dominating graph from a subdivision of T_ω : by taking \mathfrak{b} disjoint copies of it. Note that \mathfrak{d} disjoint copies are trivially dominating, since each of them contains a ray; it can be shown, however, that \mathfrak{b} subdivisions of T_ω , which need not be isomorphic, suffice:

Proposition 1.3. [5] *If a graph is the union of \mathfrak{b} disjoint subdivisions of T_ω , then it is dominating.* \square

We have now seen all the examples of dominating graphs that will be needed for our classification theorem in Section 2: Theorem 2.3 says that a graph is dominating if and only if it contains a uniform subdivision of T_ω , \mathfrak{b} disjoint subdivisions of T_ω , or \mathfrak{d} disjoint rays.

For the classification of bounded graphs, however, there are two more prototypical unbounded graphs that may occur. One of these is the graph B obtained from a ray $v_0v_1\dots$ by adding, for each $n \in \omega$, a countably infinite set of independent $v_{3n+1}v_{3n+3}$ paths of length 2 (Fig. 1).

To show that B is unbounded, let ℓ be any injective labelling of $V(B)$, and assume we are given a function $f: \omega \rightarrow \omega$. Let us specify a ray $R = u_0u_1\dots$ as follows. Put $u_0 := v_0$ and $u_1 := v_1$. For u_2 , we have an infinite choice of neighbours of v_1 , all labelled differently. We may thus choose as u_2 a neighbour of v_1 whose label is greater than $f(2)$. We continue with $u_3 := v_3$ and $u_4 := v_4$, where again we have an infinite choice for u_5 . Proceeding in this manner, we




FIGURE 1. The prototype bundle graph B

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

In a similar way one can show that every subdivision of B is unbounded; any subdivision of B will be called a *bundle graph*.

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 $n \in \omega$ (Fig. 2).




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 and any function $f: \omega \rightarrow \omega$, we can easily find a ray R through F whose labels exceed the corresponding values of f 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.)

In a similar way one can show that every subdivision of F is unbounded; any subdivision of F will be called a *fan graph*.

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 [6]; the reader may find it amusing to prove its boundedness without using Theorem 2.1 below.

2. The classification theorems

In this section we state the results from [3] and [5] which characterize the bounded and the dominating graphs by forbidden configurations, and give some indications of how these theorems are proved. The full proof of the bounded graph theorem is too complex to be sketched in detail, but we shall give an outline of the main ideas involved. For the dominating graph theorem, we shall be able to give a fairly accurate sketch of what may be the most typical case.

Let us begin with the bounded graph theorem, which was proved in [3]. The result was conjectured by Halin almost 30 years ago, but only few partial results used to be known. Recall that T_ω , B and F denote the ω -regular tree, the prototype bundle graph and the prototype fan graph, respectively, and let $I_{\mathfrak{b}}$ denote the union of \mathfrak{b} disjoint rays.

Theorem 2.1. (Bounded graph theorem)

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

The bounded graph theorem has some fundamental implications for the concept of boundedness for graphs. Let us mention two of these. First, the theorem implies that the boundedness or unboundedness of a graph depends only on its countable subgraphs: unless the graph contains \mathfrak{b} disjoint rays—in which case it is trivially unbounded—it is bounded if and only if all its countable subgraphs are bounded.

Secondly, the translation of the original problem of bounding $\omega \rightarrow \omega$ functions to a problem on graphs has been successful, in that the unboundedness of a graph is shown to be a truly structural graph property, not one of the existence of particularly intricate labellings: if a (countable) graph is unbounded then, by the bounded graph theorem, this is witnessed by any injective labelling.

As for the proof of the bounded graph theorem, we have already seen the forward implication: the graphs $I_{\mathfrak{b}}$, T_ω , B and F and their subdivisions are unbounded, and so they cannot be subgraphs of a bounded graph. For the converse implication, let us first get some intuition for why these four types of subgraph might be natural ones to occur inside an arbitrary unbounded graph G .

If G has \mathfrak{b} or more components each containing a ray, we have $G \supseteq I_{\mathfrak{b}}$. If not, it suffices to show that every such component is bounded: by definition of \mathfrak{b} , there will then be a bounding function for all of G .

We may thus assume that G is connected. As we saw earlier, locally finite connected graphs are bounded, so G will have vertices of infinite degree. Let v be such a vertex, and assume that G contains infinitely many rays starting at v . (This is not an enormous assumption, given that G contains enough rays to make it unbounded.) Could it be that these rays can be chosen independent, i.e. so that any two of them meet only in v ? If so, their union might be viewed

as the beginning of a subdivided T_ω : unless v is an atypical vertex of G (in which case the unboundedness of G would not depend on it), there should be another vertex of infinite degree on one of the rays and a similar set of independent rays issuing from it. If this process continues for long enough, it is not unreasonable to assume that the union of all the rays involved contains a subdivision of a graph in which *every* vertex has infinite degree—a graph which is easily shown to contain a copy of T_ω .

Suppose now that there is no infinite set of independent rays starting at v . Then, by Ramsey's theorem, there is an infinite set of rays from v such that every two of them meet also in some other vertex. Now if v is the only vertex that lies on infinitely many of these rays, then it is easy to find a 'fan' in their union. Indeed, let R be one of the rays, and notice that for every finite set of vertices we may find another ray from v avoiding this set. We may thus inductively select an infinite sequence of initial segments of further rays from v , which avoid R except for their two endvertices, and which meet pairwise only in v .

We may therefore assume that there exists a vertex w which lies on infinitely many rays from v . We thus have an infinite set of v - w paths in G whose second vertices are distinct. Using König's infinity lemma, it is not difficult to select an infinite subset of these paths forming either a 'bundle' or a (backward) 'fan' (Fig. 3).




FIGURE 3. Infinitely many v - w paths forming a bundle or a fan

The main purpose of the above considerations was to see that fans and bundles, as well as disjoint or infinitely branching rays, 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 theorem. 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. The difficulty in proving the theorem lies in the task of finding some structure in the graph that enables us to construct an infinite sequence of disjoint bundles or fans, as tidily threaded on a ray as in the graphs B and F .

To get to the heart of the proof, let us recall the proofs of the unboundedness of bundle graphs, fan graphs and subdivisions of T_ω , and see what these

proofs have in common.

In the case of a subdivision T of T_ω , we just took any injective labelling of T and observed that, for any given function $f: \omega \rightarrow \omega$, we could easily find a ray R through T which was not dominated by f . Indeed, no matter how we had chosen an initial segment of R , we would be able to get to another branch vertex of T , where we would have an infinite choice of labels for the next vertex of R . This label could thus be chosen larger than the corresponding term in f .

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 G is central to the proof of the bounded graph theorem, so let us give it a precise form. Let us call a family \mathcal{P} of finite paths in G a *good family* if

*for every path $P \in \mathcal{P}$ there exists an $n \in \omega$ such that P has
infinitely many extensions in \mathcal{P} of length n .* (*)

Note that the existence of a good family of paths in G implies that G is unbounded. Indeed, if in condition (*) we choose n minimal, then there exist infinitely many extensions $P' \in \mathcal{P}$ of P that 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 function f (as in our discussion of B and F), 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 f . In this manner we may construct R inductively as the limit of a nested sequence of paths from \mathcal{P} .

Moreover, if G contains a bundle graph, a fan graph or a subdivision of T_ω , then clearly G contains a good family of paths. Thus *if* the bounded graph theorem is true, then any graph G not containing I_b satisfies the following implications (which in turn imply the bounded graph theorem):

$$G \text{ is unbounded} \quad \Rightarrow \quad G \text{ contains a good family of paths} \quad \Rightarrow \quad G \text{ contains } T_\omega, B \text{ or } F \text{ as a topological minor}$$

It is therefore reasonable to expect that the proof of the bounded graph theorem could be carved up into two chunks, verifying these two implications separately. This is indeed the basic structure of the proof as given in [3]. In practice it turns out to be convenient for the proof of the second implication to strengthen the definition of a good family considerably, but for simplicity we shall here work with the definition given above. The only visible effect of this simplification will be that the first implication will appear to be easier to prove for graphs of cardinality $< \mathfrak{b}$ than for arbitrary graphs. This is not actually

the case, i.e. the ideas outlined below for the ‘case’ of $|G| \geq \mathfrak{b}$ would in fact be essential ingredients even of a proof for countable graphs only.

In order to prove the first implication, let us try to find a good family of paths in G by induction: starting with the family \mathcal{P}_0 of *all* finite paths in G , let us recursively discard any paths P from this family that violate the defining condition (*) for a good family, and hope that eventually we will be left with a family that is indeed good. More precisely, let us define subfamilies \mathcal{P}_α of \mathcal{P}_0 , for all ordinals $\alpha > 0$, 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.e., a path P such that, for each $n \in \omega$, \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 $\mathcal{P}_\alpha := \bigcap_{\beta < \alpha} \mathcal{P}_\beta$.

Clearly, there exists an α of cardinality at most $|\mathcal{P}_0|$ such that $\mathcal{P}_{\alpha+1} = \mathcal{P}_\alpha$; let α^* be the least such α , and set $\mathcal{P}_{\alpha^*} =: \mathcal{P}$. Clearly \mathcal{P} satisfies condition (*), and is therefore a good family if and only if it is non-empty. To complete the proof of the first implication, we thus have to show that $\mathcal{P} \neq \emptyset$.

We now prove that $\mathcal{P} \neq \emptyset$, *under the additional assumption that $\alpha^* < \mathfrak{b}$* . The problem of how \mathcal{P}_0 , and hence α^* , might be reduced to cardinality $< \mathfrak{b}$ without the risk of ending up with $\mathcal{P} = \emptyset$, will be addressed afterwards.

Suppose $\mathcal{P} = \emptyset$; we show that G is bounded, contrary to our assumption. Let $\ell: V(G) \rightarrow \omega$ be an arbitrary labelling of G . We shall define functions $f_\alpha: \omega \rightarrow \omega$, one for each $\alpha < \alpha^*$, so that any function f^* that dominates every f_α will bound G ; such a function f^* exists by our assumption that $\alpha^* < \mathfrak{b}$.

For each $\alpha < \alpha^*$, let the k th term of f_α be defined as

$$f_\alpha(k) := \begin{cases} \ell(v_k) & \text{if } k \leq |P_\alpha| \\ \max\{\ell(t(P)) \mid P \in \mathcal{P}_\alpha \setminus \mathcal{P}_{\alpha+1} \text{ and } |P| = k\} & \text{if } k > |P_\alpha|, \end{cases}$$

where v_k is the k th vertex of P_α and $t(P)$ 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.

To show that f^* bounds G , let $R \subset G$ be any ray. Since every initial segment of R is in \mathcal{P}_0 , but $\mathcal{P} = \emptyset$ 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, $\mathcal{P}_\alpha \setminus \mathcal{P}_{\alpha+1}$ contains every initial segment of R of length $k \geq |P_\alpha|$. Therefore f_α dominates R , by definition of f_α , and so f^* dominates R as claimed.

To complete the proof of our ‘first implication’ (i.e. that any unbounded $G \not\geq I_{\mathfrak{b}}$ contains a good family of paths), let us now see how we can replace \mathcal{P}_0

in the above recursion with a smaller starting family. This family should still be large enough to contain a good family (provided that G is unbounded), but small enough to allow the recursion to be completed in fewer than \mathfrak{b} steps.

To achieve this, we make use of a structure theorem for connected graphs not containing a topological K_ω minor; note that we may assume this for G , as otherwise G contains a subdivision of T_ω (and, in particular, a good family of paths). Let $T \subset G$ be a tree, with root r say, and let us call a ray $R \subset T$ *normal* if it starts at r . The tree T will be called a *skeleton* of G if, for each normal ray R , we can assign to every vertex of G a vertex in R , called its *R -height*, so that the following conditions are satisfied:

- (i) if t is a vertex on T , then the R -height of t is the vertex of R closest to t in T ;
- (ii) if v is a vertex of $G \setminus T$ and v is joined to a vertex $t \in T \setminus R$ by a path whose interior avoids T , then the R -height of v equals the R -height of t ;
- (iii) if a vertex v has R -height x and a vertex w has R -height different from x , then the path rTx separates v from w in G .

Thus, if T is a skeleton of G and $R \subset T$ is a normal ray, and if C is a component of $G \setminus R$ containing the branch B of T , say, then every vertex of C has the same R -height, namely the unique vertex of R to which B is attached in T .

Our structure theorem for graphs without a topological K_ω minor, which builds on work of Jung and Halin and is proved using simplicial decompositions of graphs (see [2]), can now be expressed as follows:

Theorem 2.2. *If G is connected and contains neither $I_{\mathfrak{b}}$ nor a subdivision of K_ω as a subgraph, then G has a skeleton T of order $< \mathfrak{b}$ such that every ray in G meets some ray in T infinitely often.*

Intuitively, Theorem 2.2 says that G contains a ‘small’ skeleton, and that the rest of G is ‘wrapped around’ this skeleton in a bounded sort of way. The reason why this will be useful to us is that, roughly speaking, it allows us to distinguish two cases. Either G is wrapped around T tightly, in which case there is a good family of paths all staying close to the skeleton; since the skeleton itself has order $< \mathfrak{b}$, these paths can be extracted recursively from a family \mathcal{P}_0 of order $< \mathfrak{b}$, as desired. Alternatively, the wrapping around T will be bulky in many places of different R -height (for some normal ray R). We may then extract a fan or a bundle at every such place, and combine them into a fan graph or a bundle graph: as the R -heights of different fans or bundles are distinct, they will be pairwise disjoint. (Recall that keeping fans or bundles disjoint is one of the main problems in the entire proof.)

Let us make these ideas more precise. When s and t are comparable vertices of T , i.e. if (say) s lies on the unique r - t path in T , let us say that there is *thick padding* around T at the pair (s, t) if G contains infinitely many (not necessarily independent) s - t paths whose interiors avoid T ; otherwise we

shall speak of *thin padding* at (s, t) . Note that the s - t paths in the padding at (s, t) need not be linked up with each other in $G \setminus T$; in particular, their vertices may have different heights with respect to a normal ray R .

Recall from our introductory discussion that any union of infinitely many s - t paths with distinct second vertices contains a bundle or a fan (Fig. 3). Using similar arguments, it is not difficult to show that if T has thick padding around it at (s, t) , then somewhere ‘inside this padding’ (i.e., using no vertices of T other than s and t), there must be a bundle; note that there cannot be a fan there since, by Theorem 2.2, $G \setminus T$ contains no ray.

We now come to the distinction between the two cases of whether G is ‘wrapped around T tightly’ or not. Suppose first that, for every normal ray R , the padding around R becomes thin eventually, i.e. there are only finitely many non-overlapping pairs of vertices of R at which there is thick padding. We may then restrict the starting family \mathcal{P}_0 for our recursion to paths which leave T only at pairs where the padding is thin; in other words, whenever $P \in \mathcal{P}_0$ is such that sPt is a path of length ≥ 2 and $sPt \cap T = \{s, t\}$, the set of *all* such s - t paths in G is finite. Each path in \mathcal{P}_0 is then determined by its sequence of vertices on T and the finite choices of connecting paths between these vertices, so we have $|\mathcal{P}_0| \leq |T| < \mathfrak{b}$ as desired.

We may therefore assume that T contains a normal ray R with thick padding at infinitely many non-overlapping pairs of its vertices. Let us select an infinite sequence of such pairs and extract a bundle from the padding at each of these pairs. Let us call such a bundle *wide* if the set of R -heights of its vertices is finite; otherwise let it be *tall*. Then there is an infinite subsequence of our pairs such that either each one of the associated bundles is wide, or each of these bundles is tall.

If every bundle is wide, we may select a further subsequence of pairs such that the R -heights of bundles associated with different pairs are disjoint; with a bit of work, these bundles may then be combined into a bundle graph. Suppose finally that every bundle in the sequence is tall. It is not difficult to show that the endvertices of such a bundle (i.e. its two vertices of infinite degree) must be on R . Thus, R has infinitely many vertices—two for each bundle—every one of which has neighbours of arbitrarily large R -height. Choosing appropriate disjoint connecting paths between these neighbours and their respective R -heights on R , we obtain a graph that is easily seen to contain a subdivision of K_ω . This contradicts our assumptions about G , completing the proof of the ‘first implication’.

The proof of the second implication, the fact that if G contains a good family of paths then it contains a bundle graph, a fan graph or a subdivision of T_ω , makes up about two thirds of the proof of the bounded graph theorem. The techniques used are largely similar to those outlined above. It is assumed that G is connected and has no topological K_ω minor, and therefore contains a skeleton. The skeleton may then be used to keep different fans and bundles

separated by their R -heights, for suitable normal rays R ; thus, if sufficiently many fans or bundles can be found, these may be combined to a bundle graph or a fan graph.

Let us then mention briefly how these fans or bundles are generated. Let \mathcal{P} be a good family of paths in G . The essential structure of \mathcal{P} can be represented by a graph \mathcal{T} , which looks very similar to a T_ω , and which acts as something like a covering tree for the paths $P \in \mathcal{P}$ in G . If G resembles \mathcal{T} closely, it can be shown to contain a subdivision of T_ω . On the other hand, if the paths of \mathcal{P} intersect a lot more (in G) than do their lifts in \mathcal{T} , they generate sufficiently many bundles or fans to make a bundle graph or a fan graph in G . The distinction between the latter two cases depends on the relationship between \mathcal{T} and the skeleton T of G .

To see how \mathcal{T} is obtained, let us first select a simple subfamily of \mathcal{P} , which will still be a good family. We start with any path $P_0 \in \mathcal{P}$. By condition (*), P_0 has infinitely many extensions in \mathcal{P} of some common length n . Among these, select infinitely many that agree up to their penultimate vertices; this can be done if n is first chosen minimal. In the same way, we then find an infinite set of extensions for each of these new paths, again (in each case) of some common length and agreeing up to their penultimate vertices. In ω steps, we have constructed a subfamily \mathcal{Q} of \mathcal{P} which is still good, and which clearly has a structure similar to a subdivision of T_ω : let \mathcal{Q}' be the closure of \mathcal{Q} under taking initial segments, set $V(\mathcal{T}) = \mathcal{Q}'$, and join vertices $P, P' \in \mathcal{Q}'$ by an edge whenever P' is an extension of P by one vertex. The tree \mathcal{T} obtained in this way is a subdivision of T_ω , except for its 'long root' made up of vertices that are proper initial segments of P_0 .

Let $\pi: V(\mathcal{T}) \rightarrow V(G)$ be the map that assigns to each $P \in V(\mathcal{T})$ its last vertex. Then if p is the first vertex of P_0 , and the trivial path $\{p\} \in V(\mathcal{T})$ is taken to be the root of \mathcal{T} , we see that π maps each path of the form $\{p\}\mathcal{I}P$ to the path $P \subset G$ in a natural way. Conversely, each path $P = v_0 \dots v_n \in \mathcal{Q}'$ lifts uniquely to the path in \mathcal{T} with vertices $Pv_i, i = 0, \dots, n$.

Now let P be a branch vertex of \mathcal{T} , with $\pi(P) = v$ say, and consider the subdivided edges at P pointing away from the root $\{p\}$. If these subdivided edges project under π to paths in G that are disjoint except for their common initial vertex v , then these paths might be used in the construction of a subdivided T_ω in G .

On the other hand, it might be the case that G has a vertex w such that every subdivided edge of \mathcal{T} at P contains a vertex Q with $\pi(Q) = w$. In this case, the subdivided edges at P have initial segments which project to an infinite set of v - w paths in G . As we saw earlier, the union of such a set of paths makes up a bundle or a fan.

Using an easy Ramsey type argument, we may now assume that all the branch vertices of \mathcal{T} behave in the same way; then either we obtain a subdivision of T_ω in G straight away, or else we obtain a bundle or a fan for each branch vertex of \mathcal{T} . Using a number of similar Ramsey type arguments, this

infinite set of configurations in G (fans or bundles) can then be streamlined in several ways. Eventually we end up with a set of configurations that are not only either all fans or all bundles, but which are compatible (with respect to their position towards the skeleton T) in many other ways too. (In particular, these fans or bundles will be pairwise disjoint.) In this way we eventually obtain a fan graph or a bundle graph in G , completing the proof of the bounded graph theorem. \square

We now turn to the dominating graph theorem, proved in [5]. Its proof is much shorter than that of the bounded graph theorem, and we shall be able to give a fairly complete sketch of the most typical case.

Theorem 2.3. (Dominating graph theorem)

A graph G is dominating if and only if it satisfies one of the following three conditions:

- (i) G contains a uniform subdivision of T_ω ;
- (ii) G contains \mathfrak{b} disjoint subdivisions of T_ω ;
- (iii) G contains \mathfrak{d} disjoint rays.

We prove the theorem under the set theoretic assumption that $\mathfrak{b} = \mathfrak{d}$. This is a comparatively weak assumption, much weaker, say, than Martin's axiom or even the continuum hypothesis. Note that case (ii) of the theorem is now redundant, since \mathfrak{d} subdivisions of T_ω contain \mathfrak{d} disjoint rays.

We have already seen that uniform subdivisions of T_ω and unions of \mathfrak{d} disjoint rays are dominating; it remains to prove that if G is dominating then G contains one of these two types of subgraph.

The basic idea of the proof is similar to the way in which we obtained a good family of paths for the bounded graph theorem. We recursively define a rank function ρ on some or all of the vertices of G , with the following property. If any vertex remains unranked, i.e. if the recursion ends before ρ is defined on all of $V(G)$, then G contains a uniform subdivision of T_ω ; if ρ gets defined for every vertex, then either G contains \mathfrak{d} disjoint rays or G is not dominating.

Set $\rho(v) = 0$ for all vertices v that have finite degree in G . Now let $\alpha > 0$ be given, and suppose that for every $\beta < \alpha$ we have assigned to some vertices the ρ -value β . Set $\rho(v) = \alpha$ for all those vertices v for which

- $\rho(v)$ is still undefined, and
- whenever \mathcal{P} is a set of finite paths from v , pairwise disjoint except for v , all of the same length, and each ending in a vertex which has not been given any ρ -value $\beta < \alpha$, then \mathcal{P} is finite.

If there is no such vertex v , we terminate the recursive definition of ρ , and leave ρ undefined for any remaining vertices of G .

It is not difficult to see that if ρ remains undefined for some vertices $v \in G$, then G contains a uniform subdivision of T_ω . Indeed, the definition of ρ implies

that for every such v there exists an infinite set of paths from v , pairwise disjoint except for v , all of the same length, and each ending in a vertex for which ρ is also undefined. It is now easy to use all these paths as subdivided edges to build a uniform subdivision of T_ω , choosing them inductively in ω steps and so that the portion of the tree constructed remains connected at all times: then at each point of the construction only finitely many vertices have been used, but there is always an infinite set of disjoint paths from which the next subdivided edge can be chosen.

Let us therefore assume that $\rho(v)$ is defined for all $v \in V(G)$, and that G contains no union of \mathfrak{d} disjoint rays; we shall show that G is not dominating. Let $\ell: V(G) \rightarrow \omega$ be any labelling; we now have to find a function $f: \omega \rightarrow \omega$ which is not dominated by any ray in G .

Let a path P from u to w in G be called *upward* if

$$\rho(w) = \max \{ \rho(v) : v \in P \}.$$

We first show the following.

(2.3.1) *For each $u \in V(G)$ and each integer m , there are only finitely many vertices w such that G contains an upward path of length m from u to w .*

Suppose the contrary, and consider a vertex u , an integer m , and an infinite set $\{w_n : n \in \omega\}$ such that for each n there is an upward path P_n of length m from u to w_n . Choose $k \leq m$ maximal so that there exist a vertex v and an infinite set $\mathcal{P} \subset \{P_n : n \in \omega\}$ such that v is the k th vertex in every $P \in \mathcal{P}$. (Note that k exists, because every P_n starts at u .) We shall now select an infinite sequence $\{P_{n_i} : i \in \omega\}$ of paths from \mathcal{P} so that any two of these paths are disjoint after v ; since each P_n is an upward path, and hence $\rho(v) \leq \rho(w_n)$ for every n , this will contradict the definition of $\rho(v)$.

Let P_{n_0} be any path from \mathcal{P} . Now suppose P_{n_0}, \dots, P_{n_i} have been chosen, and let U be the union of their vertex sets. By the maximality of k , there are at most finitely many paths in \mathcal{P} that contain a vertex from U after v ; let $P_{n_{i+1}}$ be any other path from \mathcal{P} . It is then clear that the full sequence $\{P_{n_i} : i \in \omega\}$ has the required disjointness property. This completes the proof of (2.3.1).

We now define the function f which will show that G is not dominating. Let U be the vertex set of the union of some maximal set of disjoint rays in G . By assumption there are fewer than \mathfrak{d} rays in this set, so $|U| < \mathfrak{d}$. Note also that every ray in G meets U in infinitely many vertices. Using (2.3.1) we may define, for each $u \in U$, a function $f_u: \omega \rightarrow \omega$ such that $f_u(m) > \ell(w)$ for every $m \in \omega$ and every vertex w to which u can be linked by an upward path of length m . By our hypothesis that $\mathfrak{b} = \mathfrak{d} > |U|$, there exists an $\omega \rightarrow \omega$ function which dominates each of the functions f_u ; let f be such a function. Redefining $f(n)$ as $\max \{ f(k) : k \leq n \}$ if necessary, we may assume that f is increasing.

Now let $R = v_0 v_1 \dots$ be any ray in G ; we have to show that $f(k) > \ell(v_k)$ for infinitely many $k \in \omega$. Using the fact that the ordinal sequence $\rho(v_0), \rho(v_1), \dots$ cannot contain an infinite decreasing sequence, it is not difficult to see that we can find an infinite increasing sequence $\{k_i : i \in \omega\}$ such that $\rho(v_{k_i}) \leq \rho(v_{k_{i+1}})$ for each i , and $\rho(v_j) < \rho(v_{k_i})$ whenever $k_i < j < k_{i+1}$. Now pick $k^* \geq k_0$ so that $u := v_{k^*} \in U$. Note that, for each $k_i > k^*$, the path $u R v_{k_i}$ is an upward path of length $k_i - k^*$.

Since f dominates f_u , there is some $K \in \omega$ such that $f_u(k) \leq f(k)$ for all $k \geq K$. But then

$$\ell(v_{k_i}) < f_u(k_i - k^*) \leq f(k_i - k^*) \leq f(k_i)$$

for all i with $k_i - k^* \geq K$, by definition of $f_u = f_{v_{k^*}}$. Thus R fails to dominate f , as required. \square

3. Domination games

As is well known, Adam and Eve used to play the following game. For a given graph G , Adam first chooses a labelling $\ell: V(G) \rightarrow \omega$. Then the two players move alternately: Eve, who moves first, plays a natural number, Adam a vertex of G . In this way, Eve creates a function $e: \omega \rightarrow \omega$, while Adam creates an ω -sequence of vertices.

In the *bounding game*, Adam tries to escape domination by Eve: he wins if and only if he succeeds in constructing a ray $A \subset G$ that is not dominated by Eve's function e . Thus, playing Adam's part in the bounding game is like an interactive version of trying to prove that G is unbounded. (The difference is that now Adam does not know what Eve will play in future moves; for a proof that G is unbounded, it would be sufficient to be able to construct an undominated ray with respect to any $\omega \rightarrow \omega$ function given complete at the start.)

It is clear that if G is bounded, then Eve has a winning strategy: once Adam has chosen his labelling of G , she just plays a bounding function with respect to this labelling, without ever paying attention to what Adam is doing. On the other hand, Eve may still have a winning strategy when G is unbounded. For example, Eve has a winning strategy for the graph $I_{\mathfrak{b}}$: she plays any number in her first move, waits to see from which component of $I_{\mathfrak{b}}$ Adam picks his first vertex, and then plays a winning strategy for that component. (Recall that each component of $I_{\mathfrak{b}}$ is a ray, and is therefore bounded.) Since Adam can only construct a ray if he stays in that component, Eve will win the game.

The following result from [4] shows that the example of $I_{\mathfrak{b}}$ does indeed mark the difference between the graphs that are unbounded and those for which

Adam has a winning strategy in the bounding game. Its proof in [4] gives an explicit winning strategy for either Adam or Eve, as appropriate.

Theorem 3.1. *Adam has a winning strategy in the bounding game on G if and only if G contains one of the graphs T_ω , B and F as a topological minor.*

The *dominating game* is defined like the bounding game, except that now Adam tries to construct his ray A in such a way that it actually dominates the function e created by Eve (and is not just not dominated by it). Again, it is clear that if Adam has a winning strategy in the dominating game on G then G must be a dominating graph. However, unlike the similarity between the bounding game and the bounded graph theorem, it turns out to be much harder for Adam to win the dominating game on G than it is to prove that G is a dominating graph: the following result from [4] implies that the dominating game can be won by Eve on ‘most’ subdivisions of T_ω , including uniform ones, as well as on disjoint unions of these (cf. Theorem 2.3).

Theorem 3.2. *Adam has a winning strategy on a graph G if and only if $G \supseteq T_\omega$. Otherwise Eve has a winning strategy.*

Proof. It is clear that Adam has a winning strategy if $G \supseteq T_\omega$: he chooses a labelling that is injective on this T_ω , and is then able to beat Eve in every move. We shall assume that $G \not\supseteq T_\omega$, and show that Eve has a winning strategy.

As in the proof of the dominating graph theorem, we start by recursively defining a rank function ρ on some or all of the vertices of G . For each ordinal α , give rank $\alpha =: \rho(v)$ to all vertices v such that all but finitely many neighbours of v have rank $< \alpha$. If any vertex remains unranked, then each of these vertices has infinitely many unranked neighbours, and we may construct a $T_\omega \subset G$ from these vertices by induction (in ω steps, as in the proof of Theorem 2.3).

Thus, since $G \not\supseteq T_\omega$ by assumption, ρ gets defined for every vertex of G . We may now choose a winning strategy for Eve as follows. Let ℓ be the labelling of G chosen by Adam at the start of the game, and let Eve’s first move be arbitrary. Later, if Adam’s last chosen vertex is v , let Eve play the number

$$1 + \max \{ \ell(w) \mid w \text{ is a neighbour of } v \text{ and } \rho(w) \geq \rho(v) \};$$

by definition of $\rho(v)$, the maximum above is just one over a finite set.

Now consider a run of the game in which Eve plays the above strategy. If Adam fails to construct a ray, then Eve wins by definition. So assume that Adam does indeed construct a ray $A \subset G$. Since there is no infinite descending sequence of ordinals, A has infinitely many vertices whose rank is at most that of its successor on A . But Eve beats Adam on all these successors, so A fails to dominate her sequence e . Thus, Eve’s strategy is indeed a winning strategy. \square

References

- [[1]] B. Bollobás, *Extremal Graph Theory*, Academic Press, London 1978.
- [[2]] R. Diestel, *Graph Decompositions—a study in infinite graph theory*, Oxford University Press, Oxford 1990.
- [[3]] R. Diestel and I. Leader, A proof of the bounded graph conjecture, *Invent. math.* (to appear).
- [[4]] R. Diestel and I. Leader, Domination games on infinite graphs, submitted.
- [[5]] R. Diestel, J. Steprans and S. Shelah, Characterizing dominating graphs, submitted.
- [[6]] R. Halin, Bounded graphs, in: (R. Diestel, Ed.) *Directions in infinite graph theory and combinatorics* (Annals of Discrete Mathematics), to appear.