Discrete Mathematics 95 (1991) 303–319
North-Holland

Excluding infinite minors

Neil Robertson *

Department of Mathematics, Ohio State University, 231 West 18th Ave., Columbus, OH, 43210, USA

Paul Seymour

Bellcore, 445 South St., Morristown, NJ 07960, USA

Robin Thomas

School of Mathematics, Georgia Institute of Technology, Atlanta, GA 30332, USA

Received 15 December 1989
Revised 4 April 1990

Abstract

Robertson, N., P. Seymour and R. Thomas, Excluding infinite minors, Discrete Mathematics 95 (1991) 303–319.

Let κ be an infinite cardinal, and let H be either a complete graph with κ vertices, or a tree in which every vertex has valency κ. What can we say about graphs G which (i) have no minor isomorphic to H, or (ii) contain no subgraph which is a subdivision of H?

These four questions are answered for each infinite cardinal κ. In each case we find that there corresponds a necessary and sufficient structural condition (or, in some cases, several equivalent conditions) for G not to contain H in the appropriate way. We survey these results and a number of related theorems.

1. Introduction

Intuitively, a graph H is a *minor* of a graph G if H can be obtained from a subgraph of G by contraction. (Graphs in this paper may be infinite, and may have loops or multiple edges.) More precisely, let us say that *H is isomorphic to a minor of G* (or G has an *H-minor,* for brevity) if for each vertex $v \in V(H)$ there is a non-null connected subgraph $\alpha(v)$ of G, and for each edge $e \in E(H)$ there is an edge $\alpha(e)$ of G, such that:

(i) for distinct $v_1, v_2 \in V(H)$, $\alpha(v_1)$ and $\alpha(v_2)$ are disjoint,
(ii) for distinct $e_1, e_2 \in E(H)$, $\alpha(e_1) \neq \alpha(e_2)$,
(iii) for $v \in V(H)$ and $e \in E(H)$, $\alpha(e) \notin E(\alpha(v))$,
(iv) if $e \in E(H)$ has distinct ends $v_1, v_2 \in V(H)$ then $\alpha(v_1)$, $\alpha(v_2)$ both contain

* This research was carried out under a consulting agreement with Bellcore.

ends of $\alpha(e)$, while if $e \in E(H)$ is a loop with end v then $\alpha(v)$ contains both ends of $\alpha(e)$.

(In particular, the subgraphs $\alpha(v)$ may be infinite, and so the intuitive definition 'obtainable from a subgraph by contracting edges' must be handled with caution.)

The first two authors proved, in a long series of papers culminating in [12, 13], the following conjecture of Wagner.

1.1. *For any infinite sequence $G_1, G_2, \ldots$ of finite graphs there exist $j > i \geq 1$ such that G_i is isomorphic to a minor of G_j.*

The third author disproved the extension of 1.1 to infinite graphs [19], in the following.

1.2. *There is an infinite sequence $G_1, G_2, \ldots$ of graphs such that for all $j > i \geq 1$, G_i is not isomorphic to a minor of G_j.*

However, the graphs in this counterexample are uncountable; and it remains open to decide the following.

1.3. Problem. *Is there an infinite sequence $G_1, G_2, \ldots$ of countable graphs such that for all $j > i \geq 1$, G_i is not isomorphic to a minor of G_j?*

The results we wish to survey here were motivated by an attempt to answer 1.3 negatively. Although the attempt was unsuccessful, we have discovered a great number of new decomposition theorems for infinite graphs of independent interest.

The method of proof of 1.1 was the following: Suppose that $G_1, G_2, \ldots$ is as in 1.1. We may assume that none of $G_2, G_3, \ldots$ has a G_1-minor, and so there is a finite complete graph H say such that $G_2, G_3, \ldots$ all have no H-minor. But there is a theorem that for any finite complete graph H, all finite graphs with no H-minor have a restricted structure (they are tree-structures of pieces which more or less have bounded genus). It suffices then to show that if $G_2, G_3, \ldots$ is an infinite sequence of finite graphs each with this restricted structure, then there exist $j > i \geq 2$ such that G_i is isomorphic to a minor of G_j, and this can be done. Thus, the heart of the proof is the theorem about the structure of graphs with no H-minor. The analogous approach to 1.3 would require a theorem about the structure of graphs with no $K_{\aleph_0}$-minor; and we have indeed been able to obtain such a theorem, and a generalization to complete graph minors of all other cardinalities.

For κ a regular uncountable cardinal, a theorem of Jung [5] says that G has a K_κ-minor if and only if G topologically contains K_κ (defined later). For κ singular or $\kappa = \aleph_0$ this is not true, but in all cases there is a similar structure theorem for graphs not topologically containing K_κ. We also study the structure of graphs not containing the tree T_κ of valency κ (either as a minor or topologically). These problems are closely related to a cops-and-robber game played on a graph, where $<\kappa$ cops try to corner a robber. It turns out, for κ uncountable, that the robber can survive if and only if the graph has a T_κ-minor; and there is a particularly simple kind of survival strategy if and only if there is a K_κ-minor.

The paper is organized as follows: In Section 2 we introduce the game and certain survival strategies (called 'escapes' and 'havens'), and describe their connection with minors of the graph. In Section 3 we observe that the non-existence of escapes is equivalent with the graph having a certain kind of decomposition. Sections 4–6 describe different kinds of decompositions, appropriate to different excluded objects. In Section 7 we discuss the structure corresponding to excluding $K_{\aleph_0}$-minors, which slipped through the net of the previous chapters. Then in Section 8, all the results are summarized.

2. Cops and a robber

Here is a game, played on a graph G. There are two players, one controlling the cops and the other the robber. There is a fixed cardinal κ (finite or infinite) and at any time the cardinality of the set of cops in play is constrained to be less than κ. The object of the cop player is to corner the robber, and the robber attempts to survive uncaptured. The robber occupies a vertex of the graph, and may at any time run at great speed along a finite path of the graph to another vertex. He is not permitted to run through a cop, however. Each cop is placed on a vertex of the graph or is temporarily removed from the graph. In the cop player's turn, he may either remove cops from the graph, or place new cops on arbitrary vertices of the graph, subject only to the constraint that fewer than κ cops may be in place in the graph at any time. The robber makes his response after the cop player has declared where he intends to place his new cops if any and before the new cops are actually in position. This is a full knowledge game, and in particular the cops know where the robber is; the problem is to capture him by landing a cop on the same vertex. (There is another version of the game on finite graphs, investigated by La Paugh [8] and Kirousis and Papadimitriou [6], where the robber is invisible, but that is different, and will not concern us. Our game can also be played transfinitely, where 'limit moves' are permitted, but we shall not consider this. For us, each turn is the tth turn for some integer t.)

For instance, if the graph G is a finite tree, then the cop player can win with only two cops at his disposal. He places cop 1 on some vertex v_1, and examines which component C of $G \backslash v_1$ contains the robber. He chooses a vertex v_2 of C adjacent to v_1 and places cop 2 there, and removes cop 1. By repeating this process the robber is eventually trapped in a leaf of the tree and captured. In general, if the cop player wins using $<\kappa$ cops, we say that '$<\kappa$ cops can search the graph'. Thus, <3 cops can search any finite tree.

On the other hand, if G is an infinite graph, containing a ray R (a *ray* is a 1-way infinite path) then $<\aleph_0$ cops cannot search the graph. For the robber can survive by remaining in the infinite part of R, further along R than any of the finitely many vertices occupied by cops. (When new cops land, he just runs further along R.)

We remark that when κ is infinite, there is no point in the cop player ever removing from a graph any cop which has already been positioned; for it costs nothing to create a new cop instead. In the finite case, however, it is evidently important to reuse the same cops.

Let G be a graph and $X \subseteq V(G)$. We call the vertex set of a component of $G \backslash X$ an X-*flap*. Let us state the game more formally. A *position* is a pair (X, F) where $X \subseteq V(G)$ with $|X| < \kappa$, and F is an X-flap. We set $X_1 = \emptyset$ and the game starts by the robber choosing a $\emptyset$-flap F_1. In general, at the beginning of the tth turn we have a position (X_{t-1}, F_{t-1}). (X_{t-1} is the set of vertices occupied by cops; and F_{t-1} is the X_{t-1}-flap containing the robber. Since he can run arbitrarily fast, all that matters is which X_{t-1}-flap contains him.) The cop player chooses a set $X_t \subseteq V(G)$ with $|X_t| < \kappa$ such that $X_t \subseteq X_{t-1}$ or $X_{t-1} \subseteq X_t$, and then the robber chooses an X_t-flap F_t such that $F_t \cap F_{t-1} \neq \emptyset$ if possible. (If there is no such choice of F_t, the cops have captured the robber.) Thus if $X_t \subseteq X_{t-1}$ then $F_t \supseteq F_{t-1}$ and F_t is uniquely determined; but if $X_t \supseteq X_{t-1}$ then the robber makes a choice. This completes the tth turn.

Of interest to us at the moment are survival strategies for the robber. (Later we shall study search strategies for the cops.) We denote by $[V(G)]^{<\kappa}$ the set of all subsets of $V(G)$ of cardinality $<\kappa$. An *escape of order* κ is a function β with domain $[V(G)]^{<\kappa}$, such that:

(i) if $X, Y \in [V(G)]^{<\kappa}$ and $X \subseteq Y$ then $\beta(X)$ is the union of all X-flaps which intersect $\beta(Y)$, and

(ii) $\beta(\emptyset) \neq \emptyset$ (and hence $\beta(X) \neq \emptyset$ for all $X \in [V(G)]^{<\kappa}$).

2.1. *$<\kappa$ cops can search G if and only if there is no escape of order κ.*

Proof. If $<\kappa$ cops cannot search G, then for each $X \in [V(G)]^{<\kappa}$, let $\beta(X)$ be the set of all vertices $v \in V(G) - X$ such that the robber can guarantee to survive if the game begins at position (X, F), where F is the X-flap containing v. Then β is an escape of order κ. Conversely, if there is an escape β of order κ, the robber can survive by always choosing $F_t \subseteq \beta(X_t)$, and so $<\kappa$ cops cannot search G. □

An escape β of order κ is a *haven of order* κ if $\beta(X)$ is an X-flap for every $X \in [V(G)]^{<\kappa}$. A haven β is *convex* if for all $X, Y \in [V(G)]^{<\kappa}$,

$$(X \cap \beta(Y)) \cup (Y \cap \beta(X)) \cup (\beta(X) \cap \beta(Y)) \neq \emptyset.$$

For G, κ finite these three concepts are closely related, as we see from the following rather difficult theorem of [18].

2.2. *Let G be a finite graph and κ a finite cardinal. Then the following are equivalent*:

(i) *there is an escape of order κ in G,*
(ii) *there is a haven of order κ in G,*
(iii) *there is a convex haven of order κ in G.*

The equivalence of (ii) and (iii) also holds if G, κ are infinite; indeed, if $\kappa \geqslant \aleph_0$ then it is easy to prove that *every* haven of order κ is convex. However, the equivalence of (i) and (ii) fails for $\kappa \geqslant \aleph_1$, as we shall see. For any cardinal $\kappa \geqslant 2$, let T_κ be the tree in which every vertex has valency κ. The connection between the game and the minors of interest to us is given by the following (from which in particular we see that T_κ has an escape of order κ but no haven of order κ, when $\kappa \geqslant \aleph_1$).

2.3 [14, 17]. *For $\kappa \geqslant \aleph_1$, G has an escape of order κ if and only if G has a T_κ-minor , and G has a haven of order κ if and only if G has a K_κ-minor.*

Proof. The 'only if' parts are difficult, and we omit them, but the 'if' parts are easy. Suppose that G has a T_κ-minor. For each $X \in [V(G)]^{<\kappa}$, let $\beta(X)$ be the union of all X-flaps F such that the restriction of G to F has a T_κ-minor; then β is an escape of order κ. If G has a K_κ-minor, let α be as in the definition of 'minor'. For each $X \in [V(G)]^{<\kappa}$, let $\beta(X)$ be the X-flap which includes $V(\alpha(v))$ for some $v \in V(K_\kappa)$. (This exists, because X is too small to intersect every $V(\alpha(v))$, and is unique, because any two distinct $V(\alpha(v))$'s are joined by an edge of G.) Then β is a haven of order κ. □

The significance of this result is that in practice it is often easier to construct a haven or an escape than to construct the desired minor directly; and conversely by exhibiting a search strategy for the cops we can sometimes prove that no T_κ-minor exists.

What about escapes and havens of order $\aleph_0$? That is answered by the following.

2.4 [17]. *For a graph G, the following are equivalent*:

(i) *G has an escape of order $\aleph_0$,*
(ii) *G has a haven of order $\aleph_0$,*
(iii) *G has a ray.*

Proof. Let R be a ray of G and for each finite $X \subseteq V(G)$ define $\beta(X)$ to be the unique F with X-flap $F \cap V(R)$ infinite; then β is a haven of order $\aleph_0$. Thus (iii)$\Rightarrow$(ii)$\Rightarrow$(i). To see (i)$\Rightarrow$(iii), here is a strategy for the cops: place a new cop on every vertex just visited by the robber (including those on the paths he runs along). This forces the robber either to trace out a ray, or to be captured, and so (i)$\Rightarrow$(iii). □

In fact, a stronger version of the equivalence of (ii) and (iii) above holds. Let us say two rays R_1, R_2 in G are *parallel* if for every finite $X \subseteq V(G)$, the unique X-flap F with $F \cap V(R_1)$ infinite has $F \cap V(R_2)$ infinite. This is an equivalence relation on the rays of G, and its equivalence classes are called the *ends* of G (see [4]). We see that, if we construct a haven as in the proof of 2.4 starting from a ray R, then two rays yield the same haven if and only if they belong to the same end. Moreover, since every haven of order $\aleph_0$ arises from a ray, we have in this sense the following.

2.5. *There is a natural* $1-1$ *correspondence between the ends of* G *and the havens in* G *of order* $\aleph_0$.

Escapes and havens of finite order in finite graphs do not correspond so closely to minors, despite 2.2. The $n \times n$ *grid* is the graph with vertex set $\{(i, j)$: $1 \leqslant i, j \leqslant n\}$ where (i, j) and (i', j') are adjacent if $|i'-i| + |j'-j| = 1$. 2.6 follows from the theorems of [11, 16,18].

2.6. *Let* G *be a finite graph. For* $n \geqslant 1$, *if* G *has an* $n \times n$ *grid minor, then* G *has a haven of order* n. *Conversely, if* G *has a haven of order* 2000^{n^5} *then* G *has an* $n \times n$ *grid minor.*

Thus, we have two differences between the finite and infinite case. First, 2.3 gives a more exact relationship between escapes, havens and minors than 2.6 does; but perhaps more surprisingly, in the finite case the minors corresponding to the existence of escapes and havens are neither trees nor complete graphs but grids.

3. Escapes and rayless tree-decompositions

A *graph-theoretic tree-decomposition* of a graph G is a pair (T, W), where T is a tree and $W = (W_t: t \in V(T))$ is a family of subsets of $V(G)$, such that:

(i) $\bigcup(W_t: t \in V(T)) = V(G)$, and for every edge $e \in E(G)$ some W_t contains both ends of e,

(ii) if $t_1, t_2, t_3 \in V(T)$, and t_2 lies on the path of T between t_1 and t_3, then $W_{t_1} \cap W_{t_3} \subseteq W_{t_2}$.

It has *width* $<\kappa$, where κ is a cardinal, if $|W_t| < \kappa$ for each $t \in V(T)$, and

$$\left|\bigcup_{j \geqslant 1} \bigcap_{i \geqslant j} W_{t_i}\right| < \kappa$$

for every ray of T with vertices $t_1, t_2, \ldots$ in order. (There need not be any κ' such that the decomposition has width $<\kappa$ if and only if $\kappa' < \kappa$, and so we shall not attempt to define the 'width' of a decomposition.) A *rayless tree-decomposition* is a graph-theoretic tree-decomposition (T, W) such that T has no ray. There is the following nice connection between rayless tree-decompositions and escapes.

3.1 [17]. *For any cardinal κ, finite or infinite, a graph G has no escape of order κ if and only if it has a rayless tree-decomposition of width $<\kappa$.*

Proof. The 'if' part is easy. For if (T, W) is a rayless tree-decomposition of width $<\kappa$ then $<\kappa$ cops can search the graph, as follows: Choose $t_1 \in V(T)$, and place cops on W_{t_1}. Choose $t \in V(T)$ such that the robber lies in W_t, and let t_2 be a neighbour of t_1 on the path between t_1 and t. Remove the cops in $W_{t_1} - W_{t_2}$, and place cops on $W_{t_2} - W_{t_1}$. Now repeat the process. Since T is rayless, the robber is eventually captured. The 'only if' part for κ finite is a difficult theorem of [18] for finite graphs G, and can be extended to infinite graphs G by a compactness argument, using the result of [8, 20]. For κ infinite, however, the 'only if' part is easy. For suppose that G has no escape of order κ. Define $\beta(X)$, for each $X \in [V(G)]^{<\kappa}$, to be the union of all X-flaps F such that the restriction of G to F has no rayless tree-decomposition of width $<\kappa$. Since κ is infinite it is easy to see that β satisfies the first escape axiom, and hence does not satisfy the second. Hence $\beta(\emptyset) = \emptyset$, and G has a rayless tree-decomposition of width $<\kappa$, as required. □

3.1 has several corollaries. From 3.1, 2.2, and 2.6, we deduce a sharpening of a result of [11].

3.2 [16]. *For any integer $n \geq 1$, every finite graph with no $n \times n$ grid minor has a graph-theoretic tree-decomposition of width $<2000^{n^5}$.*

From 3.1 and 2.4 we obtain a theorem of Halin [3].

3.3. *G has no ray if and only if G has a rayless tree-decomposition of width $<\aleph_0$.*

From 3.1 and 2.3 we obtain the first of our main theorems.

3.4 [17]. *For $\kappa \geq \aleph_1$, G has no T_κ-minor if and only if G has a rayless tree-decomposition of width $<\kappa$.*

What about $T_{\aleph_0}$-minors? Excluding them does not give us any kind of rayless tree-decomposition, but a more complicated structure, and we postpone it.

4. Dissections

A *separation* of G is a pair (A, B) of subsets of $V(G)$ such that $A \cup B = V(G)$ and no edge joins a vertex of $A - B$ to a vertex of $B - A$. Two separations (A, B) (C, D) *cross* unless either $A \subseteq C$ and $D \subseteq B$, or $A \subseteq D$ and $C \subseteq B$, or $B \subseteq C$ and $D \subseteq A$, or $B \subseteq D$ and $C \subseteq A$. A *dissection* of G is a set $\mathscr{D}$ of separations of G, such that:

(i) if $(A, B) \in \mathscr{D}$ then $(B, A) \in \mathscr{D}$,
(ii) if $(A, B) \in \mathscr{D}$ then $A \neq V(G)$,
(iii) if (A, B), $(C, D) \in \mathscr{D}$ and $A \neq C$ then $B \neq D$,
(iv) no two members of $\mathscr{D}$ cross.

Let (T, W) be a graph-theoretic tree-decomposition of G. For each $e \in E(T)$, let T_1, T_2 be the components of $T \backslash e$, and let $A_e = \bigcup (W_t\colon t \in V(T_1))$, $B_e = \bigcup (W_t\colon t \in V(T_2))$. The union over all $e \in E(T)$ of the sets $\{(A_e, B_e), (B_e, A_e)\}$ is certainly a set of mutually noncrossing separations, but it may not be a dissection because conditions (ii) and (iii) above may be violated. However, it is easy to show that any dissection of a *finite* graph G arises in this way from some graph-theoretic tree-decomposition of G, and we find that all the relevant results about tree-decompositions of finite graphs, like 3.1 and 3.2, can be reformulated in terms of dissections. Thus, dissections provide another way to generalize tree-decompositions of finite graphs to infinite graphs. It is not equivalent, for there are dissections which do not arise from any graph-theoretic tree-decomposition. (Indeed, one can characterize those dissections $\mathscr{D}$ which do arise from graph-theoretic tree-decompositions; it is necessary and sufficient that $\bigcup (B_i\colon i \geq 1) = V(G)$ for every infinite sequence (A_i, B_i) $(i = 1, 2, \ldots)$ of distinct members of $\mathscr{D}$ with $A_1 \supseteq A_2 \supseteq \cdots$ and $B_1 \subseteq B_2 \subseteq \cdots$.)

Let $\mathscr{D}$ be a dissection. An *orientation* of $\mathscr{D}$ is a subset $\mathscr{P} \subseteq \mathscr{D}$, such that:

(i) if $(A, B) \in \mathscr{D}$ then one of (A, B), $(B, A) \in \mathscr{P}$,
(ii) if (A, B), $(C, D) \in \mathscr{P}$ then $B \nsubseteq C$.

If $\mathscr{D}$ arises from a graph-theoretic tree-decomposition (T, W) with T finite, one can verify that the orientations of $\mathscr{D}$ are in $1-1$ correspondence with $V(T)$; the orientation corresponding to $t_0 \in V(T)$ is the set of all separations

$$\left(\bigcup (W_t\colon t \in V(T_1)), \bigcup (W_t\colon t \in V(T_2))\right)$$

taken over all $e \in E(T)$, where T_1, T_2 are the components of $T \backslash e$ and $t_0 \in V(T_2)$. If $\mathscr{D}$ arises from (T, W) and T is infinite, then $\mathscr{D}$ may have orientations not corresponding to the vertices; in fact, it has one for each vertex and one for each end of T. Nevertheless, orientations provide us with a generalization of the concept of a vertex of a tree, expressed in the language of dissections.

Motivated by this, let us say a dissection $\mathscr{D}$ has *width* $<\kappa$ if for every orientation $\mathscr{P}$ of $\mathscr{D}$,

$$\left|\bigcup (B\colon (A, B) \in \mathscr{P})\right| < \kappa.$$

If the dissection arises from a graph-theoretic tree-decomposition (T, W), one can verify that this definition agrees with our definition of 'width $<\kappa$' for such a decomposition. Indeed, if $\mathcal{P}$ is determined by a vertex $t \in V(T)$, then $\bigcap(B\colon (A, B) \in \mathcal{P}) = W_t$, while if $\mathcal{P}$ is determined by an end of T containing a ray with vertices $t_1, t_2, \ldots$ in order, then

$$\bigcap (B\colon (A, B) \in \mathcal{P}) = \bigcup_{i \geq 1} \bigcap_{j \geq i} W_{t_j}.$$

Dissections and havens work well together, because of the following easy lemma. (The *order* of a separation (A, B) is $|A \cap B|$.)

4.1 [14]. *Let β be a convex haven in G of order κ, and let $\mathcal{D}$ be a dissection of G each member of which has order $<\kappa$. Then the set $\{(A, B) \in \mathcal{D}\colon \beta(A \cap B) \subseteq B\}$ is an orientation of $\mathcal{D}$.*

Now we can begin to discuss the second of our structure theorems, concerning the structure of graphs not topologically containing K_κ. A graph G *topologically contains* H if there is a function α with domain $V(H) \cup E(H)$, such that:

(i) for each $v \in V(H)$, $\alpha(v) \in V(G)$ and $\alpha(v_1) \neq \alpha(v_2)$ for distinct $v_1, v_2 \in V(H)$,

(ii) for each $e \in E(H)$, $\alpha(e)$ is a finite path of G with $E(\alpha(e)) \neq \emptyset$ and with ends $\alpha(v_1)$, $\alpha(v_2)$ where e has ends v_1, v_2 in H,

(iii) for distinct $e_1, e_2 \in E(H)$, the paths $\alpha(e_1)$ and $\alpha(e_2)$ are disjoint except possibly for their ends,

(iv) for each $v \in V(H)$ and $e \in E(H)$, $\alpha(v)$ is not an internal vertex of $\alpha(e)$.

A preliminary form of our result is the following.

4.2 [14, 15]. *For any graph G and $\kappa \geq \aleph_0$, G does not topologically contain K_κ if and only if there is a dissection of G of width $<\kappa$.*

Proof. 'Only if' is complicated and we omit it, but let us prove 'if'. Let $\mathcal{D}$ be a dissection of width $<\kappa$; then every member of $\mathcal{D}$ has order $<\kappa$, as is easily seen, by considering the orientation $\{(A, B) \in \mathcal{D}\colon A_0 \subseteq B$ or $B_0 \subseteq B$, and $A_0 \neq B\}$, for each $(A_0, B_0) \in \mathcal{D}$. Suppose, for a contradiction, that G topologically contains K_κ, and let α be the corresponding function. Let $Z = \{\alpha(v)\colon v \in V(K_\kappa)\}$. For each $X \in [V(G)]^{<\kappa}$, let $\beta(X)$ be the X-flap containing some member of Z. (This exists since X is too small to include Z, and is unique since any two members of Z are joined by κ mutually disjoint paths, not all of which meet X.) Then β is a convex haven of order κ. By 4.1,

$$\mathcal{P} = \{(A, B) \in \mathcal{D}\colon \beta(A \cap B) \subseteq B\}$$

is an orientation of $\mathcal{D}$. Let $X = \bigcap(B\colon (A, B) \in \mathcal{P})$; then $|X| < \kappa$ since $\mathcal{D}$ has width $<\kappa$. Choose $v \in Z - X$. Since $v \notin X$, there exists $(A, B) \in \mathcal{P}$ with $v \in A - B$. But since $v \notin A \cap B$ it follows from the definition of β that $v \in \beta(A \cap B)$, and so $\beta(A \cap B) \nsubseteq B$, a contradiction. $\square$

Given that G does not topologically contain K_κ, can we prove the existence of a more concrete structure than just a dissection of width $<\kappa$? An example of [7] shows that G need not have a graph-theoretic tree-decomposition of width $<\kappa$, if $\kappa > \aleph_0$. Nevertheless, we can improve 4.2 by the use of 'well-founded tree-decompositions', which are midway in generality between graph-theoretic tree-decompositions and dissections.

A *well-founded tree* is a non-null partially ordered set $T = (V(T), \leqslant)$ such that:
(i) for each $s \in V(T)$, the set $\{t \in V(T)\colon t \leqslant s\}$ is well-ordered by $\leqslant$,
(ii) every non-empty $X \subseteq V(T)$ has an *infimum*, that is, an element $z \in V(T)$ such that for all $t \in V(T)$, $t \leqslant z$ if and only if $t \leqslant x$ for all $x \in X$.

It follows that there is a unique element $s \in V(T)$ such that $s \leqslant t$ for all $t \in V(T)$, which we call the *root*. If $t_1, t_2, t_3 \in V(T)$, we say that t_2 is *between* t_1 and t_3 if $t_0 \leqslant t_2$, where t_0 is the infimum of $\{t_1, t_3\}$, and either $t_2 \leqslant t_1$ or $t_2 \leqslant t_3$. A *well-founded tree-decomposition of a graph* G is a pair (T, W), where $T = (V(T), \leqslant)$ is a well-founded tree and $W = (W_t\colon t \in V(T))$ is a family of subsets of $V(G)$, such that:
(i) $\bigcup (W_t\colon t \in V(T)) = V(G)$, and every edge of G has both ends in some W_t,
(ii) if $t_1, t_2, t_3 \in V(T)$ and t_2 is between t_1 and t_3 then $W_{t_1} \cap W_{t_3} \subseteq W_{t_2}$,
(iii) if C is a chain of T with supremum s (that is, C is totally ordered by $\leqslant$, and for all $t \in V(T)$, $c \leqslant t$ for all $c \in C$ if and only if $s \leqslant t$), then $\bigcap (W_c\colon c \in C) \subseteq W_s$.

We say that (T, W) has *width* $<\kappa$ if for every chain C of T,

$$\left| \bigcup_{c \in C} \bigcap_{t \geqslant c,\, t \in C} W_t \right| < \kappa.$$

(We observe, by taking $|C| = 1$, that this implies that $|W_t| < \kappa$ for all $t \in V(T)$; but it is somewhat stronger, for some chains may not have suprema.) If (T, W) is a graph-theoretic tree-decomposition and we regard T as a well-founded tree by choosing a root in the natural way, this definition of 'width $<\kappa$' agrees with the earlier one.

A well-founded tree T has *height* $<\kappa$ if every chain in T has order type $<\kappa$, regarded as an ordinal. Then we can vary 4.2 as follows.

4.3 [15]. *For any graph G and cardinal $\kappa \geqslant \aleph_0$, the following are equivalent:*
(i) *G does not topologically contain K_κ,*
(ii) *G has a well-founded tree-decomposition of width $<\kappa$.*
If κ is regular with $\kappa > \aleph_0$, these are equivalent to
(iii) *G has a well-founded tree-decomposition of width $<\kappa$ and height $<\kappa$.*
If $\kappa = \aleph_0$, (i), (ii) *are equivalent to*
(iv) *G has a graph-theoretic tree-decomposition of width $<\aleph_0$.*

Diestel [2] has independently proved a structural characterization of the graphs not topologically containing K_κ, in the case when $\kappa > \aleph_0$ and is regular. His result uses a different generalization of finite tree-decompositions, but is similar to (and interderivable with) the equivalence of (i) and (iii) in 4.3.

The need for our rather curious definition of width is shown by the following example: Let $G = K_{\aleph_0}$, with $V(G) = \{v_0, v_1, v_2, \ldots\}$. Let T be a ray with vertices $t_0, t_1, t_2, \ldots$ in order, and for each $i \geq 0$ let $W_{t_i} = \{v_0, \ldots, v_i\}$. Let $W = (W_t\colon t \in V(T))$; then (T, W) is a graph-theoretic tree-decomposition of G and each $|W_t| < \aleph_0$, yet G topologically contains $K_{\aleph_0}$.

We could replace the condition about chains in the definition of the width of a well-founded tree-decomposition by the weaker condition that each $|W_t| < \kappa$, together with an extra condition that every chain has a supremum; for one can always add to T any suprema that are missing. However, doing so can increase the height of T, and on this definition graph-theoretic tree-decompositions would not be well-founded tree-decompositions unless they were rayless; so we prefer the version given.

5. Adhesion

So far, we have discussed excluding T_κ-minors, and excluding K_κ topologically, because these are the simpler theorems to state. Next, let us consider excluding K_κ-minors. For regular $\kappa > \aleph_0$, this is equivalent to excluding K_κ topologically, for a theorem of Jung [5] states the following.

5.1. *For $\kappa > \aleph_0$ regular, a graph has a K_κ-minor if and only if it topologically contains K_κ.*

There is a similar result for T_κ when κ is regular and uncountable. However, for κ singular or $\kappa = \aleph_0$, these are false; indeed, in those cases there is a graph with a K_κ-minor in which all vertices have valency $<\kappa$. In particular then, if κ is singular or $\kappa = \aleph_0$, the structure of 4.2 (ii) or 4.3 (ii) is not sufficiently restrictive to exclude K_κ-minors, and we need to impose an additional condition on the dissection.

Let $\mathcal{D}$ be a dissection of a graph G. We say that $\mathcal{D}$ has *adhesion* $<\kappa$ if for every orientation $\mathcal{P}$ of $\mathcal{D}$ there exists $\kappa' < \kappa$ such that for every $(A, B) \in \mathcal{P}$ there exists $(A', B') \in \mathcal{P}$ of order $\leq \kappa'$ with $A \subseteq A'$ and $B' \subseteq B$. Then we have the following.

5.2 [14]. *For all $\kappa > \aleph_0$, a graph has no K_κ-minor if and only if it has a dissection of width $<\kappa$ and adhesion $<\kappa$.*

Even the 'easy' part of this, the 'if' half, is not very easy, and so we omit the proof completely.

There is a version of 5.2 in terms of well-founded tree-decompositions, analogous to 4.3 but we omit details—see [14]. Also, 5.2 requires that $\kappa > \aleph_0$. We do have a structure theorem for excluding $K_{\aleph_0}$-minors, but postpone it because it is quite different.

6. Linear decompositions

Our results so far were motivated by attempts to find infinite analogues of 3.2. However, 3.2 concerns excluding finite grids, and we really want to exclude T_κ or K_κ. Thus, one might think that the following result is more likely to have an interesting infinite analogue.

6.1 [1, 10]. *For any finite tree H, every finite graph with no H-minor has a tree-decomposition (T, W) of width $<|V(H)|$ such that T is a path.*

Let us define a *linear decomposition* (T, W) of G to consist of a Dedekind complete, linearly ordered set $T = (V(T), \leqslant)$, and a family $W = (W_t\colon t \in V(T))$ of subsets of $V(G)$, such that:

(i) $\bigcup (W_t\colon t \in V(T)) = V(G)$, and every edge of G has both ends in some W_t,

(ii) if $t_1, t_2, t_3 \in V(T)$ and $t_1 \leqslant t_2 \leqslant t_3$ then $W_{t_1} \cap W_{t_3} \subseteq W_{t_2}$,

(iii) if $s \in V(T)$ is the supremum or infimum of a non-null subset $C \subseteq V(T)$ then $\bigcap (W_c\colon c \in C) \subseteq W_s$. It has *width* $<\kappa$ if $|W_t| < \kappa$ for each t.

Seese [21] extended 6.1 to infinite graphs by a compactness argument, to yield the following.

6.2 [21]. *For any finite tree H, every graph with no H-minor has a linear decomposition of width $<|V(H)|$.*

For countable trees a similar result holds, because of the following.

6.3 [17]. *A graph does not topologically contain $T_{\aleph_0}$ if and only if it has a linear decomposition of width $<\aleph_0$.*

But for larger cardinals, there is no similar result because T_κ has a linear decomposition of width $<\aleph_1$ (enumerate the rays from left to right). For larger regular cardinals, the appropriate excluded object is not a tree but a clique.

6.4 [15]. *For $\kappa > \aleph_0$, regular, a graph does not topologically contain K_κ if and only if it has a linear decomposition of width $<\kappa$.*

Proof. We only show the easy 'if' half. Suppose that (T, W) is a linear decomposition of width $<\kappa$ of G, and yet G topologically contains K_κ. Thus there exists $X \subseteq V(G)$ with $|X| = \kappa$ such that any two members of X are joined by κ internally disjoint paths. For each $x \in X$, let $I_x = \{t \in V(T)\colon x \in W_t\}$. Then I_x

is a closed interval. For distinct x, x', $I_x \cap I_{x'} \neq \emptyset$, for otherwise there exists $t_0 \in V(t)$ such that $t < t_0$ for all $t \in I_x$ and $t > t_0$ for all $t \in I_{x'}$ (or vice versa); and this is impossible, because then every path of G between x and x' meets W_{t_0}, and $|W_{t_0}| < \kappa$. Thus $I_x \cap I_{x'} \neq \emptyset$, and so from Dedekind completeness there exists $t \in V(T)$ such that $t \in I_x$ for all $x \in X$, that is, $X \subseteq W_t$. But that is impossible since $|W_t| < \kappa$. □

For $\kappa = \aleph_\omega$, 6.4 is false. Indeed, we give a graph in [15] which has no $K_{\aleph_\omega}$-minor, and yet has no linear decomposition of width $<\aleph_\omega$.

A *well-ordered decomposition* is a linear decomposition (T, W) such that T is a well-order. The third of our main results is the following.

6.5 [17]. *For all $\kappa \geq \aleph_0$, a graph does not topologically contain T_κ if and only if it has a well-ordered decomposition of width $<\kappa$.*

Proof. Again, we shall only prove the 'if' part. Let (T, W) be a well-ordered decomposition of G of width $<\kappa$, and suppose that G topologically contains T_κ. Thus, there exists $X \subseteq V(G)$ with $|X| = \kappa$ such that each $x \in X$ is joined by κ paths, mutually disjoint except for x, to κ other members of X. For each $x \in X$, choose $s(x) \in V(T)$ maximal such that $x \in W_{s(x)}$. Choose $x \in X$ with $s(x)$ minimal. (This is possible since T is a well-order.) Certainly $s(x)$ is not the supremum of $V(T)$, and so there exists t say with $s(x) < t$ such that there is no $t' \in V(T)$ with $s(x) < t' < t$. Since there are κ paths from x to other members of X, and $|W_{s(x)} \cup W_t| < \kappa$, there is a path P between x and some $x' \in X - \{x\}$ such that $V(P) \cap (W_{s(x)} \cup W_t) = \{x\}$. It follows that $s(x') < t$, and so $s(x') \leq s(x)$, and equality does not hold since $x' \notin W_{s(x)}$. But this contradicts the choice of x. □

Let us say a linear decomposition (T, W) has *adhesion* $<\kappa$ if for every $s \in V(T)$ there exists $\kappa' < \kappa$ such that for every $t \in V(T)$ with $t \neq s$ there exist $t_1, t_2 \in V(T)$ with $t_1 < t_2$ such that there is no $t_3 \in V(T)$ with $t_1 < t_3 < t_2$ and $|W_{t_1} \cap W_{t_2}| \leq \kappa'$, where $s \leq t_1 < t_2 \leq t$ if $s < t$ and $t \leq t_1 < t_2 \leq s$ if $t < s$.

6.6 [17]. *For all $\kappa \geq \aleph_0$, a graph has no T_κ-minor if and only if it has a well-ordered decomposition of width $<\kappa$ and adhesion $<\kappa$.*

Thus, 6.6 is related to 6.5 as 5.2 is to 4.2. Corresponding to 6.3, we have the following.

6.7 [17]. *A graph has no $T_{\aleph_0}$-minor if and only if it has a linear decomposition of width $<\aleph_0$ and adhesion $<\aleph_0$.*

We also have variant forms of these results. A linear decomposition (T, W) is *scattered linear* if the set of rational numbers cannot be monotonely embedded into T.

6.8 [17]. *For all $\kappa \geqslant \aleph_0$, a graph does not topologically contain T_κ if and only if it has a scattered linear decomposition of width $<\kappa$.*

6.9 [17]. *For all $\kappa \geqslant \aleph_0$, a graph has no T_κ-minor if and only if it has a scattered linear decomposition of width $<\kappa$ and adhesion $<\kappa$.*

6.9 fills a gap left by 3.4, which did not cater for excluding $T_{\aleph_0}$. Another way is given by 6.11 below. A graph-theoretic tree-decomposition (T, W) is *scattered* if T has no $T_{\aleph_0}$-minor (or equivalently, does not topologically contain T_3).

6.10 [17]. *For any cardinal κ with $\mathrm{cf}(\kappa) = \omega$, a graph does not topologically contain T_κ if and only if it has a scattered graph-theoretic tree-decomposition of width $<\kappa$.*

6.10 does not hold in general if $\mathrm{cf}(\kappa) > \omega$. (There is a form which holds in general using 'scattered' well-founded tree-decompositions—see [17].) But its analogue for T_κ minors does hold in general, as follows.

6.11 [17]. *For all $\kappa \geqslant \aleph_0$, a graph has no T_κ-minor if and only if it has a scattered graph-theoretic tree-decomposition of width $<\kappa$ and adhesion $<\kappa$.*

(A graph-theoretic tree-decomposition (T, W) has *adhesion* $<\kappa$ if $|W_s \cap W_t| < \kappa$ for every pair s, t of adjacent vertices of T, and for every ray of T with vertices $t_1, t_2, \ldots$ in order, there exists $\kappa' < \kappa$ such that $|W_{t_i} \cap W_{t_{i+1}}| \leqslant \kappa'$ for infinitely many integers $i \geqslant 1$.)

7. Countable clique and grid minors

There remains one case not yet covered; what is the structure of graphs with no $K_{\aleph_0}$-minor? By analogy with 5.2 one might guess that it was necessary and sufficient that the graph have a dissection of width $<\aleph_0$ and adhesion $<\aleph_0$. That turns out to be sufficient but not necessary; indeed, that structure is equivalent to excluding a 'half-grid' minor. The *half-grid* is the graph with vertex set all pairs of integers (x, y) with $y \geqslant 0$, where (x, y) and (x', y') are adjacent if $|x - x'| + |y - y'| = 1$.

7.1 [14]. *A graph has no half-grid minor if and only if it has a dissection of width $<\aleph_0$ and adhesion $<\aleph_0$.*

Let us say a graph-theoretic tree-decomposition (T, W) is *narrow* if $|W_s \cap W_t|$ is finite for every pair of adjacent vertices s, t of T, and for every ray of T with vertices $t_1, t_2, \ldots$ in order, infinitely many of the numbers $|W_{t_i} \cap W_{t_{i+1}}|$ are equal.

(This is equivalent to the decomposition having adhesion $<\aleph_0$.) We can sharpen 7.1 as follows.

7.2 [14]. *A graph has no half-grid minor if and only if it has a narrow graph-theoretic tree-decomposition of width* $<\aleph_0$.

This is an extension of a result of Halin [4]. We have not been able to determine the structure corresponding to excluding the 'full' grid.

So what then is the appropriate structure for excluding $K_{\aleph_0}$-minors? Let (T, W) be a graph-theoretic tree-decomposition of G. For each $t \in V(T)$, we define the *torso at* t to be the graph with vertex set W_t in which u, v are adjacent if either they are adjacent in G or $u, v \in W_t \cap W_{t'}$ for some neighbour t' of t in T.

7.3 [14]. *A graph has no* $K_{\aleph_0}$*-minor if and only if it has a narrow graph-theoretic tree-decomposition* (T, W) *such that for each* $t \in V(T)$ *there is an integer* k *such that the torso at* t *has no* K_k*-minor.*

8. Summary

Let us summarize all these results.

8.1. *The following are equivalent*:
(i) *G has no ray,*
(ii) *$<\aleph_0$ cops can search G,*
(iii) *G has a rayless tree-decomposition of width $<\aleph_0$.*

8.2. *For all $\kappa \geqslant \aleph_0$, the following are equivalent*:
(i) *G does not topologically contain T_κ,*
(ii) *G has a well-ordered decomposition of width $<\kappa$,*
(iii) *G has a scattered linear decomposition of width $<\kappa$.*
If $\mathrm{cf}(\kappa) = \omega$, *these are equivalent to*
(iv) *G has a scattered graph-theoretic tree-decomposition of width $<\kappa$.*
If $\kappa = \aleph_0$, these are equivalent to
(v) *G has a linear decomposition of width $<\kappa$.*

8.3. *The following are equivalent*:
(i) *G has no $T_{\aleph_0}$-minor,*
(ii) *G has a scattered graph-theoretic tree-decomposition of width $<\aleph_0$ and adhesion $<\aleph_0$,*
(iii) *G has a well-ordered decomposition of width $<\aleph_0$ and adhesion $<\aleph_0$,*
(iv) *G has a linear decomposition of width $<\aleph_0$ and adhesion $<\aleph_0$,*
(v) *G has a scattered linear decomposition of width $<\aleph_0$ and adhesion $<\aleph_0$.*

8.4. *For all $\kappa > \aleph_0$, the following are equivalent*:
(i) *G has no T_κ-minor,*
(ii) *$<\kappa$ cops can search G,*
(iii) *G has a rayless tree-decomposition of width $<\kappa$,*
(iv) *G has a scattered tree-decomposition of width $<\kappa$ and adhesion $<\kappa$,*
(v) *G has a well-ordered decomposition of width $<\kappa$ and adhesion $<\kappa$,*
(vi) *G has a scattered linear decomposition of width $<\kappa$ and adhesion $<\kappa$.*

8.5. *For all $\kappa \geqslant \aleph_0$, the following are equivalent*:
(i) *G does not topologically contain K_κ,*
(ii) *G has a dissection of width $<\kappa$,*
(iii) *G has a well-founded tree-decomposition of width $<\kappa$.*
If $\kappa > \aleph_0$ and is regular, then (i)–(iii) *are equivalent to*
(iv) *G has a well-founded tree-decomposition of width $<\kappa$ and height $<\kappa$,*
(v) *G has a linear decomposition of width $<\kappa$.*
If $\kappa = \aleph_0$ then (i)–(iii) *are equivalent to*
(vi) *G has a graph-theoretic tree-decomposition of width $<\aleph_0$.*

8.6. *A graph has no $K_{\aleph_0}$-minor if and only if it has a narrow graph-theoretic tree-decomposition (T, W), such that for each $t \in V(T)$ there is an integer k such that the torso at t has no K_k-minor.*

8.7. *For all $\kappa > \aleph_0$, G has no K_κ-minor if and only if G has a dissection of width $<\kappa$ and adhesion $<\kappa$.*

8.8. *The following are equivalent*:
(i) *G has no half-grid minor,*
(ii) *G has a dissection of width $<\aleph_0$ and adhesion $<\aleph_0$,*
(iii) *G has a narrow graph-theoretic tree-decomposition of width $<\aleph_0$.*

References

[1] D. Bienstock, N. Robertson, P.D. Seymour and R. Thomas, Quickly excluding a forest, J. Combin. Theory Ser. B. 52 (1991) 274–283.
[2] R. Diestel, The structure of TK_a-free graphs, submitted.
[3] R. Halin, Graphen ohne unendliche Wege, Math Nachr. 31 (1966) 111–123.
[4] R. Halin, Über die Maximalzahl fremder unendlicher Wege in Graphen, Math. Nachr. 30 (1965) 63–85.
[5] H.A. Jung, Zusammenzüge und Unterteilungen von Graphen, Math. Nachr. 35 (1967) 241–268.
[6] L.M. Kirousis and C.H. Papadimitriou, Interval graphs and searching, Discrete Math. 55 (1985) 181–184.
[7] I. Kříž and R. Thomas, Clique-sums, tree-decompositions and compactness, Discrete Math. 81 (1990) 177–185.
[8] I. Kříž and R. Thomas, The Menger-like property of tree-width of infinite graphs and related compactness results, J. Combin. Theory, Ser. B. 52 (1991) 86–91.

[9] A. La Paugh, Recontamination does not help to search a graph, 1982, manuscript.
[10] N. Robertson and P.D. Seymour, Graph minors I. Excluding a forest, J. Combin. Theory Ser. B. 35 (1983) 39–61.
[11] N. Robertson and P.D. Seymour, Graph minors V. Excluding a planar graph, J. Combin. Theory Ser. B 41 (1986) 92–114.
[12] N. Robertson and P.D. Seymour, Graph minors XV. Wagner's conjecture, manuscript.
[13] N. Robertson and P.D. Seymour, Graph minors XVI. Well-quasi-ordering on a surface, manuscript.
[14] N. Robertson, P.D. Seymour and R. Thomas, Excluding infinite clique minors, manuscript.
[15] N. Robertson, P.D. Seymour and R. Thomas, Excluding infinite clique subdivisions, Trans. Amer. Math. Soc., to appear.
[16] N. Robertson, P.D. Seymour and R. Thomas, Quickly excluding a planar graph, submitted.
[17] P.D. Seymour and R. Thomas, Excluding infinite trees, Trans. Amer. Math. Soc., to appear.
[18] P.D. Seymour and R. Thomas, Graph searching, and a minimax theorem for tree-width, submitted.
[19] R. Thomas, A counterexample to Wagner's conjecture for infinite graphs, Math. Proc. Cambridge Philos. Soc. 103 (1988) 55–57.
[20] C. Thomassen, Configurations in graphs of large minimum degree, connectivity or chromatic number, in: G.S. Bloom, R.L. Graham and J. Malkevitch, eds., Combinatorial Mathematics, Proc. 3rd International Conference 1985, Ann. New York Acad. Sci., Vol. 555 (New York Acad. Sci., New York, 1989) 402–411.
[21] D. Seese, Ordered tree representations of infinite graphs, 1988, preprint.