Cycles, Minors and Trees

DISSERTATION

zur Erlangung des Doktorgrades

des Fachbereichs Mathematik

der Universität Hamburg

vorgelegt von

Daniela Kühn

aus Chemnitz

Hamburg

2001

Als Dissertation angenommen vom Fachbereich
Mathematik der Universität Hamburg

auf Grund der Gutachten von Prof. R. Diestel, PhD
und Prof. Dr. T. Andreae

Hamburg, den 27.06.2001

Prof. Dr. R. Hass
Dekan des Fachbereichs Mathematik

Contents

Introduction

The purpose of this chapter is to summarize the results obtained in this thesis. A more detailed discussion of these results is deferred to the relevant chapters. In Chapter 1 we develop the notion of an infinite cycle and prove several infinite analogues of well-known facts about the cycle space of a finite graph. Chapters 2 and 3 both deal with substructures of graphs of large average degree. In Chapter 4 we prove several results about partitions of graphs, one of which will enable us to strengthen the main result of Chapter 2 for graphs of high connectivity. Rich substructures are also the theme of Chapter 5. However, instead of considering graphs of high average degree or connectivity, we are there concerned with graphs which contain a large externally highly connected subgraph. Finally, in Chapter 6 we give a short proof of the result of Nash-Williams that the infinite trees are well-quasi-ordered by the topological minor relation, i.e. that for every infinite sequence $T_1, T_2, \ldots$ of infinite trees there exist indices $i < j$ such that T_i is a topological minor of T_j. The terminology we use in this thesis is that of [6].

Infinite cycles

It is well-known that for every spanning tree T of a finite graph G the fundamental cycles (those consisting of a chord xy of T together with the path in T joining x to y) generate the entire cycle space of G—every element of the cycle space can be expressed as a sum mod 2 of fundamental cycles. Richter [39] asked in which way this fact might generalize to locally finite infinite graphs.

If we define the cycle space of a locally finite graph G in the same way as for a finite graph, i.e. as the set of all finite sums mod 2 of finite cycles, then of course, for any spanning tree of G, the fundamental cycles generate the entire cycle space. But an infinite graph may have infinitely many finite cycles. So it is natural to allow also well-defined infinite sums of finite cycles—those for which every edge lies in only finitely many summands and for which it is therefore decidable whether any given edge belongs to the sum or not. This already shows that Richter's question leads to something new: now the elements of the cycle space can be infinite, but the fundamental cycles are still finite.

Furthermore, one might try to generalize the notion of a cycle to obtain infinite cycles as well. Then the cycle space would consist of well-defined sums mod 2 of finite or infinite cycles. But what should these infinite cycles be? If we look at a 2-way infinite ladder, then it is natural to say that its two sides form a cycle together with the two points at infinity (Fig. 1). Similarly, the three sides of a 3-way infinite ladder form a cycle together with the three points at infinity. In Chapter 1 we generalize these examples by defining an infinite cycle

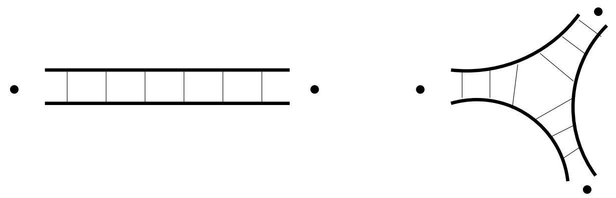

Figure 1: Infinite cycles in 2- and 3-way infinite ladders

to be a homeomorphic image of the unit circle in the graph compactified by its ends. We show that for a locally finite connected graph G, the spanning trees of G for which the fundamental cycles generate every element of the cycle space of G are precisely the end-respecting (equivalently: end-faithful) spanning trees of G:

Theorem *Let G be a locally finite connected graph and let T be any spanning tree of G. Then the following are equivalent:*

(i) *Every cycle of G is a sum of fundamental cycles.*

(ii) *Every element of the cycle space of G is a sum of fundamental cycles.*

(iii) *T is end-respecting.*

It turns out that this theorem does not hold if we no longer restrict ourselves to locally finite graphs. However, even for general infinite graphs we will still be able to characterize the spanning trees for which the fundamental cycles generate every cycle, and those for which the same holds for all elements of the cycle space. Surprisingly, there are graphs for which these two sets of spanning trees differ. We also prove an infinite analogue of the fact that every element of the cycle space of a finite graph is an edge-disjoint union of finite cycles:

Theorem *Every element of the cycle space of an infinite graph is an edge-disjoint union of (finite and infinite) cycles.*

The results of Chapter 1 are joint work with R. Diestel [10, 11].

Induced subdivisions in $K_{s,s}$-free graphs of large average degree

A classical theorem of Mader states that for any given graph H every graph of sufficiently large average degree contains a subdivision of H. (The *average degree* of a graph G is equal to the average number of edges incident with a vertex of G. A *subdivision* of a graph G is a graph obtained from G by replacing the edges of G by internally disjoint paths.) Clearly, if we ask for an *induced* subdivision of H, then a statement analogous to Mader's theorem does not hold: for example, the complete bipartite graph $K_{s,s}$ has average degree s but it does not contain an induced subdivision of a cycle of length six. (A subgraph

G' of G is *induced* if every edge of G joining two vertices of G' is also contained in G'.) The purpose of Chapter 2 is to show that if we restrict our attention to $K_{s,s}$-free graphs then an analogue of Mader's theorem is indeed true:

Theorem *For every graph H and every $s \in \mathbb{N}$ there exists $d = d(H, s)$ such that every graph of average degree at least d contains either $K_{s,s}$ as a subgraph or an induced subdivision of H.*

Kierstead and Penrice [21] showed that a stronger statement is true when H is a tree: any given tree can be found as an induced *subgraph* in every $K_{s,s}$-free graph of sufficiently large average degree. Clearly, the stronger statement does not hold if H contains a cycle. In Chapter 2 we also include an alternative proof of the result of Kierstead and Penrice as well as an elementary proof of the special case of the above theorem when H is a cycle. The results of this chapter are joint work with D. Osthus [26, 27].

Subgraphs of large average degree containing no cycle of length less than six

By a classical theorem of Erdős there exist graphs which have both arbitrarily large average degree and arbitrarily high girth. (The *girth* of a graph is the length of its shortest cycle.) Such graphs G are locally 'sparse'—a ball of small radius around a vertex of G always induces just a tree in G—but globally they are 'dense'. A conjecture of Thomassen [52] states that such graphs not only exist but occur in *all* graphs of large average degree: for all integers k, g there exists $f = f(k, g)$ such that every graph of average degree at least f has a subgraph of average degree at least k and girth at least g. The aim of Chapter 3 is to prove this conjecture for the case of $g \leq 6$:

Theorem *For every k there exists $d = d(k)$ such that every graph of average degree at least d contains a subgraph of average degree at least k and girth at least six.*

This result is joint work with D. Osthus [25].

Partitions of graphs with high minimum degree or connectivity

Hajnal [19] and Thomassen [50] independently proved that the vertex set of every highly connected graph G can be partitioned into sets S and T such that the graphs $G[S]$ and $G[T]$ induced by these sets still have high connectivity. In Chapter 4 we strengthen this result by showing that we can additionally require that every vertex in S has many neighbours in T:

Theorem *For every ℓ there exists $k = k(\ell)$ such that the vertex set of every k-connected graph G can be partitioned into non-empty sets S and T such that both $G[S]$ and $G[T]$ are ℓ-connected and every vertex in S has at least ℓ neighbours in T.*

This is best possible in the sense that we cannot additionally require the entire bipartite subgraph of G between S and T to have large minimum degree. If we apply Mader's theorem on subdivisions to $G[S]$ we obtain

Corollary *For every ℓ and every graph H there exists $k = k(\ell, H)$ such that every k-connected graph G contains a subdivision TH of H such that $G - V(TH)$ is ℓ-connected.*

Similarly, applying our result on induced subdivisions from Chapter 2, we derive an induced version of the above corollary for $K_{s,s}$-free graphs. We also prove several related results about partitions of graphs. The results of this chapter are joint work with D. Osthus [28].

Forcing complete minors by high external connectivity

In Chapter 2 we looked for induced subdivisions in graphs of large average degree. In Chapter 5 we will be concerned with sufficient conditions for the existence of complete minors. A graph H is a *minor* of a graph G if H can be obtained from a subgraph of G by contracting edges. So in particular, if G is a subdivision of H then H is a minor of G. Given a graph G, we say that $X \subseteq V(G)$ is *externally k-connected in G* if $|X| \geq k$ and, for all $Y, Z \subseteq X$ with $|Y| = |Z| \leq k$, there are $|Y|$ disjoint Y–Z paths in G which have no inner vertices or edges in $G[X]$. A subgraph H of G is *externally k-connected in G* if $V(H)$ is externally k-connected in G.

It it easy to see that if G is a graph which has an externally highly connected set X such that $G[X]$ contains a large grid minor, then G has a large complete minor. Thus large externally highly connected grids force large complete minors. Grohe [18] asked whether there are graphs H which are substantially thinner than grids, but which still force large complete minors in the same way. In Chapter 5 we investigate what minimum amount of structure is required for such graphs H. We observe that they must contain large binary trees with some small additions, and prove that some canonical instances of this structure are also sufficient to force large complete minors. Chapter 5 is based on [23].

On well-quasi-ordering infinite trees

The notion of externally highly connected sets, which was the subject of Chapter 5, was introduced by Diestel et al. [9] in order to give a short proof of the fundamental result of Robertson and Seymour [40] that every graph of large tree-width contains a large grid minor. This result is one of the cornerstones in Robertson's and Seymour's proof of the Graph Minor Theorem: the finite graphs are well-quasi-ordered by the minor relation, i.e. for every infinite sequence $G_1, G_2, \ldots$ of finite graphs there exist indices $i < j$ such that G_i is a minor of G_j. An example of Thomas [46] shows that the Graph Minor Theorem fails for uncountable graphs. It remains an open problem whether it holds for countable graphs.

If we restrict our attention to the set of all finite trees, then the Graph Minor Theorem becomes much easier and was first proved by Kruskal [22]. He showed that for trees even a stronger statement is true: the finite trees are well-quasi-ordered by the *topological* minor relation. This had been conjectured by Vázsonyi in the late 1930's. (A graph H is a *topological minor* of a graph G if G contains a subdivision of H.) Nash-Williams [33] generalized Kruskal's theorem to infinite trees. The proof of this result (for which Nash-Williams introduced the stronger notion of a better-quasi-ordering) is much harder than that of the finite case. In Chapter 6 we combine ideas of several authors into a simpler and considerably shorter proof [24] of Nash-Williams's result.

Acknowledgements

I am very grateful to Reinhard Diestel for supervising my work. I am indebted to the Studienstiftung for enabling me to write this thesis by providing me with a scholarship. I thank Deryk for his love and support.

Chapter 1

Infinite cycles

1.1 Introduction

One of the basic and well-known facts about finite graphs is that their *fundamental* cycles C_e (those consisting of a chord $e = xy$ on some fixed spanning tree T together with the path xTy joining the endvertices of e in T) generate their entire cycle space: every cycle of the graph can be written as a sum mod 2 of fundamental cycles. Richter [39] asked if and how this fact might generalize to locally finite infinite graphs. Our first and main aim is to show that this question, if viewed in the right way, admits a surprisingly elegant positive answer. As a spin-off, we obtain infinite generalizations of most of the other basic facts concerning the cycle space of a finite graph too. We shall also consider graphs that may have infinite degrees, in which case the situation turns out to be more complicated.

Of course, the finite fundamental cycle theorem transfers verbatim to infinite graphs as long as we consider only the usual finite cycles, and stick to the usual definition of the cycle space as the subspace of the edge space generated by these cycles. However, there is also a very natural notion of an infinite cycle, and the above question becomes interesting when these are admitted too—especially, since fundamental cycles continue to be finite.

Before we make all this precise, let us look at an informal example. Let L be the 2-way infinite ladder (viewed as a 1-complex) and compactify it by adding two points ω, ω' at infinity, one for each end of the ladder (Fig. 1.1).

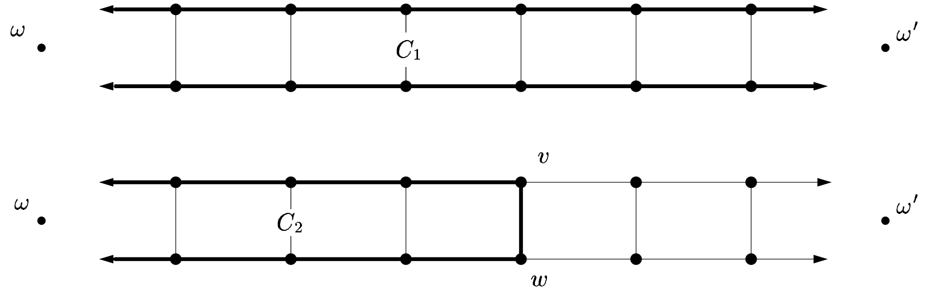

Figure 1.1: Two infinite circles in the double ladder plus ends

Let a *circle* in the resulting topological space $\overline{L}$ be any homeomorphic image of the unit circle in the Euclidean plane. Then every cycle of L is a circle in $\overline{L}$, but there are more circles than these. For example, the two sides of the ladder (each a 2-way infinite path) form a circle C_1 together with the points ω and ω', and for every rung vw the two horizontal 1-way infinite paths from v or w towards ω form a circle C_2 together with ω and the edge vw. Both these circles contain infinitely many edges, and they are determined by these edges as the closure of their union.

Now consider a spanning tree T of L. If T consists of the bottom side of L and all the rungs, then every edge e of the top side of L induces a fundamental cycle C_e. The sum (mod 2) of all these fundamental cycles is precisely the edge set of C_1, the set of horizontal edges of L. Similarly, the edge set of C_2 is the sum all the fundamental cycles C_e with e left of v (Fig. 1.2).

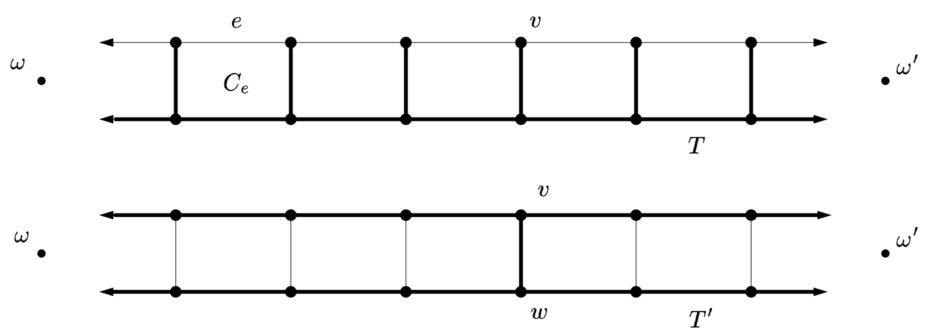

Figure 1.2: Two spanning trees of the double ladder

However, for the spanning tree T' consisting of the two sides of L and the one rung vw, neither C_1 nor C_2 can be expressed as a sum of fundamental cycles. Indeed, as every fundamental cycle contains the edge vw, any sum of infinitely many fundamental cycles will be ill-defined: it is simply not clear whether the edge vw should belong to this sum or not.

So even this simple example shows that our task is interesting: while it is possible and natural to extend the usual cycle space of a finite graph to infinite graphs in a way that allows for both infinite (topological) cycles and infinite sums generating such cycles, the answer to the question of whether all infinite cycles and their sums can be generated from fundamental cycles is by no means clear and will, among other things, depend on the spanning tree considered.

Here is an overview of the layout of this chapter and its main results. In Section 1.2 we identify some minimum requirements which any topology on an infinite graph with its ends—in the ladder example, these are the points ω and ω'— should satisfy in order to reflect our intuitive geometric picture of ends as distinct points at infinity. We then define the cycle space of an infinite graph more formally.

In Section 1.3 we introduce *end-respecting* spanning trees. It turns out that, in a locally finite graph, these are precisely the spanning trees for which infinite sums of fundamental cycles are always well-defined.

In Section 1.4 we consider the question of how best to choose the topology on an infinite graph with ends to obtain the most natural notion of a circle and the strongest possible version of our theorem. For locally finite graphs we prove that this topology is essentially unique.

In Section 1.5 we prove our main result for locally finite graphs and the topology chosen in Section 1.4: for precisely the spanning trees identified in Section 1.3, every infinite cycle of a locally finite graph is the (infinite) sum of fundamental cycles, and so are the other elements of its cycle space.

In Section 1.6 we ask to what extent this result depends on the concrete topology assumed. We find an abstract condition on the topology of a locally finite graph such that our theorem holds, for any end-respecting spanning tree, if and only if this condition is met.

In Section 1.7, we show that, for any such topology and any end-respecting spanning tree, a subgraph of a locally finite graph lies in its cycle space if and only if it meets every finite cut in an even number of edges. As a corollary, we obtain an extension of Nash-Williams's theorem that a graph is an edge-disjoint union of (finite) cycles if and only if all its cuts are even or infinite.

In the remainder of this chapter we then restrict ourselves to the topology chosen in Section 1.4 and consider graphs that may have infinite degrees. We shall characterize the spanning trees whose fundamental cycles generate every cycle (Section 1.8) or the entire cycle space (Section 1.9).

In Section 1.10 we prove an infinite analogue of the fact that every element of the cycle space of a finite graph is the edge-disjoint union of cycles.

We conclude the chapter by mentioning an open problem (Section 1.11).

1.2 Basic facts and concepts

In this chapter, we shall freely view a graph either as a combinatorial object or as the topological space of a 1-complex. (So every edge is homeomorphic to the real interval $[0,1]$, the basic open sets around an inner point being just the open intervals on the edge. The basic open neighbourhoods of a vertex x are the unions of half-open intervals $[x,z)$, one from every edge $[x,y]$ at x; note that we do not require local finiteness here.)

We shall frequently use the following well-known fact [3, Thm. 3.7]:

Lemma 1.1 *Every continuous injective map from a compact space X to a Hausdorff space Y is a topological embedding, i.e. a homeomorphism between X and its image in Y under the subspace topology.*

A homeomorphic image (in the subspace topology) of the unit interval in a topological space X will be called an *arc* in X; a homeomorphic image of the unit circle in X is a *circle* in X.

When A is an arc in X, we denote the set of all inner points of A by $\mathring{A}$. Similarly, when E is a set of edges, we write $\mathring{E}$ for the set of all inner points of edges in E. The following lemma can be proved by elementary topological arguments.

Lemma 1.2 *Every arc in G between two vertices is a graph-theoretical path. If X is an open subset of G, then the set of points in X that can be reached by an arc in X from a fixed point $x \in X$ is open. The topological components of X coincide with its arc-connected components.*

Given a spanning tree T in a graph G, every edge $e \in E(G) \setminus E(T)$ is a *chord* of T, and the unique cycle C_e in $T + e$ is a *fundamental cycle* with respect to T. A rooted spanning tree T of G is *normal* if the endvertices of every edge of G are comparable in the tree order induced by T. Countable connected graphs have normal spanning trees, but not all uncountable ones do; see [14] for details. We will use the following simple lemma, a proof can be found in [13].

Lemma 1.3 *Let $x_1, x_2 \in V(G)$, and let T be a normal spanning tree of G. For $i = 1, 2$ let P_i denote the path in T joining x_i to the root of T. Then $V(P_1) \cap V(P_2)$ separates x_1 from x_2 in G.*

We refer to 1-way infinite paths as *rays*, to 2-way infinite paths as *double rays*, and to the subrays of rays or double rays as their *tails*. If we consider two rays in a graph G as *equivalent* if no finite set of vertices separates them in G, then the equivalence classes of rays are known as the *ends* of G. (The ladder, for example, has two ends, the grid has one, and the binary tree has continuum many; see [7] for more background.) We shall write $\overline{G}$ for the union of G (viewed as a space, i.e. a set of points) and the set of its ends.

We shall consider various topologies on $\overline{G}$ in this chapter. But they will all satisfy the following two minimum requirements, without which we feel the resulting notion of a circle would seem unnatural and contrived.

The topology on $\overline{G}$ is Hausdorff, and the subspace topology which it induces on G is the topology of G when viewed as a 1-complex. (1)

Moreover, every ray should converge to the end it belongs to:

If $R \subseteq G$ is a ray and ω is the end of G containing R, then every neighbourhood of ω contains a tail of R. (2)

Together, conditions (1) and (2) imply that a subset of G is open in $\overline{G}$ if and only if it is open in G; we shall use this fact freely throughout the chapter. The following lemma, which summarizes some properties of arcs and circles in topologies satisfying (1) and (2), can be proved by elementary (though not completely trivial) topological arguments.

Lemma 1.4 *Let $\overline{G}$ be endowed with a topology satisfying (1) and (2). Then every arc A in $\overline{G}$ whose endpoints are vertices or ends, and every circle C in $\overline{G}$, includes every edge of G of which it contains an inner point. If v is a vertex in $\mathring{A}$ (respectively on C), then A (respectively C) contains precisely two edges of G at v.*

Thus in particular, every circle in $\overline{G}$ 'has' a unique set of edges, and we may define a *cycle* to be the subgraph of G consisting of all edges contained in a given circle in $\overline{G}$ (and the vertices incident with these edges). Note that these include the usual finite cycles in G, and in particular its fundamental cycles (with respect to any given spanning tree).

Another condition on the topology on $\overline{G}$ that seems natural in a context where the circles of a graph are to be represented by subgraphs consisting of their edges is that every circle is uniquely determined by its edges, as the closure of their union. Equivalently:

$$\textit{For every circle } C \subseteq \overline{G}, \textit{ the set } C \cap G \textit{ is dense in } C. \tag{3}$$

Although we shall not formally require (3), the topologies we shall consider for $\overline{G}$ will turn out to satisfy this condition, too. However, (3) does not follow from (1) and (2): in Section 1.4 we shall construct a graph with a topology satisfying (1) and (2) that contains a circle consisting entirely of ends.

A family $(G_i)_{i \in I}$ of subgraphs of a graph G will be called *edge-thin* if no edge of G lies in G_i for infinitely many i, and the *sum* $\sum_{i \in I} G_i$ of this family is the subgraph of G consisting of all edges that lie in G_i for an odd number of indices i (and the vertices incident with these edges).

Based on the concept of a cycle and our definition of 'sum', we may now define the *cycle space* $\mathcal{C}(G)$ of a locally finite graph G as the set of sums of edge-thin families of cycles. If G is finite, then this definition is compatible with the standard one (except that we now consider subgraphs of G rather than edge sets).

However, in order to define the cycle space for graphs which may have infinite degrees, we shall also have to take account of multiplicities of vertices if we want at least some spanning trees to exist whose fundamental cycles generate the cycle space. Indeed, let G be the graph obtained from two distinct vertices v and w by adding new vertices $x_1, x_2 \ldots$ and joining them to both v and w. Then the path $P := vx_1w$ is a well-defined sum of finite cycles according to the above definition (and hence an element of the cycle space), but there is no spanning tree T of G whose fundamental cycles sum to P: any such sum would consist of infinitely many fundamental cycles each containing v and w, and so the two edges of the path vTw would lie in infinitely many summands (contradiction). Hence there is no spanning tree of G for which the fundamental cycles generate its cycle space.

To overcome this problem we sharpen the requirements on the sums making up the cycle space, as follows. Call a family $(G_i)_{i \in I}$ of subgraphs of a graph G *thin* if no vertex of G lies in G_i for infinitely many i. So in particular, every thin family is edge-thin. Let the *sum* $\sum_{i \in I} G_i$ of a thin family be the subgraph of G consisting of all edges that lie in G_i for an odd number of indices i (and the vertices incident with these edges), and let the *cycle space* $\mathcal{C}(G)$ of G be the set of all sums of thin families of cycles. If G is locally finite, then the edge-thin families of cycles are precisely the thin families (since by Lemma 1.4 every cycle containing a vertex v must also contain some edge incident with v), and so our definition reduces to that given earlier.

Clearly $\mathcal{C}(G)$ is closed under finite sums, and as a consequence of our main results (and hence assuming a concrete topology for $\overline{G}$) we shall see later that

$\mathcal{C}(G)$ is closed also under taking infinite sums. This does not appear to be obvious from the definition, and we have not pursued the question of whether the definition implies it (independently of the topology assumed).

1.3 Choosing the spanning tree for a locally finite graph

Our ladder example seemed to suggest that choosing the right spanning tree might be an essential and difficult part of our problem. For locally finite graphs however, this is not the case: as we shall see, there is a canonical kind of spanning tree that will always do the job, and none other will. Before we define these spanning trees, let us recall a standard lemma about infinite graphs; the proof is not difficult and is included in [8, Lemma 1.2].

Lemma 1.5 *Let U be an infinite set of vertices in a connected graph G. Then G contains either a ray R with infinitely many disjoint U–R paths or a subdivided star with infinitely many leaves in U.*

If H is a subgraph of G, then clearly every end ω of H is a subset of a unique end ω' of G. The map $\pi_{HG} : \overline{H} \to \overline{G}$ which sends every end ω of H to this end ω' of G and which is the identity on H is called the *canonical projection of $\overline{H}$ to $\overline{G}$*.

A spanning tree T of G is *end-faithful* if the canonical projection π_{TG} is bijective, i.e. if every end of G contains rays from exactly one end of T. A spanning tree T of G is *end-respecting* if the canonical projection π_{TG} is injective. Using Lemma 1.5 it is easy to show that for locally finite G every end-respecting spanning tree is even end-faithful, but this is not true in general. Note that every connected countable graph has an end-faithful spanning tree; for example, the *normal* spanning trees defined in Section 1.2 are end-faithful. See [7] for further details.

The following observation shows that we shall want to restrict our attention to end-respecting spanning trees: any other spanning tree T would always contain a non-empty cycle, which is not only counter-intuitive but would also put an end to our hopes of showing that all cycles are sums of fundamental cycles. (Clearly, in any such sum each fundamental cycle present could be taken to occur exactly once, but then the sum would contain its chord and hence not lie in T.)

Lemma 1.6 *Let T be a spanning tree of a graph G, and assume that T contains no non-empty cycle of G. Then T is end-respecting.*

Proof. Suppose T contains two rays R_1, R_2 from a common end ω of G which are inequivalent in T. Then these rays can be chosen so as to meet precisely in their common first vertex. It is now straightforward to show that, by (1) and (2), $R_1 \cup R_2 \cup \{\omega\}$ is a circle in $\overline{G}$, and so $R_1 \cup R_2 \subseteq T$ is a cycle. □

If G is locally finite, then the converse of Lemma 1.6 will follow, for a concrete topology on $\overline{G}$ we shall consider, from Theorem 1.13 which says that whenever T is end-respecting in G all the cycles in G are sums of fundamental cycles (and hence, in particular, not contained in T). Conditions (1) and (2) do not imply the converse of Lemma 1.6 for arbitrary topologies, though: in Section 1.4 we shall construct a topology for the binary tree which satisfies (1) and (2), but under which a double ray in the tree occurs as a cycle. A more complicated example at the end of Section 1.6 will show that the converse of Lemma 1.6 does not even follow from (1), (2) and (3).

For locally finite graphs, end-respecting spanning trees have the pleasant property that every sum of fundamental cycles is well-defined:

Lemma 1.7 *Let T be an end-respecting spanning tree of a locally finite graph G. Then the fundamental cycles of G with respect to T form a thin family.*

Proof. Suppose that the lemma is false. As every egde-thin family of (finite) cycles of a locally finite graph G is thin, it follows that there are infinitely many fundamental cycles $C_1, C_2, \ldots$ all containing the same edge $e = xy$. Then xy is an edge of T; let T_x and T_y be the components of $T - e$ containing x and y, respectively. For $i = 1, 2, \ldots$ let $e_i = x_i y_i$ be the edge of C_i not on T; since $e \in C_i$, we may assume that $x_i \in T_x$ and $y_i \in T_y$.

Applying Lemma 1.5 to T_x with $U = \{x_1, x_2, \ldots\}$, we obtain a ray $R_x \subseteq T_x$ and an infinite index set $I \subseteq \mathbb{N}$ such that the paths $P_i \subseteq T_x$ from x_i to R_x are disjoint for different $i \in I$. Applying Lemma 1.5 to T_y with $U = \{y_i \mid i \in I\}$, we likewise obtain a ray $R_y \subseteq T_y$ and an infinite index set $I' \subseteq I$ such that the paths $Q_i \subseteq T_y$ from y_i to R_y are disjoint for different $i \in I'$. As the rays R_x and R_y are disjoint, they belong to different ends of T. But each of the paths $P_i e_i Q_i$ with $i \in I'$ links R_x to R_y in G, and these are infinitely many disjoint paths. Therefore R_x and R_y belong to a common end of G, so T is not end-respecting. $\square$

We remark that the converse of Lemma 1.7 holds too: if T is not end-respecting, we can always find a family of fundamental cycles that is not (edge-) thin.

1.4 Choosing the topology on $\overline{G}$

Since the meaning of our intended result (that the fundamental cycles of a graph generate its cycle space) depends on the notion of a circle and hence on the topology considered for $\overline{G}$, we have to fix some such topology at some point. But which topology should we choose? In Section 1.2 we laid down two minimum requirements for any topology on $\overline{G}$ that we might consider as natural, conditions (1) and (2). However, these two conditions do not determine the topology on $\overline{G}$.

For example, the following topology satisfies (1) and (2) and would not seem unnatural. Given an end ω and a finite set S of vertices of G, there is exactly one component $C = C_G(S, \omega)$ of $G - S$ which contains a tail of every ray in

ω. We say that ω *belongs* to C. Writing $E_G(S,\omega)$ for the set of all the C–S edges in G, let us consider the topology on $\overline{G}$ that is generated by the open sets of G (as a 1-complex) and all sets of the form $\{\omega\} \cup C_G(S,\omega) \cup \mathring{E}'_G(S,\omega)$, where $\mathring{E}'_G(S,\omega)$ is any union of half-edges $(x,y] \subset e$, one for every $e \in E_G(S,\omega)$, with $x \in \mathring{e}$ and $y \in C$. Then the circles in this topology resemble those of our ladder example:

Proposition 1.8 *Let G be any infinite graph. Under the above topology, every circle in $\overline{G}$ is either a finite cycle or the union of finitely many double rays with their ends.*

Proposition 1.8 is not difficult to prove, and it easily implies that every cycle is the sum of finite cycles. If G is locally finite then this implies by Lemma 1.7 that, given any end-respecting spanning tree, the fundamental cycles do indeed generate the cycle space.

Although the topology for $\overline{G}$ on which Proposition 1.8 is based may be a natural one to consider, it does not yield the strongest possible theorem. For note that if TOP_1 and TOP_2 are topologies on $\overline{G}$ such that TOP_2 is Hausdorff and coarser than TOP_1, then every TOP_1-circle is also a TOP_2-circle (cf. Lemma 1.1). Thus reducing the collection of open sets increases the set of circles, and so we can strengthen our theorem by proving it for a coarser topology. Our next aim, therefore, is to introduce a topology on $\overline{G}$ that is coarser than that considered above. For locally finite G, this topology will turn out to be coarsest possible with (1), and therefore yield the best possible result.

Given an end ω and a finite set S of vertices of G, let $\overline{C}_G(S,\omega)$ denote the union of $C := C_G(S,\omega)$ with the set of all ends belonging to C. There is an obvious correspondence between these ends of G and those of C, and we shall not normally distinguish between them. (Thus, $\overline{C}$ will be treated as a subset of $\overline{G}$ when this simplifies the notation.) Let TOP denote the topology on $\overline{G}$ generated by the open sets of the 1-complex G and all sets of the form

$$\widehat{C}_G(S,\omega) := \overline{C}_G(S,\omega) \cup \mathring{E}'_G(S,\omega),$$

where again $\mathring{E}'_G(S,\omega)$ is any union of half-edges $(x,y] \subset e$, one for every $e \in E_G(S,\omega)$, with $x \in \mathring{e}$ and $y \in C$. So for each end ω, the sets $\widehat{C}_G(S,\omega)$ with S varying over the finite subsets of $V(G)$ are the basic open neighbourhoods of ω. The topology which TOP induces on the end space $\overline{G} \setminus G$ of G is the standard topology there as studied in the literature.

The following observation is not difficult to prove; see e.g. [7].

Lemma 1.9 *If G is connected and locally finite, then $\overline{G}$ is compact in* TOP.

By Lemma 1.1, the topology of a compact Hausdorff space cannot be made coarser without loss of the Hausdorff property. So for G locally finite, there is no Hausdorff topology on $\overline{G}$ which is strictly coarser than TOP, and in this sense proving the theorem for TOP (as we shall do in the next section) will be best possible.

Lemma 1.10 *For every infinite graph G, the topology* TOP *satisfies (1), (2) and (3).*

Proof. Conditions (1) and (2) hold trivially. Suppose there is a circle C in $\overline{G}$ such that $C \cap G$ is not dense in C. Then some point on C has a neighbourhood N in C that consists entirely of ends. We may assume that $N \not\supseteq C$, and that $N = O \cap C$ for some basic open set O in $\overline{G}$. Then $O = \widehat{D}$ for some component D of $G - S$ with $S \subseteq V(G)$ finite, and $N = (\overline{D} \setminus G) \cap C$ is the intersection of two closed sets in $\overline{G}$. So N is closed in $\overline{G}$ and hence in C. Since $N = O \cap C$ is also open in C, the homeomorphism between C and the unit circle takes N to an open and closed proper subset there, contradicting its connectedness. □

Note that if $\overline{G}$ is endowed with TOP, then Lemmas 1.4 and 1.10 together imply that every infinite cycle C is a disjoint union of double rays, and its defining circle is the closure of C in $\overline{G}$.

Let us return to the case when G is locally finite. Since TOP is best possible for our purposes among all topologies comparable with it, the question arises whether there are topologies which satisfy our minimum requirements (1) and (2) but are incomparable with TOP. In the remainder of this section, we first construct such an example. However we then show that a slight strengthening of (2) will rule out such (pathological) examples and imply that every topology on $\overline{G}$ satisfying this condition and (1) is indeed comparable with TOP, making our theorem best possible also in a more global sense.

So we are looking for a locally finite graph G with a topology that satisfies (1) and (2) but is incomparable with TOP. Since every Hausdorff topology comparable with TOP refines TOP and hence inherits (3) from it (Lemma 1.1), it suffices to construct a topology for $\overline{G}$ that violates (3). As a spin-off, we thus obtain that conditions (1) and (2) do not imply (3):

Proposition 1.11 *There exists a locally finite graph G with a topology that satisfies (1) and (2) but not (3), and hence is incomparable with* TOP.

Proof. Our graph G will be the infinite binary tree T. We label its vertices with finite 0–1 sequences in the obvious way: the root (which is considered as the *lowest* point in T) is labelled with the empty sequence, and if a vertex has label ℓ then its two successors are labelled $\ell 0$ and $\ell 1$. Then the rays from the root (and hence the ends of T) correspond bijectively to the infinite 0–1 sequences and may be thought of as elements of the real interval $[0, 1]$ in their binary expansion. Let J be the set of all rationals in $(0, 1)$ with a finite binary expansion, and let $J' := [0, 1] \setminus J$. Then each $r \in J'$ comes from exactly one ray R_r, while every $q \in J$ comes from two: a ray R_q labelled eventually 0, and a ray R'_q labelled eventually 1. (For example, the rays $1011000\ldots$ and $10101111\ldots$ both correspond to $11/16$.) Let ω_r, ω_q and ω'_q denote the ends containing R_r, R_q and R'_q, respectively. Let M' be the set of the ends ω'_q, and let M be the set of all the other ends.

We now define the topology on $\overline{T}$ so as to turn the bijection between $[0, 1]$ and M into a topological embedding. Every point in T will have the same basic open neighbourhoods as it does in T viewed as a 1-complex. The basic open neighbourhoods of an end $\omega_p \in M$ are constructed as follows. Choose an open neighbourhood I of p in $[0,1]$. For each $s \in I$ choose a point z on R_s, and let N_s be the set consisting of ω_s and all the points of T (not of $\overline{T}$) above z. Now take the union of the N_s over all $s \in I$ to be a basic open neighbourhood of ω_p. The basic open neighbourhoods of the ends $\omega'_q \in M'$ will be as small as possible given (1) and (2), consisting just of the 'open final segments' of R'_q: pick a point $z \in R'_q$, choose for every vertex t_i on R'_q above z a point z_i in the interior of the edge at t_i that does not lie on R'_q, and take the union of the set of points on R'_q above z with all the partial edges $[t_i, z_i)$ to be a basic open neighbourhood of ω'_q.

It is straightforward to check that the topology generated in this way satisfies (1) and (2). But it violates (3), since the union of M and the rays $R_0 = 000\ldots$ and $R_1 = 111\ldots$ forms a circle. □

With only a slight modification, the above example will even contain a circle consisting entirely of ends. Indeed, if we identify 0 and 1 in $[0,1]$ to form a circle (and put ω_1 in M' rather than M), then in the analogously defined topology the set M of ends becomes a circle in $\overline{T}$.

In Section 1.6 we shall construct a graph G with a topology that satisfies (3) as well as (1) and (2) but is still incomparable with Top.

The topology on the binary tree constructed in the proof of Proposition 1.11 will hardly be considered as natural. But how unnatural did it *have* to be? Our next Proposition answers this question in an unexpectedly clear-cut way: any topology that witnesses Proposition 1.11 has to violate an only slight and still pretty natural sharpening of condition (2). Or put another way, every topology that satisfies (1) and this new condition is comparable with Top.

In order to state the new condition we need a definition. A *comb* C with *back* R is obtained from a ray R and a sequence $x_1, x_2, \ldots$ of distinct vertices by adding for each $i = 1, 2, \ldots$ a (possibly trivial) x_i–R path P_i so that all the P_i are disjoint and R meets P_{i+1} after P_i. The vertices x_i will be called the *teeth* of C. When we speak of a comb in G, however, we wish to admit inner points of edges as teeth. We therefore call C a *comb in* G if C is a comb in some subdivision of G (in which every edge may be assumed to be subdivided at most once).

Clearly, condition (2) is a special case of the more general requirement that the teeth of every comb converge to the end of its back:

Every neighbourhood of an end ω contains all but finitely many of the teeth of every comb in G whose back lies in ω. (4)

Note that for locally finite graphs (4) holds in Top and (1) and (4) together imply (3). Indeed, we have the following stronger assertion:

Proposition 1.12 *Let G be a locally finite connected graph, and let* Top$'$ *be any topology on $\overline{G}$ that satisfies (1) and (4). Then* Top$'$ *is a refinement of* Top.

Proof. For a proof that the open sets of Top are open in Top$'$, it suffices to show that for every end ω of G and every finite $S \subset V(G)$ there exists a TOP'-neighbourhood of ω contained in $\overline{C}$, where $C := C_G(S,\omega)$. Suppose not, let $s_1, s_2, \dots$ be an enumeration of the vertices outside S and C, and put $S_i := S \cup \{s_1, \dots, s_i\}$. Now consider any i. By (1), the vertices of S_i together with all their incident edges in G form a compact set $\overline{S}_i = G[S_i \cup N(S_i)]$ in TOP'. So every TOP'-neighbourhood of ω contains a neighbourhood that avoids $\overline{S}_i$, and hence meets one of the components of $G - S_i$ other than C. (Recall that no such neighbourhood is contained in $\overline{C}$, and apply (2).) As $G - S_i$ has only finitely many components, this implies that there is a component $D_i \neq C$ of $G - S_i$ met by every TOP'-neighbourhood of ω.

Choosing these components $D_1, D_2, \dots$ in turn, we can ensure that $D_1 \supset D_2 \supset \dots$. Now pick a sequence of distinct points $x_i \in D_i$ so that no two of them lie inside the same edge. By Lemma 1.5 (applied to the subdivision of G obtained by making every x_i a vertex) there is a comb C in G with teeth among the x_i and back R, say. Since each D_i contains all but finitely many of the x_j and is separated from the rest of G by the finite set S_i, this ray R has a tail in every D_i. But any two rays with this property are equivalent (because the S_i eventually contain any finite separator of D_1), so the end ω' of R is independent of the choice of $x_1, x_2, \dots$ but depends only on the sequence $D_1 \supset D_2 \supset \dots$.

Since TOP' is Hausdorff, there are disjoint TOP'-neighbourhoods N of ω and N' of ω'. As N meets every D_i but no vertex or edge lies in more than finitely many D_i, we may choose all our points x_i inside N and hence outside N'. Hence our comb C contradicts (4). □

1.5 The generating theorem for locally finite graphs

Let us now prove our main theorem for locally finite graphs G: under TOP, the fundamental cycles with respect to any end-respecting spanning tree generate the entire cycle space of G. Although this result is a special case of Theorem 1.27 below, we prove it separately, as the proof is much simpler than that of the general case.

Theorem 1.13 *Let G be a locally finite connected graph, let $\overline{G}$ be endowed with* TOP*, and let T be any spanning tree of G. Then the following are equivalent:*

(i) *Every cycle of G is a sum of fundamental cycles.*

(ii) *Every element of the cycle space $\mathcal{C}(G)$ of G is a sum of fundamental cycles.*

(iii) *T is end-respecting.*

Proof. Clearly, (ii) implies (i). Lemma 1.6 and the remark preceding it show that (i) implies (iii). To show that (iii) implies (ii) it suffices to prove the following:

If T is end-respecting, then every cycle C of G is equal to the sum of all the fundamental cycles C_e with $e \in E(C) \setminus E(T)$. (∗)

Indeed, every element Z of $\mathcal{C}(G)$ is by definition the sum of a thin family $\mathcal{F}$ of cycles. By $(*)$, each $C \in \mathcal{F}$ is the sum of fundamental cycles C_e with $e \in E(C)$. Since $\mathcal{F}$ is thin, none of these edges e lies on more than finitely many cycles in $\mathcal{F}$, so all these fundamental cycles together form a family in which none of them occurs infinitely often. By Lemma 1.7 this is again a thin family, so it has a well-defined sum. Clearly, this sum equals Z.

To prove $(*)$ let C be given, and let C' be its defining circle in $\overline{G}$. (Since Top satisfies (3), C' is the closure of C, so in particular C' is uniquely determined.) Let C'' be the unit circle, and pick a homeomorphism $\sigma : C'' \to C'$. We have to show that an edge f of G lies in C if and only if it lies in an odd number of the cycles C_e in $(*)$. This is clear when f is a chord of T, as in that case f lies on C_e only if $f = e$.

So consider an edge $f \in T$. Let G_1 and G_2 denote the subgraphs of G induced by the two components of $T - f$, and let E_f be the set of G_1–G_2 edges in G (including f). Note that the edges $e \neq f$ in E_f are precisely the chords e of T with $f \in E(C_e)$. Since the family of these C_e is thin by Lemma 1.7, the set E_f is finite. Hence for $i = 1, 2$, G_i is a component of $G - S$ for the finite set $S = N(G_i)$ of its neighbours outside, every set of the form $\widehat{G}_i$ is open in Top, and $\overline{G}_i = \widehat{G}_i \setminus \mathring{E}_f$ is open in $\overline{G} \setminus \mathring{E}_f$.

As σ is a homeomorphism, $C'' \setminus \sigma^{-1}(\mathring{E}_f \cap C')$ consists of finitely many intervals, $I_1, \dots, I_k$ say. Each $\sigma(I_i)$ is a connected subset of $C' \setminus \mathring{E}_f$ and hence cannot meet both of the disjoint open subsets $\overline{G}_1$ and $\overline{G}_2$ of $\overline{G} \setminus \mathring{E}_f$. Our circle C' therefore contains an even number of edges from E_f. Hence, C contains f if and only if it contains an odd number of other edges from E_f, which it does if and only if f lies on an odd number of the cycles C_e with $e \in E(C)$ and hence in the sum of $(*)$. $\square$

By the argument that showed $(*)$ to be sufficient for a proof of Theorem 1.13, the theorem and Lemma 1.7 imply the following:

Corollary 1.14 *If G is locally finite, then its cycle space in* Top *is closed under taking sums.*

We conclude this section with an example illustrating how unlike our initial ladder examples the cycles covered by Theorem 1.13 can become. Adding just a few edges to the binary tree, we obtain a graph in which all the fundamental cycles sum up to a single circle containing continuum many ends and a 'dense' set of double rays (so that between any two double rays there lies another).

Consider again the infinite binary tree T, and let J, J', R_q, R'_q, ω_q and ω'_q be defined as in the proof of Proposition 1.11. Let D_0 be the double ray formed by the rays $R_0 = 000\dots$ and $R_1 = 111\dots$. For every $q \in J$ add an edge $e_q = t_q t'_q$ between disjoint tails of R_q and R'_q, so that if ℓ is the label of the last common vertex of R_q and R'_q, the vertex t_q is labelled $\ell 011$ and t'_q is labelled $\ell 100$. Then the double rays $D_q = (t_q R_q \cup t'_q R'_q) + e_q$ are disjoint from D_0 and from each other, and T is an end-respecting spanning tree of the resulting graph G.

Let us show that the union C of all the D_q (for $q \in J \cup \{0\}$) and the set of ends of G is a circle in Top. Denote the unit circle by C', and let I_0 be a closed interval on C'. Let $\sigma : I_0 \mapsto D_0 \cup \{\omega_0, \omega_1\}$ be a homeomorphism, and put $x_0 := \sigma^{-1}(\omega_0)$ and $x_1 := \sigma^{-1}(\omega_1)$. Our aim is to extend σ to a homeomorphism between C' and C.

Let $I := C' \setminus \mathring{I}_0$, and think of x_0 as the *left* and x_1 as the *right* endpoint of I. Assign to the points $q \in J$ disjoint closed subintervals $I_q = [x_q, x'_q]$ of $\mathring{I}$, so that I_{q_1} lies left of I_{q_2} whenever $q_1 < q_2$, and I is the closure of $U := \bigcup_q I_q$. (For example, this could be done inductively in ω steps.) Then the points of $\mathring{I} \setminus U$ correspond bijectively to the points in $J' \cap (0,1)$ of the completion $[0,1]$ of J; let x_r be the point of $\mathring{I} \setminus U$ corresponding to $r \in J' \cap (0,1)$. Finally, let $\sigma : C' \to C$ map each I_q continuously onto $D_q \cup \{\omega_q, \omega'_q\}$ so that $\sigma(x_q) = \omega_q$ and $\sigma(x'_q) = \omega'_q$, and put $\sigma(x_r) = \omega_r$ for all $r \in J' \cap (0,1)$. Then $\sigma : C' \to C$ is a homeomorphism, so C is indeed a circle.

Theorem 1.13 now says that all the fundamental cycles in G together sum to an infinite cycle: the cycle defined by our circle C. Once observed, this can also easily be checked directly.

1.6 A topological condition equivalent to the generating theorem for locally finite graphs

Let G be a locally finite connected graph. In this section we shall identify a condition just in terms of the topology on $\overline{G}$ that is equivalent to the validity of Theorem 1.13 (ii). This has some interesting consequences.

First, since the elements of the cycle space of G are the same—the sums of fundamental cycles of any end-respecting spanning tree—whenever this condition holds, we find that the cycle space is independent of the topology used as long as it satisfies this condition. In particular, refinements of Top (which will all satisfy the condition) may have fewer cycles than Top (recall Prop. 1.8) but will have the same cycle space. Second, since the new condition will not follow from (1), (2) and (3), we also obtain a negative answer to the question of whether these three (rather natural) conditions alone can guarantee the validity of Theorem 1.13.

We will need the following lemma from elementary topology [20, p. 208]. A continuous image of $[0,1]$ in a topological space X is a (topological) *path* in X; the images of 0 and 1 are its *endpoints*. By Lemma 1.1, a path in a Hausdorff space is an arc if and only if the corresponding map $[0,1] \to X$ is injective.

Lemma 1.15 *Every path with distinct endpoints x, y in a Hausdorff space X contains an arc in X between x and y.*

As always, we consider only topologies on $\overline{G}$ that satisfy our two minimum requirements (1) and (2). Then if E is any finite set of edges and C is a component of $G - E$, the subspace $\overline{C}$ of $\overline{G}$ is path-connected. Our new condition says that these $\overline{C}$ are in fact the whole path components of $\overline{G} \setminus \mathring{E}$:

Whenever E is a finite set of edges of G, every path component of the topological space $\overline{G} \setminus \mathring{E}$ is of the form $\overline{C}$, for some component C of the graph $G - E$. (5)

Note that (5) implies (3). Indeed if (3) fails, then $\overline{G}$ contains an arc $[\omega, \omega']$ consisting entirely of ends. Let S be a finite set of vertices separating a ray in ω from a ray in ω', and let E be the set of edges incident with S. Then ω and ω' do not both lie in $\overline{C}$ for any component C of $G - E$, although they do lie in the same path component of $\overline{G} \setminus \mathring{E}$.

Although Top clearly satisfies (5), it is not difficult to construct locally finite graphs with topologies that satisfy (1), (2), (3) and (5) but are incomparable with Top. (By Proposition 1.12, these topologies must violate (4).) The following result may thus be viewed as a topologically best-possible strengthening of Theorem 1.13:

Theorem 1.16 *Let G be a locally finite connected graph, let $\overline{G}$ carry any topology satisfying (1) and (2), and let T be an end-respecting spanning tree in G. Then the following two assertions are equivalent:*

(i) *$\overline{G}$ satisfies (5);*

(ii) *every element of the cycle space of G is the sum of fundamental cycles.*

Proof. The proof of Theorem 1.13 shows that (i) implies (ii). Indeed, as $\overline{G}$ satisfies (5), both $\overline{G_1}$ and $\overline{G_2}$ as considered there are path components of $\overline{G} \setminus \mathring{E}_f$ and thus again each $\sigma(I_i)$ lies in either $\overline{G_1}$ or $\overline{G_2}$. So let us prove the converse implication.

If (5) fails, then for some finite set $E \subseteq E(G)$ there are components $D_1 \neq D_2$ of $G - E$ such that $\overline{D_1}$ and $\overline{D_2}$ are contained in the same path component of $\overline{G} \setminus \mathring{E}$. By making E smaller, we may assume that D_1 and D_2 are the only components of $G - E$. Let $f_1, \ldots, f_k$ be the D_1–D_2 edges contained in T.

For each $i = 1, \ldots, k$, let E_i be the set of the edges of G between different components of $T - f_i$. Since the edges $e \neq f_i$ in E_i are precisely the chords e of T with $f_i \in C_e$, Lemma 1.7 implies that each E_i is finite.

By definition, $\overline{D_1}$ and $\overline{D_2}$ are joined in $\overline{G} \setminus \mathring{E}$ by a topological path π; since they are path-connected, we may assume that the endpoints of π are vertices, and by Lemma 1.15 we may assume that π is an arc.

Since by Lemma 1.4 an arc between two vertices includes every edge of which it contains an inner point, and since the E_i are finite, $\pi \setminus (\mathring{E}_1 \cup \cdots \cup \mathring{E}_k)$ consists of finitely many closed segments whose endpoints are vertices. One of these, π' say, is again an arc from a vertex $v_1 \in D_1$ to a vertex $v_2 \in D_2$. (For since π contains no D_1–D_2 edge, the endpoints of every missing edge lie in the same D_j.)

Pick an edge $f_i \in \{f_1, \ldots, f_k\}$ from the path $v_1 T v_2$, let P be the segment of $v_1 T v_2$ that includes f_i and meets π' only in its endpoints, and let π'' be the segment of π' between these points. Then $P \cup \pi''$ is a circle in $\overline{G}$ that contains f_i but no other edge from E_i. Its cycle is therefore not a sum of fundamental cycles, so (ii) fails as required. □

Corollary 1.17 *The cycle space of a locally finite graph G is independent of the topology chosen for $\overline{G}$, as long as the topology satisfies (1), (2) and (5). In* *particular, $\mathcal{C}(G)$ is the same for all refinements of* Top *(with (1) and (2)), and thus uniquely determined for all topologies that satisfy (1) and (4).*

Proof. For the second statement, let us first verify (5) for an arbitrary refinement Top$'$ of Top that satisfies (1) and (2). Since rays converge to their ends by (2), every $\overline{C}$ as in (5) is path-connected. So any path component D of $\overline{G} \setminus \mathring{E}$ is a disjoint union of such $\overline{C}$. But every $\overline{C}$ is open in $\overline{G} \setminus \mathring{E}$, in Top and hence also in Top$'$. Hence D, being connected, consists of a single $\overline{C}$.

For the last statement, recall that every topology satisfying (1) and (4) is a refinement of Top (Prop. 1.12). $\square$

In the remainder of the section we construct an example which shows that (5) does not follow from (1), (2) and (3). Hence, by Theorem 1.16, these three conditions alone cannot guarantee the validity of Theorem 1.13.

Consider again the infinite binary tree T, and let J, J', R_q, R'_q, ω_q and ω'_q be as in the proof of Proposition 1.11. Add a new double ray D with ends $\tau \neq \tau'$ which meets T exactly in its root t. Moreover for each $q \in J$ add a new double ray D_q with ends $\tau_q \neq \tau'_q$, together with an edge $e_q = v_q v'_q$ joining a vertex v_q on D_q to a vertex v'_q on D; choose these edges e_q independent. Denote the graph thus obtained by G.

Let us now define the topology on $\overline{G}$. G itself will carry the topology of a 1-complex. The basic open neighbourhoods of an end ν of the form τ, τ', ω_q or ω'_q with $q \in J$ will consist of an 'open final segment' of the ray $R \in \nu$ starting at t: pick a point z on R, as well as an inner point z_e of every edge $e \notin R$ incident with a vertex $v \in \mathring{z}R$; then take the union of $\mathring{z}R$ with all the partial edges $[v, z_e) \subseteq e$ to be a basic open neighbourhood of ν. The basic open neighbourhoods of an end ω_r for $r \in J'$ are constructed as follows. Choose an open neighbourhood I of r in $[0,1]$. For each $s \in I \cap J'$ choose a point z on R_s, and let N_s be the set consisting of ω_s and all the points of T above z. For each $q \in I \cap J$ pick an inner point z_q of e_q. Take the union of the N_s over all $s \in I \cap J'$ together with the union of $\{\tau_q, \tau'_q\} \cup D_q \cup [v_q, z_q)$ over all $q \in I \cap J$ to be a basic open neighbourhood $N(I)$ of ω_r. To construct a basic open neighbourhood of an end τ_q (respectively τ'_q) for $q \in J$, choose an open neighbourhood I of q in $(0, q]$ (respectively $[q, 1)$) and take the union of $N(I)$ (defined as before) together with τ_q (respectively τ'_q) and an open final segment of the subray of $D_q - v_q$ contained in τ_q (respectively τ'_q) to be a basic open set.

Clearly, the topology generated in this way satisfies (1) and (2). As in the proof of Lemma 1.10 one can show (3). (Indeed, the open set N considered there can again be chosen so that it is also closed.) Furthermore, using similar arguments as in the example in Section 1.5 one can show that all the ends ω_r ($r \in J'$) and all the sets $\{\tau_q\} \cup D_q \cup \{\tau'_q\}$ with $q \in J$ together form an arc π whose endpoints belong to different components of $G - t$. Hence, this topology does not satisfy (5).

Finally, since G is a tree and π forms a circle together with the rays R_0 and R_1, our example shows also that the converse of Lemma 1.6 can fail even for topologies satisfying (1), (2) and (3).

1.7 Cuts and cycle decompositions

Let G be a locally finite connected graph, and let $\overline{G}$ carry any topology (such as TOP) for which G satisfies the generating theorem, ie. which satisfies (1), (2) and (5).

Our next theorem characterizes the elements of $\mathcal{C}(G)$ in terms of the (edge) cuts in G. Recall that a *cut* in G is the set of all the edges of G between the two classes of some bipartition of $V(G)$. When G is finite, the elements of its cycle space are precisely those sets of edges that are 'orthogonal' to every cut in G, ie. contain an even number of edges from every cut [6]. For arbitrary locally finite G, this generalizes as follows:

Theorem 1.18 *Let G be a locally finite connected graph, let $\overline{G}$ be endowed with any topology satisfying (1), (2) and (5), and let $H \subseteq G$ be any subgraph without isolated vertices. Then the following statements are equivalent:*

(i) *$H \in \mathcal{C}(G)$;*

(ii) *$|E(H) \cap F|$ is even for every finite cut F of G.*

Proof. The fact that every cycle in G meets every finite cut in an even number of edges is proved as in the proof of Theorem 1.13, using (5) instead of the definition of TOP if desired. Since sums (mod 2) of even sets are even, this implies (i)$\Rightarrow$(ii).

For the converse implication, assume (ii) and let T be an end-respecting spanning tree of G. We show that H is equal to the sum $Z \in \mathcal{C}(G)$ of all the fundamental cycles C_e with $e \in E(H) \setminus E(T)$. For every chord e of T in G, clearly $e \in H$ if and only if $e \in Z$. So consider an edge $f \in T$. Let E_f be the set of edges $e \neq f$ of G between the two components of $T - f$. Since T is end-respecting and $f \in C_e$ for precisely those chords e of T that lie in E_f, Lemma 1.7 implies that E_f is finite, and $f \in Z$ if and only if $|E_f \cap E(H)|$ is odd. By (ii), the latter holds if and only if $f \in H$, as desired. $\square$

A classical theorem of Nash-Williams [32] says that a connected graph (of any cardinality) is an edge-disjoint union of finite cycles if and only if each of its cuts is either infinite or even. By Theorem 1.18, the locally finite such graphs are characterized by a further property, namely the fact that they by themselves are contained in their cycle space:

Corollary 1.19 *Let G be a locally finite connected graph and let $\overline{G}$ be endowed with any topology satisfying (1), (2) and (5). Then the following statements are equivalent:*

(i) *every cut in G is either infinite or even;*

(ii) *G is an edge-disjoint union of finite cycles;*

(iii) *G is an edge-disjoint union of cycles;*

(iv) *$G \in \mathcal{C}(G)$.*

Proof. (i)⇒(ii) is Nash-Williams's theorem. (In fact, as Nash-Williams observed [32], this result is easy for countable graphs.) The implications (ii)⇒(iii) and (iii)⇒(iv) are trivial. The implication (iv)⇒(i) follows from Theorem 1.18. □

In Section 1.10 we shall extend the equivalence between (iii) and (iv) in Corollary 1.19 to arbitrary elements of the cycle space of G. This more general equivalence, which holds in arbitrary infinite graphs, is much harder than Corollary 1.19 and cannot easily be reduced to it: although the elements of the cycle space of an arbitrary infinite graph G still form locally finite subgraphs of G, they are not normally elements of the cycle space of that subgraph (just as a single infinite cycle C in G is, by itself, merely a disjoint union of double rays containing no cycle at all).

1.8 The generating theorem for cycles

Throughout the remainder of this chapter $\overline{G}$, $\overline{H}$ etc. will always be endowed with Top. In this section we characterize the spanning trees of an arbitrary graph G whose fundamental cycles generate every cycle of G.

In Section 1.5 we showed that if G is locally finite, then these are precisely its end-respecting (equivalently: end-faithful) spanning trees. In general, however, this need not be true. Consider the graph G obtained from two rays $R = x_0x_1x_2\dots$ and $Q = x_0y_1y_2\dots$ that meet only in their first point x_0 by adding the edges x_iy_i for all $i \geq 1$, and adding a new vertex z joined to all the x_i (Fig. 1.3). Then R and Q belong to the same end ω of G. Thus $R \cup Q \cup \{\omega\}$ is a circle in $\overline{G}$, and so $R \cup Q$ is a cycle in G. But if T is the spanning tree of G consisting of Q together with all edges at z, then T is end-respecting (even end-faithful), but $R \cup Q$ is not a sum of fundamental cycles: since all these contain z, any sum of them would have to be finite.

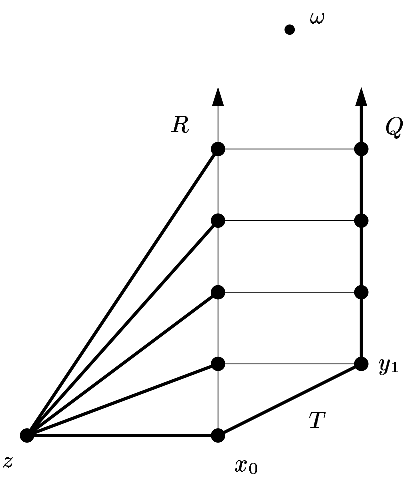

Figure 1.3: The infinite cycle $R \cup Q$ is not a sum of fundamental cycles

The above example motivates the consideration of the following subclass of the end-respective spanning trees.

Definition Given a graph G, let $\mathcal{T}(G)$ denote the class of all end-respecting spanning trees T of G which do not contain a subdivided infinite star S whose leaves lie on a ray $R \subseteq G$ such that G contains another ray R' which is equivalent to but disjoint from R.

Note that there are graphs G for which $\mathcal{T}(G)$ is empty; $K_{\aleph_1}$ and $K_{\aleph_0,\aleph_1}$ are obvious examples. On the other hand, using Lemma 1.3 one can easily show that $\mathcal{T}(G)$ contains every normal spanning tree of G. We do not know whether there are graphs G for which $\mathcal{T}(G)$ is non-empty but which have no normal spanning tree.

Theorem 1.20 *Let G be an infinite connected graph and let $\overline{G}$ be endowed with* Top. *Let T be a spanning tree of G. Then every cycle of G is the sum of fundamental cycles if and only if $T \in \mathcal{T}(G)$.*

For the proof of this theorem we first need some lemmas.

Lemma 1.21 *Given any spanning tree T of G, every finite cycle C of G is the sum of fundamental cycles. More precisely, C is equal to the sum Z of all the fundamental cycles C_e with $e \in E(C) \setminus E(T)$.*

Proof. Clearly $C + Z$ is a finite subgraph of T with all degrees even. Hence $C + Z = \emptyset$, i.e. $C = Z$. $\square$

Lemma 1.22 *Let T be a spanning tree of G, let Z be a sum of fundamental cycles, and let $\mathcal{D}$ be a set of finite cycles in $Z \cup T$. If no two elements of $\mathcal{D}$ share an edge outside T, then $\mathcal{D}$ is thin.*

Proof. Suppose that x is a vertex that lies on infinitely many cycles $D \in \mathcal{D}$. By Lemma 1.21, each of these D is a sum of fundamental cycles C_e with $e \in E(D) \setminus E(T)$, so x lies on some C_e with $e \in E(D) \setminus E(T)$. By assumption these edges e differ for different D, so x lies on infinitely many C_e. As each e lies in Z, all these C_e are among the fundamental cycles whose sum is Z (indeed, Z must be the sum of the fundamental cycles C_e with $e \in E(Z) \setminus E(T)$), which contradicts the definition of sum. $\square$

Lemma 1.23 *Let A be an arc in $\overline{G}$, and let $x \neq y$ be vertices on A. Let X be a closed subset of $\overline{G}$ which avoids the subarc of A between x and y. Then G contains an x–y path P with $P \cap X = \emptyset$.*

Proof. Let A' be the subarc of A between x and y. Choose a cover $\mathcal{N}$ of A' by basic open sets of $\overline{G}$ each avoiding X. As A' is compact, $\mathcal{N}$ contains a finite subcover of A', $\{N_1, \dots, N_k\}$ say, where we may assume that $N_\ell \cap A' \neq \emptyset$ for all ℓ.

Let us show that $H := (N_1 \cup \cdots \cup N_k) \cap G$ is a connected subspace of G. If not, then H is the union of two disjoint non-empty open subsets H_1 and H_2 of H. Since each N_ℓ is a basic open set, $N_\ell \cap G$ is connected and hence lies in either H_1 or H_2. Let U_1 be the union of all N_ℓ with $N_\ell \cap G \subseteq H_1$, and define U_2 similarly. Since two N_ℓ cannot share an end if their intersections with G are disjoint, U_1 and U_2 are disjoint. Thus A' is the union of the two disjoint non-empty open sets $A' \cap U_1$ and $A' \cap U_2$, contradicting its connectedness.

So H is connected. Lemma 1.2 together with the fact that H contains both x and y imply that H also contains a (graph-theoretical) path P between these two vertices. Clearly, P is as required. $\square$

An *orientation* of an arc A is a linear ordering of its points which is induced by a homeomorphism $\sigma : [0,1] \to A$ (i.e. if $a, b \in A$ then $a < b$ if $\sigma^{-1}(a) < \sigma^{-1}(b)$ in $[0,1]$). Given an oriented arc $\vec{A}$ and $a \in A$, we will refer to the points $b \in A$ with $b < a$ as the points *left* of a, and analogously we will speak of points to the *right* of a. We write $a\vec{A}$ for the (oriented) subarc of A consisting of all the points $a' \geq a$, and define $\vec{A}a$ and $a\vec{A}b$ analogously. A sequence $(e_i)_{i=1}^\infty$ of distinct edges or vertices on A is *monotone* if there is an orientation of A such that each e_i lies *between* e_{i-1} and e_{i+1}, i.e. on the right of e_{i-1} and on the left of e_{i+1}. A sequence $(e_i)_{i=1}^\infty$ of distinct edges or vertices on a circle C is *monotone* if there is a subarc A of C containing each e_i and $(e_i)_{i=1}^\infty$ is monotone on A. An *orientation* of C is a choice of one of the two orientations of every arc $A \subseteq C$ such that all these orientations are compatible on their intersections. Given an oriented circle $\vec{C}$ and $a, b \in C$ with $a \neq b$ we define $a\vec{C}b$ to be the (oriented) subarc of C between a and b.

Lemma 1.24 *Let A be an arc in $\overline{G}$. Let $(e_i)_{i=1}^\infty$ and $(f_i)_{i=1}^\infty$ be monotone sequences of distinct edges on A converging from different sides to an end ω of G lying on A. Then ω contains two disjoint rays R and R' such that R contains every e_i while R' contains every f_i.*

Proof. First fix an orientation of A. We may assume that $(e_i)_{i=1}^\infty$ converges to ω from the left, and $(f_i)_{i=1}^\infty$ converges to ω from the right. Let $e_i =: x_i^1 x_i^2$ and $f_i =: y_i^1 y_i^2$ where x_i^1 lies on the left of x_i^2 and y_i^1 lies on the right of y_i^2. Let $A_i := x_i^2 \vec{A} x_{i+1}^1$ and $A_i' := \vec{A} x_1^1 \cup x_{i+1}^2 \vec{A}$, and let $B_i := y_{i+1}^1 \vec{A} y_i^2$ and $B_i' := y_1^1 \vec{A} \cup \vec{A} y_{i+1}^2$.

We will construct the rays R and R' inductively, extending in each step the initial segments of R and R' already defined. Thus suppose that for some $i \geq 0$ we have constructed finite disjoint paths R_i and R_i' which are empty if $i = 0$, and for $i > 0$ are such that R_i joins x_1^2 to x_{i+1}^1, contains each e_j with $1 < j \leq i$ and avoids A_i', while R_i' joins y_1^2 to y_{i+1}^1, contains each f_j with $1 < j \leq i$ and avoids B_i'.

Let us now extend R_i and R_i'. By Lemma 1.23 there is an x_{i+1}^2–x_{i+2}^1 path P in G which avoids the closed set $R_i \cup R_i' \cup A_{i+1}'$. Put $R_{i+1} := R_i e_{i+1} P$. Applying Lemma 1.23 again, we find a y_2^{i+1}–y_1^{i+2} path P' which avoids $R_{i+1} \cup R_i' \cup B_{i+1}'$. Put $R_{i+1}' := R_i' f_{i+1} P'$. Continuing inductively, we obtain rays $R := \bigcup_{i=1}^\infty R_i$ and $R' := \bigcup_{i=1}^\infty R_i'$. But then $e_1 R$ and $f_1 R'$ are as required. $\square$

Lemma 1.25 *Let T be a spanning tree of G, and let T_1, T_2 be subtrees of T with finite intersection. Suppose that G has an end ω which, for each $i = 1, 2$, contains disjoint rays R_i and R'_i such that R_i has infinitely many vertices in T_i. Then $T \notin \mathcal{T}(G)$.*

Proof. For $i = 1, 2$, apply Lemma 1.5 to T_i with $U := V(R_i \cap T_i)$. If the lemma returns a star in one of the T_i then $T \notin \mathcal{T}(G)$ by definition of $\mathcal{T}(G)$. But if it returns a ray in each T_i then both these rays lie in ω, and so T is not end-respecting. Thus again $T \notin \mathcal{T}(G)$. $\square$

Proof of Theorem 1.20. To prove the forward implication, suppose that $T \notin \mathcal{T}(G)$. By Lemma 1.6 and the remark preceding it we may assume that T is end-respecting. Thus there are disjoint equivalent rays R and R' in G such that T contains a subdivision S of an infinite star whose leaves lie on R. Clearly, we may assume that R meets S only in its leaves. Let ω be the end of G containing R and R'. Let $P = x \dots x'$ be an R–R' path in G. Let C' be the circle in $\overline{G}$ consisting of ω together with P, xR and $x'R'$. Let C be the cycle of C'. Thus $C = P \cup xR \cup x'R'$. Let $\mathcal{D}$ be the (infinite) set of all finite cycles which consist of a finite subpath of xR between two consecutive leaves of S on xR together with the path in S joining these leaves. Then $\mathcal{D}$ is not thin, since the centre of S lies in all cycles in $\mathcal{D}$. Lemma 1.22 now implies that C cannot be a sum of fundamental cycles, as required.

To prove the converse implication, we now assume that $T \in \mathcal{T}(G)$. Let C be a cycle of G; we shall prove that C is the sum of all the fundamental cycles C_e of T with $e \in C$. Let $\mathcal{C}$ denote the set of these C_e. Let C' be the defining circle of C, let C'' be the unit circle, and let $\sigma : C'' \to C'$ be a homeomorphism.

We first show that $\mathcal{C}$ is a thin family. Suppose not, and let x be a vertex that lies on C_e for infinitely many chords e of T on C. Since C' is compact, these edges e have an accumulation point ω on C' (which must be an end), and we may choose a monotone sequence $e_1, e_2, \dots$ from among these edges that converges to ω. Since $x \in C_{e_i}$, the endvertices of e_i never lie in the same component of $T - x$. Partitioning the components of $T - x$ suitably into two sets, we may write T as the union of two subtrees T_1 and T_2 that meet precisely in x and are joined by infinitely many e_i. Applying Lemma 1.24 to a suitable subarc of C' containing all the e_i as well as a monotone sequence of edges on C' converging to ω from the other side, we obtain disjoint rays R and R' both belonging to ω and such that R contains every e_i. Then R meets both T_1 and T_2 infinitely often, and we may apply Lemma 1.25 with $R_1 := R =: R_2$ and $R'_1 := R' =: R'_2$ to conclude that $T \notin \mathcal{T}(G)$, contrary to our assumption.

It remains to prove that the cycles in $\mathcal{C}$ sum to C. We thus have to show that an edge f of G lies on an odd number of the cycles in $\mathcal{C}$ if and only if $f \in C$. This is clear when f is a chord of T (and C_f is a fundamental cycle), so we assume that $f \in T$. Let G_1 and G_2 be the subgraphs of G induced by the components of $T - f$, and let E_f be the set of all G_1–G_2 edges of G (including f). Note that the edges $e \neq f$ in E_f are precisely the chords e of T with $f \in C_e$. As $\mathcal{C}$ is thin, C contains only finitely many edges from E_f.

Let us show that the number of edges of C in E_f is even. Since σ is a homeomorphism, $C'' \setminus \sigma^{-1}(\mathring{E}_f \cap C)$ consists of finitely many closed intervals,

$I_1, \ldots, I_k$ say. Since each $\sigma(I_i) \subseteq C'$ is path-connected, it suffices to show that G_1 and G_2 belong to different path-components X_1 and X_2 of $\overline{G} \setminus \mathring{E}_f$: then each $\sigma(I_i)$ lies inside either X_1 or X_2, and thus $E(C) \cap E_f$ is even. Suppose then that G_1 and G_2 are contained in the same path-component of $\overline{G} \setminus \mathring{E}_f$. By Lemma 1.15, there is an arc A in $\overline{G} \setminus \mathring{E}_f$ from a vertex of G_1 to one in G_2. Let ω be the supremum of the points on A that lie in G_1; this can only be an end. Choose monotone edge sequences $(e_i)_{i=1}^{\infty}$ and $(f_i)_{i=1}^{\infty}$ on A with all e_i in G_1 and all f_i in G_2, and so that $(e_i)_{i=1}^{\infty}$ and $(f_i)_{i=1}^{\infty}$ converge to ω from different sides. Apply Lemma 1.24 to obtain disjoint rays R and R' in ω such that R contains every e_i while R' contains every f_i. Now Lemma 1.25 applied with $R_1 := R =: R_2'$ and $R_2 := R' =: R_1'$ implies that $T \notin \mathcal{T}(G)$, a contradiction.

So we have proved that C contains an even number of edges from E_f. As $f \in E_f$, this means that $f \in C$ if and only if C contains an odd number of the edges $e \neq f$ from E_f, which it does if and only if f lies on an odd number of fundamental cycles $C_e \in \mathcal{C}$. □

1.9 Generating arbitrary elements of the cycle space

In this section we characterize the spanning trees whose fundamental cycles generate not only each individual cycle but the entire cycle space of an arbitrary graph. It turns out that these include all normal spanning trees. We shall need this fact in the proof of our characterization theorem below, so let us prove it first:

Lemma 1.26 *Let G be a graph with a normal spanning tree T and let $\overline{G}$ be endowed with* Top. *Then every element Z of the cycle space of G is the sum of fundamental cycles.*

Proof. Write Z as the sum $\sum_{i \in I} Z_i$ of cycles of G. Since $\mathcal{T}(G)$ contains the normal spanning tree T, Theorem 1.20 implies that each Z_i is a sum $\sum_{j \in J_i} C_i^j$ of fundamental cycles. We may assume that the C_i^j are distinct for different $j \in J_i$. To prove the lemma, it suffices to show that the family $\mathcal{C} := (C_i^j)_{i \in I, j \in J_i}$ is thin, since then clearly Z is the sum of all the cycles in $\mathcal{C}$. So suppose that $\mathcal{C}$ is not thin. Then there is a vertex v which lies in the fundamental cycles C_i^j for an infinite set J of pairs (i, j). Since T is normal, every vertex set of a fundamental cycle C_e is a chain in T, its minimum and maximum being joined by e. Thus choosing v minimal in T and possibly discarding finitely many pairs from J, we may assume that v is the lowest vertex (in T) of each C_i^j with $(i, j) \in J$ and thus incident with its chord e_i^j. As C_i^j is the only cycle in the family $(C_i^j)_{j \in J_i}$ that contains e_i^j and this family sums to Z_i, we have $v \in e_i^j \in Z_i$ for all $(i, j) \in J$. But each Z_i has only finitely many summands C_i^j containing v, so $v \in Z_i$ for infinitely many i. Thus $(Z_i)_{i \in I}$ is not thin, contradicting the fact that $Z = \sum_{i \in I} Z_i$. □

We remark that Lemma 1.26 does not extend to arbitrary spanning trees in $\mathcal{T}(G)$. For example, consider the graph G obtained from infinitely many disjoint finite cycles $C_1, C_2, \dots$ by adding a new vertex s and joining it to two vertices of each C_i. Let T be a spanning tree of G containing all the edges of G incident with s. Then $T \in \mathcal{T}(G)$. But as each fundamental cycle contains s, the element $Z = \sum_{i=1}^{\infty} C_i$ of the cycle space of G is not a sum of fundamental cycles.

Let us then determine the subclass $\mathcal{T}'(G) \subseteq \mathcal{T}(G)$ of those spanning trees of G whose fundamental cycles generate all of $\mathcal{C}(G)$. Recall that a *comb* C with *back* R is obtained from a ray R and a sequence $x_1, x_2, \dots$ of distinct vertices (to be called the *teeth* of C) by adding for each $i = 1, 2, \dots$ a (possibly trivial) x_i–R path P_i so that all the P_i are disjoint.

Definition Let $\mathcal{T}'(G)$ be the class of all spanning trees $T \in \mathcal{T}(G)$ such that G does not contain infinitely many disjoint finite cycles $C_1, C_2, \dots$ for which one of the following conditions holds (Fig. 1.4):

- T contains two subdivided infinite stars S_1 and S_2 such that S_1 and S_2 meet at most in the centre of S_1 which is then also the centre of S_2, and such that each C_i contains a leaf of both S_1 and S_2 ($i = 1, 2, \dots$).

- T contains a subdivided infinite star S and a comb C such that S and C are disjoint and each C_i contains both a leaf of S and a tooth of C ($i = 1, 2, \dots$).

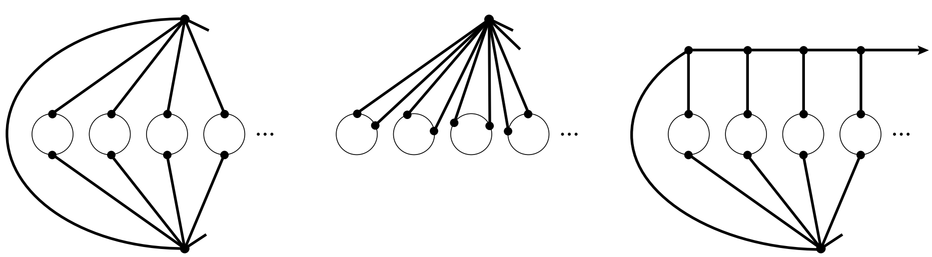

Figure 1.4: The additional forbidden configurations for $\mathcal{T}'(G)$

As before, one can easily show using Lemma 1.3 that $\mathcal{T}'(G)$ contains every normal spanning tree of G.

Theorem 1.27 *Let G be an infinite connected graph and let $\overline{G}$ be endowed with* TOP. *Let T be a spanning tree of G. Then every element of the cycle space of G is a sum of fundamental cycles if and only if $T \in \mathcal{T}'(G)$.*

For the proof of this theorem we need three lemmas.

Lemma 1.28 *For every cycle C in a graph G there exists a countable subgraph H of G such that C is a cycle in H.*

Proof. We may assume that C is not a finite cycle. Let C' be the defining circle of C, and fix an orientation of C'. Note that C' contains only countably many edges (as the inverse images on the unit circle of the interiors of the edges on C' are disjoint open intervals each containing a rational), and thus, by Lemmas 1.4 and 1.10, C' contains only countably many vertices. Let $(x_1, y_1), (x_2, y_2), \ldots$ be an enumeration of the ordered pairs of vertices on C'. Applying Lemma 1.23 to the arcs $x_i \vec{C}' y_i$ we may inductively define paths $P_1, P_2, \ldots$ of G so that each P_i joins x_i to y_i, meets C' only in $x_i \vec{C}' y_i$ and avoids all the previously chosen paths P_j with $x_i \vec{C}' y_i \cap P_j = \emptyset$ $(j < i)$. Thus in particular, P_i avoids all P_j with $j < i$ and $x_i \vec{C}' y_i \cap x_j \vec{C}' y_j = \emptyset$. Let H be the subgraph of G consisting of C and all the paths P_i $(i = 1, 2, \ldots)$. Then H is countable. We will show that C is a cycle in H. Thus we have to find a circle C^* in $\overline{H}$ such that $C^* \cap H = C$.

Let us first show that every end ω of G on C' contains a ray $R_\omega \subseteq H$. We will construct R_ω inductively. Fix a monotone sequence $e_1, e_2, \ldots$ of independent edges on C' converging to ω from the left. Let $e_i =: v_i w_i$ be such that $w_i \in v_i \vec{C}' \omega$. Let Q_1 be the path of the form P_i which was chosen for the pair (w_1, v_2). Then there exists $j_1 > 2$ such that the pair (w_2, v_{j_1}) succeeds (w_1, v_2) in the enumeration of all pairs of vertices on C'. Let Q_2 be the path of the form P_i chosen for (w_2, v_{j_1}). From the choice of j_1 and the fact that $w_1 \vec{C}' v_2 \cap w_2 \vec{C}' v_{j_1} = \emptyset$ it follows that Q_1 and Q_2 are disjoint. Now let $j_2 > j_1$ be such that (w_{j_1}, v_{j_2}) succeeds (w_2, v_{j_1}), and define Q_3 to be the path of the form P_i chosen for (w_{j_1}, v_{j_2}). Again, Q_3 is disjoint from both Q_1 and Q_2. Continue inductively to obtain disjoint paths $Q_1, Q_2, \ldots$. Then the ray $R_\omega := e_1 Q_1 e_2 Q_2 e_{j_1} Q_3 e_{j_2} \ldots$ lies in H. As $e_1, e_2, \ldots$ converges to ω, this ray R_ω also belongs to ω, as required. Let ω' be the end of H containing R_ω. Let us now prove the following claim, which in particular implies that ω' does not depend on the choice of R_ω.

If both $(g_i)_{i=1}^\infty$ and $(h_i)_{i=1}^\infty$ are monotone sequences of independent edges on C' converging to an end $\omega \in C'$, then H contains infinitely many disjoint paths joining endvertices of the g_i to endvertices of the h_i. $(*)$

First note that $(*)$ follows similarly as in the construction of R_ω if $(g_i)_{i=1}^\infty$ and $(h_i)_{i=1}^\infty$ converge to ω from the same side. So we may assume that $(g_i)_{i=1}^\infty$ converges to ω from the left while $(h_i)_{i=1}^\infty$ converges to ω from the right. Let $g_i =: a_i b_i$ and $h_i =: a_i' b_i'$ be such that $b_i \in a_i \vec{C}' \omega$ and $b_i' \in \omega \vec{C}' a_i'$. Let P_1' be the path of the form P_i chosen for (a_1, a_1'). Since P_1' is a closed subset of $\overline{G}$ and does not contain ω, we can find $i_1 > 1$ such that P_1' avoids $a_{i_1} \vec{C}' a_{i_1}'$ and (a_{i_1}, a_{i_1}') succeeds (a_1, a_1') in the enumeration of all pairs of vertices on C'. Let P_2' be the path of the form P_i chosen for (a_{i_1}, a_{i_1}'). Then P_1' and P_2' are disjoint. Now let $i_2 > i_1$ be such that P_2' avoids $a_{i_2} \vec{C}' a_{i_2}'$ and (a_{i_2}, a_{i_2}') succeeds (a_{i_1}, a_{i_1}'), and define P_3' to be the path of the form P_i chosen for (a_{i_2}, a_{i_2}'). Then P_3' is disjoint from both P_1' and P_2'. Continue inductively. Then $P_1', P_2', \ldots$ are disjoint paths as desired in $(*)$.

Let $\varphi : C' \to \overline{H}$ be the map which sends every end ω on C' to ω' and which is the identity on C. Since by definition, every end ω on C' contains the ray R_ω from ω', we have $\omega \supseteq \omega'$; and thus φ is injective. We will show that φ is a topological embedding, and thus that $\varphi(C')$ is a circle in $\overline{H}$ and C is a cycle in H. By Lemma 1.1 it suffices to prove that φ is continuous. This trivially holds in points of $C' \cap H$. So let us consider a basic open neighbourhood N of an end $\omega' \in \varphi(C')$ of $\overline{H}$. Then N is of the form $\widehat{D}$ for some component D of $H - S$ with $S \subseteq V(H)$ finite. We shall find vertices v and w on C' such that $\omega \in v\vec{C}'w$ and $\varphi(v\vec{C}'w) \subseteq \widehat{D}$. This will follow if there are vertices v and w on C' such that $\omega \in v\vec{C}'w$ and $(v\vec{C}'w \cap C) \subseteq D$. Indeed, every end $\tau \in \varphi(v\vec{C}'w)$ contains a ray of H meeting C' in a sequence of edges converging to $\varphi^{-1}(\tau)$. Hence all but finitely many of these edges lie in $(v\vec{C}'w \cap C) \subseteq D$, and thus $\tau \in \widehat{D}$.

So suppose that there are no such vertices v and w. Then there is a monotone sequence $f_1, f_2, \dots$ of independent edges on C' converging to ω such that $f_i \notin D$ for all i. Exactly as in the construction of R_ω we can find a ray $R \subseteq H$ in ω containing infinitely many of the f_i; say $f_{i_1}, f_{i_2}, \dots$. Let $(e_{j_k})_{k=1}^{\infty}$ be the edge-sequence from the construction of R_ω. From $(*)$ it follows that there are infinitely many disjoint paths joining endvertices of the e_{j_k} to endvertices of the f_{i_k}. Thus R and R_ω are equivalent in H, contradicting the fact that $\omega' \in \widehat{D}$.

□

Lemma 1.29 *Let $H_1 \subseteq H_2$ be subgraphs of G, and let C be a cycle in H_1. If C is a cycle in G, then it is also a cycle in H_2.*

Proof. Let C' be the defining circle of C in $\overline{H}_1$. We show that the restriction to C' of the canonical embedding $\pi_{H_1 H_2}$ is injective; then by Lemma 1.1 it is a topological embedding (since $\pi_{H_1 H_2}$ is continuous), and so $C = \pi_{H_1 H_2}(C') \cap H_2$ will be a cycle in H_2.

Note first that $\pi_{H_1 G}$ maps C' onto the defining circle C'' of C in $\overline{G}$: since $\pi_{H_1 G}(C')$ is compact (and hence closed) and contains C as a dense subset, it is the closure of C in $\overline{G}$, which we know to be C''.

Now if $\pi_{H_1 H_2}$ is not injective on C' then neither is $\pi_{H_1 G} = \pi_{H_2 G} \circ \pi_{H_1 H_2}$, so there are two ends $\omega_1, \omega_2 \in C'$ with $\pi_{H_1 G}(\omega_1) = \pi_{H_1 G}(\omega_2)$. Pick $x, y \in C$ so that ω_1, ω_2 lie in distinct path-components of $C' \setminus \{x, y\}$. Then $\pi_{H_1 G}(C' \setminus \{x, y\}) = C'' \setminus \{x, y\}$ is path-connected, contradicting the fact that removing any two distinct points from a circle makes it path-disconnected. □

Lemma 1.30 *Let T be a spanning tree of G, and let $C_1, C_2, \dots \subseteq G$ be disjoint finite cycles. From each C_i pick an edge e_i not on T. If G has a vertex x that lies on each of the fundamental cycles C_{e_i}, then $T \notin \mathcal{T}'(G)$.*

Proof. As $x \in C_{e_i}$, each e_i has its endvertices in two different components of $T - x$. Partitioning these components suitably into two sets, we may write T as the union of two subtrees T_1 and T_2 that meet precisely in x and are joined by infinitely many e_i. Applying Lemma 1.5 to T_1 with U the set of endvertices of these e_i in T_1, we obtain an infinite set $I \subseteq \mathbb{N}$ and either a ray in T_1 joined to all the e_i with $i \in I$ by disjoint paths in T_1, or else a subdivided star in T_1 whose leaves are precisely the endvertices of the e_i with $i \in I$ in T_1. Now apply Lemma 1.5 to T_2 with U the set of endvertices of these e_i $(i \in I)$ in T_2 to obtain an infinite set $I' \subseteq I$ and either a ray or a subdivided star in T_2. If both applications of the lemma return a ray then these rays are equivalent, and so T

is not end-respecting. If both return a star, then these stars can be chosen so as to meet at most in their common centre (which then must be x). As $e_i \in C_i$, each C_i with $i \in I'$ contains leaves of both stars. So these stars satisfy the first condition in the definition of $\mathcal{T}'(G)$. Similarly, if the lemma returns a ray and a star, then they satisfy the second condition in the definition of $\mathcal{T}'(G)$. Thus in each case we have shown that $T \notin \mathcal{T}'(G)$, as desired. □

Proof of Theorem 1.27. To prove the forward implication, suppose that $T \notin \mathcal{T}'(G)$. By Theorem 1.20 we may assume that $T \in \mathcal{T}(G)$. Thus there are disjoint finite cycles $C_1, C_2, \dots$ in G satisfying one of the two conditions in the definition of $\mathcal{T}'(G)$. We consider only the case that T contains two subdivided infinite stars S_1 and S_2 (which are either disjoint or meet only in their common centre) such that each C_i meets both S_1 and S_2; the other case is similar. We may assume that $C_1 \cup C_2 \cup \dots$ avoids the path $P \subseteq T$ joining the centre of S_1 to that of S_2. On each C_i choose an S_1–S_2 path $P_i = x_i \dots y_i$. Since C_i is disjoint from P, the x_i–y_i path in T forms a finite cycle together with P_i. Let $\mathcal{D}$ denote the set of all these cycles, one for each i. Then $\mathcal{D}$ is not thin, as every cycle in $\mathcal{D}$ contains the centre of S_1. Thus Lemma 1.22 implies that the element $Z = \sum_{i=1}^{\infty} C_i$ of the cycle space of G cannot be the sum of fundamental cycles, as desired.

To prove the converse implication, suppose that $T \in \mathcal{T}'(G)$, and let Z be an element of the cycle space of G. Write Z as the sum $\sum_{i \in I} Z_i$ of cycles Z_i. By Theorem 1.20, each Z_i is the sum of a thin family $\mathcal{C}_i = (C_i^j)_{j \in J_i}$ of (distinct) fundamental cycles. It suffices to show that $\mathcal{C} := (C_i^j)_{i \in I, j \in J_i}$ is a thin family: then clearly Z is the sum of all the cycles in $\mathcal{C}$.

Suppose that $\mathcal{C}$ is not thin. Then some vertex x lies on infinitely many cycles in $\mathcal{C}$. Since every family $\mathcal{C}_i$ is thin, there exists an infinite set $I' \subseteq I$ such that for every $i \in I'$ the vertex x lies on some cycle in $\mathcal{C}_i$. Denoting the defining chord of this (fundamental) cycle by e_i, we thus have $x \in C_{e_i} \in \mathcal{C}_i$ for every $i \in I'$.

As the fundamental cycles in $\mathcal{C}_i$ are distinct, their defining chords do not cancel in the sum $\sum_{C \in \mathcal{C}_i} C = Z_i$, so $e_i \in Z_i$ for every i. On the other hand as the family $(Z_i)_{i \in I}$ is thin, we have $e_i \in Z_k$ for only finitely many k. In particular, $e_i \neq e_k$ for all but finitely many k. Conversely, Z_k contains only finitely many e_i (since $\mathcal{C}_k$ is thin and every C_{e_i} contains x), so $Z_k \ni e_i$ for only finitely many i. Replacing I' with an appropriate infinite subset if necessary, we may therefore assume that $e_i \in Z_k$ if and only if $i = k$ (for all $i, k \in I'$), and further that I' is countable.

For $Z' := \sum_{i \in I'} Z_i$ the above implies that $e_i \in Z'$ for all $i \in I'$. Moreover, Lemmas 1.28 and 1.29 imply that Z' lies in the cycle space of a countable subgraph H of G. Since every countable connected graph has a normal spanning tree, Lemma 1.26 thus implies that Z' is a sum of a thin family $\mathcal{C}'$ of finite cycles: of fundamental cycles of normal spanning trees of the components of H. As every e_i lies in Z' and hence in some cycle of $\mathcal{C}'$, and since each of these cycles meets only finitely many others, $\mathcal{C}'$ has an infinite subfamily of disjoint cycles each containing an edge e_i with $i \in I'$. Lemma 1.30 now implies that $T \notin \mathcal{T}'(G)$, contradicting our assumption. $\square$

1.10 The structure of the elements of the cycle space

Our main purpose in this section is to prove the infinite analogue of the fact that every element of the cycle space of a finite graph is an edge-disjoint union of finite cycles.

First, however, let us extend Theorem 1.18 to arbitrary infinite graphs that have a normal spanning tree. Let us say that a set X of vertices *covers* a set of edges if each of these edges has a vertex in X.

Theorem 1.31 *Let G be any graph that has a normal spanning tree, and let $H \subseteq G$ be any subgraph without isolated vertices. Then the following statements are equivalent:*

(i) *$H \in \mathcal{C}(G)$;*

(ii) *for every cut F of G that is covered by finitely many vertices, $|E(H) \cap F|$ is (finite and) even.*

Proof. The proof is essentially the same as that of Theorem 1.18. For the implication (i)$\Rightarrow$(ii) we now first have to prove that $|E(C) \cap F|$ is finite for every cycle C, but this is clear since F is covered by finitely many vertices and C is 2-regular by Lemmas 1.4 and 1.10. Similarly as in the proof of Theorem 1.18 it can then be shown that $|E(C) \cap F|$ is even. Since every $Z \in \mathcal{C}(G)$ is the sum of a thin family of cycles, and so only finitely many of these cycles meet F, it follows that $|E(Z) \cap F|$ is even.

For the converse implication we now use a normal spanning tree T, which has the property that the cut of G corresponding to any edge $f \in T$ is covered by the finitely many vertices that lie below the vertices of f in T. $\square$

Theorem 1.32 *Let G be an infinite graph and let $\overline{G}$ be endowed with* Top. *Then every element of the cycle space of G is an edge-disjoint union of cycles in G.*

The basic idea for the proof of Theorem 1.32 is the same as in the finite case: given $Z \in \mathcal{C}(G)$, we shall find a single cycle $C \subseteq Z$ in G and then iterate with $Z - C$, continuing until the cycles deducted from Z have exhausted it. As in the finite case, none of the cycles from the constituent sum of Z need be a subgraph of Z, so finding C is non-trivial. But while for finite Z we can find C greedily inside Z (using the fact that all degrees in Z are at least 2), this need not be possible when Z is infinite: a maximal path in Z may well be a double ray rather than define a cycle, and it is then not clear how to extend this double ray beyond its ends to a circle giving rise to the desired cycle C.

Our main lemma for the proof of Theorem 1.32 thus deals with finding C, and it does so in a countable subgraph H of G. Finding the right H in which to do this will cause a few (managable) complications later on, but the key advantage is that H, being countable, will have a normal spanning tree T. We may then write any $Z \in \mathcal{C}(H)$ as a sum of *finite* cycles (namely, of fundamental cycles with respect to T; cf. Lemma 1.26), which will make standard compactness arguments available for the construction of C.

Lemma 1.33 *Let H be a countable subgraph of G, let $Z \in \mathcal{C}(H)$, and let $e = vw \in E(Z)$. Then $\overline{H}$ contains a topological path P from v to w such that $P \cap H \subseteq Z - e$.*

Proof. As H is countable, it has a normal spanning tree. Thus Lemma 1.26 implies that Z can be written as $Z = \sum_{i=1}^{\infty} C_i$, where the C_i are finite cycles in H forming a thin family. Let $H' := \bigcup_{i=1}^{\infty} C_i$. Replacing Z with the sum Z' of those C_i that lie in the component of H' containing e, we may assume that H' is connected. (Indeed, $Z' \in \mathcal{C}(H)$ and $e \in Z' \subseteq Z$; hence a proof of the lemma for Z' implies it for Z.) Since the family $(C_i)_{i=1}^{\infty}$ is thin, H' is locally finite. Put $Z_i := \sum_{j=1}^{i} C_j$. As $e \in Z$, there exists an $i_0 > 0$ such that $e \in Z_i$ for all $i \geq i_0$. Furthermore, each Z_i is finite and hence an edge-disjoint union of finite cycles in H'. Fix such a set of finite cycles for every $i \geq i_0$, and let D_i denote the cycle containing e. Let P_i be the finite path $D_i - e$, and orient it from v to w.

Let $e_1, e_2, \dots$ be an enumeration of the edges in $E(H') \setminus \{e\}$. Let us define a sequence $X_0 \subseteq X_1 \subseteq \dots$ of finite subsets of $E(H') \setminus \{e\}$ and a sequence $I_0 \supseteq I_1 \supseteq \dots$ of infinite subsets of $\mathbb{N}$ so that the following holds for all $i = 0, 1 \dots$:

$X_i = \{e_1, \dots, e_i\} \cap E(P_j)$ for all $j \in I_i$, and all these P_j induce the same linear ordering on X_i and the same orientation on the edges in X_i. $(*)$

To do this, we begin with $X_0 = \emptyset$ and $I_0 = \{i \in \mathbb{N} \mid i \geq i_0\}$. For every $i \geq 0$ in turn, we then check whether $e_{i+1} \in P_j$ for infinitely many $j \in I_i$. If so, we put $X_{i+1} := X_i \cup \{e_{i+1}\}$ and choose $I_{i+1} \subseteq I_i$ so as to satisfy $(*)$ for $i+1$; if not, we let $X_{i+1} := X_i$ and put $I_{i+1} := \{j \in I_i \mid e_{i+1} \notin P_j\}$ (in which case $I_i \setminus I_{i+1}$ is finite, and $(*)$ again holds for $i + 1$). Finally, let $X := \bigcup_{i=0}^{\infty} X_i$, and write $\dot{X}$ for the subgraph of H consisting of the edges in X and their incident vertices.

The set X is linearly ordered as follows. Given $e, f \in X$, consider the least index i such that $e, f \in X_i$. If e precedes f (say) on one P_j with $j \in I_i$ then it does so on every such P_j, and hence in particular on every P_j with $j \in I_k$ and $k > i$ (since $I_k \subseteq I_i$). Similarly, every edge $e \in X$ has a unique orientation, its common orientation on every P_j with $j \in I_i$ and i large enough that $e \in X_i$.

Let us show that $\dot{X} \subseteq Z - e$. Given an edge $f \in X$, we have $f \in P_j \subseteq Z_j - e$ for infinitely many j—indeed, by $(*)$ this holds for all $j \in I_i$ with i large enough that $f \in X_i$. But then $f \in Z_j$ for all large enough j (because f lies on only finitely many C_i), and hence also $f \in Z$.

Using the local finiteness of H', it is in fact easy to show that $\dot{X} + e$ is a 2-regular subgraph of Z, in which two edges of X are adjacent if and only if they are adjacent elements in the linear ordering on X. Indeed, given a vertex $u \in \dot{X}$ choose k large enough that every edge of H' incident with u precedes e_k in the enumeration of all the edges e_i, and pick $j \in I_k$. Then the edges at u in $\dot{X}$ are precisely the edges at u in X_k, which by $(*)$ are precisely the edges at u in P_j. If $u \in \{v, w\}$ there is one such edge; otherwise there are two.

If X is finite, then $\dot{X}$ is a v–w path in $Z - e$, and thus $\dot{X}$ is a topological path P as sought in the lemma. So let us assume that X is infinite. Then $\dot{X}$ is a disjoint union of two rays R_v and R_w starting at v and w, respectively, and possibly some further double rays. We will show that the closure of $\dot{X}$ in $\overline{H}$ is a topological path P as desired.

Assign to R_v a half-open subinterval J_{R_v} of $[0,1]$ containing 0, to R_w a half-open subinterval J_{R_w} containing 1, and to each double ray $D \subseteq \dot{X}$ an open subinterval J_D, in such a way that all these intervals are disjoint, their order on $[0,1]$ (oriented from 0 to 1) reflects the order of their corresponding rays and double rays induced by the linear ordering on X, and so that $[0,1]$ is the closure of the union U of these subintervals. (Since $\dot{X}$ contains only countably many double rays, this can be done in at most ω steps.) Let $\sigma : [0,1] \to \overline{H}$ map each interval J_Q continuously and bijectively onto its ray or double ray Q so that the order of the edges of Q in X reflects that induced by σ. Thus in particular $\sigma(0) = v$ and $\sigma(1) = w$. In what follows we will show that we can extend σ to a continuous map from $[0,1]$ to $\overline{H}$ by mapping the points in $[0,1] \setminus U$ to suitable ends of H. The image of $[0,1]$ will then be a path P as desired.

So let x be a point in $[0,1] \setminus U$. Choose a sequence $(u_i)_{i=1}^{\infty}$ of vertices of $\dot{X}$ so that the sequence $(\sigma^{-1}(u_i))_{i=1}^{\infty}$ is monotone in $[0,1]$ and converges to x. Since H' is connected and locally finite, we may apply Lemma 1.5 to find a ray $Q_x \subseteq H'$ such that H' contains infinitely many disjoint Q_x–$\{u_i \mid i \in \mathbb{N}\}$ paths. Let ω_x be the end of H containing Q_x, and extend σ by setting $\sigma(x) := \omega_x$. (We remark that although formally ω_x depends on the choice of $(u_i)_{i=1}^{\infty}$, this is not in fact the case. However, we shall not need this below.)

We have to prove that $\sigma : [0,1] \to \overline{H}$ is continuous. Clearly, σ is continuous in points of U. So let $x \in [0,1] \setminus U$, and let N be a basic open neighbourhood of ω_x in $\overline{H}$. Then N is of the form $\widehat{D}$ for some component D of $H - S$ with $S \subseteq V(H)$ finite. We have to show that there is an open neighbourhood O of x in $[0,1]$ such that $\sigma(O) \subseteq \widehat{D}$.

We will first show that there are points $a \neq b$ in $[0,1]$ such that $x \in (a,b)$ and either $\sigma(a,x) \cap \dot{X} \subseteq D$ or $\sigma(a,x) \cap \dot{X} \cap D = \emptyset$, and such that the analogous assertion holds for (x,b). Let $k := |S|$, and suppose there is no such point a (say). Then we can find a monotone sequence $f_1, \ldots, f_{k+2}$ of $k+2$ distinct edges in X lying alternately inside and outside of D (and having no incident vertex in S). As the sequence $f_1, \ldots, f_{k+2}$ is monotone in the ordering on X (and this ordering is well defined), every path P_j with $j \in I_i$ and i large enough that $f_1, \ldots, f_{k+2} \in X_i$ contains all these edges in this order. But then P_j meets S in at least $k+1$ vertices, a contradiction. Hence there are points a and b as required.

Let us now show that either $\sigma(a,b) \cap \dot{X} \subseteq D$ or $\sigma(a,b) \cap \dot{X} \cap D = \emptyset$. This will follow from the choice of a and b if there are sequences $(v_i)_{i=1}^{\infty}$ and $(v_i')_{i=1}^{\infty}$ of distinct vertices of $\dot{X}$ such that $(\sigma^{-1}(v_i))_{i=1}^{\infty}$ is monotone and converges to

x from the left while $(\sigma^{-1}(v'_i))_{i=1}^{\infty}$ is monotone and converges to x from the right, and such that H contains infinitely many disjoint paths $P'_i = v_i \dots v'_i$. We will construct such paths inductively (Fig. 1.5). Let $(f_i)_{i=1}^{\infty}$ and $(f'_i)_{i=1}^{\infty}$ be monotone sequences of distinct edges of X such that $(\sigma^{-1}(f_i))_{i=1}^{\infty}$ converges to x from the left while $(\sigma^{-1}(f'_i))_{i=1}^{\infty}$ converges to x from the right, and such that f_{i+1} succeeds both f_i and f'_{i+1} in the enumeration $e_1, e_2, \dots$ of $E(H') \setminus \{e\}$, and f'_{i+1} succeeds f'_i in this enumeration (for all $i \geq 1$). Let k be such that $f_1 = e_k$, and pick $r \in I_k$. Then $f_1, f'_1 \in P_r$: if $i < k$ is such that $e_i = f'_1$, then $r \in I_k \subseteq I_i$ and hence $f'_1 = e_i \in X_i \subseteq E(P_r)$ by $(*)$. Moreover, since f_1 lies to the left of f'_1 in X, it precedes f'_1 on P_r. Let v_1 be the last vertex of f_1 on P_r, and let v'_1 be the first vertex of f'_1 on P_r. Put $P'_1 := v_1 P_r v'_1$. Now let $s > 1$ be such that f_s succeeds every edge of P'_1 in the sequence $e_1, e_2, \dots$, and such that no edge of $E(P'_1) \cap X$ lies between f_s and f'_s in X. Let k' be such that $f_s = e_{k'}$, and pick $r' \in I_{k'}$. Then $f_s, f'_s \in P_{r'}$, and f_s precedes f'_s on $P_{r'}$. Let v_2 be the last vertex of f_s on $P_{r'}$, and let v'_2 be the first vertex of f'_s on $P_{r'}$. Put $P'_2 := v_2 P_{r'} v'_2$. Since $e_{k'}$ succeeds every edge from $E(P'_1) \setminus X$ in the enumeration of the e_i, condition $(*)$ implies that $P_{r'}$ (and hence P'_2) has no edge in $E(P'_1) \setminus X$. And P'_2 has no edge in $E(P'_1) \cap X$, because none of those edges lies between f_s and f'_s in X: since $e_{k'}$ equals or succeeds f_s, f'_s and every edge from $E(P'_1) \cap X$ in the enumeration of the e_i, the position of any such edge on $P_{r'}$ relative to f_s and f'_s would be the same as in X, i.e. it would precede f_s or succeed f'_s on $P_{r'}$ and hence not lie on P'_2. Thus P'_1 and P'_2 are edge-disjoint. Continuing inductively, we obtain infinitely many edge-disjoint paths $P'_i = v_i \dots v'_i$, one for every $i \in \mathbb{N}$. As all these paths lie in the locally finite graph H', infinitely many of them are disjoint, as desired. Thus we have shown that either $\sigma(a,b) \cap \dot{X} \subseteq D$ or $\sigma(a,b) \cap \dot{X} \cap D = \emptyset$.

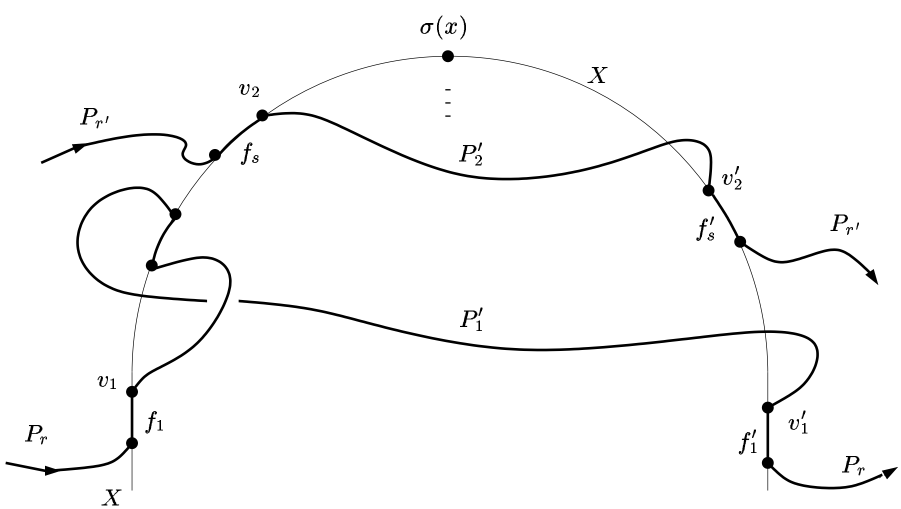

Figure 1.5: Constructing the paths P'_i

By definition, ω_x contains the ray Q_x, and Q_x was defined in such a way that there is a sequence $(u_i)_{i=1}^{\infty}$ of distinct vertices in $\dot{X}$ such that H' contains infinitely many disjoint Q_x–$\{u_i \mid i \in \mathbb{N}\}$ paths, and where $(\sigma^{-1}(u_i))_{i=1}^{\infty}$ converges to x. Then all but finitely many of the points $\sigma^{-1}(u_i)$ lie in (a,b). Since $\sigma(x) = \omega_x \in \widehat{D}$, it follows that $\sigma(a,b) \cap \dot{X} \subseteq D$. Now let $y \in (a,b)$ be such that $\sigma(y)$ is an end of H. Thus $\sigma(y) = \omega_y$, and ω_y contains the ray Q_y. As before, the definition of Q_y and the fact that $\sigma(a,b) \cap \dot{X} \subseteq D$ imply that $\sigma(y) \in \widehat{D}$. Thus $\sigma(a,b) \subseteq \widehat{D}$, and we have shown that σ is continuous. □

Proof of Theorem 1.32. Let $Z \in \mathcal{C}(G)$ be given, and let $Z = \sum_{i \in I} Z_i$ where each Z_i is a cycle in G. We first show that I may be partitioned into countable sets I_α so that for all $\alpha \neq \beta$ the graphs $\sum_{i \in I_\alpha} Z_i$ and $\sum_{i \in I_\beta} Z_i$ are edge-disjoint. To do this, consider the auxiliary graph G' with vertex set I in which $i \neq j$ are joined if Z_i and Z_j share an edge. As each Z_i has only countably many edges and each edge lies in only finitely many Z_i, each i has only countably many neighbours in G'. Thus every component of G' is countable, and so the vertex sets I_α of the components of G' form a partition of I with the desired properties. Hence, to prove the theorem, we may assume that I itself is countable. Lemmas 1.28 and 1.29 now imply that there is a countable subgraph H of G such that every Z_i is a cycle in H, and thus Z is an element of the cycle space of H.

Let us rename H as H^0 and Z as Z^0, so that from now on we may use 'H' and 'Z' as variables in Lemma 1.33. Our aim is to write Z^0 as an edge-disjoint union of cycles $C^1, C^2, \ldots$ in G. We shall find these C^n inductively inside $Z^{n-1} := \sum_{i \in I} Z_i + C^1 + \ldots + C^{n-1}$ by applying Lemma 1.33 to $Z = Z^{n-1}$ in a suitable subgraph H^{n-1} of G. (Thus $C^n \subseteq Z^{n-1}$, and hence $Z^0 \supseteq Z^1 \supseteq Z^2 \supseteq \ldots$ with $Z^n = Z^{n-1} - C^n$.)

Starting our inductive definition of the C^n at $n = 1$, let us assume that $C^1, \ldots, C^{n-1}$ (and hence $Z^0, \ldots, Z^{n-1}$) have been defined as above, and that H^{n-1} is some countable subgraph of G in which $C^1, \ldots, C^{n-1}$ and all the Z_i are cycles. To define C^n, let P be as provided by Lemma 1.33 for $H = H^{n-1}$ and $Z = Z^{n-1}$, where $e = vw$ is taken to be the first edge in Z^{n-1} from some fixed enumeration of all the edges of Z^0. (As e will lie in C^n, this choice of e ensures that all the C^n together exhaust Z^0.) The image $\pi_{HG}(P)$ of P in $\overline{G}$ under the canonical projection $\pi_{HG} : \overline{H} \to \overline{G}$ is a path in $\overline{G}$ from v to w. Apply Lemma 1.15 to find an arc $A \subseteq \pi_{HG}(P)$ in $\overline{G}$ with endpoints v and w. Then $A \cup e$ is a circle in $\overline{G}$ whose cycle (in G) is a subgraph of Z^{n-1} containing e, because $P \cap G = P \cap H \subseteq Z^{n-1} - e$; we take C^n to be this cycle.

By Lemma 1.28 there is a countable subgraph H' of G such that C^n is a cycle in H'. By Lemma 1.29 and our assumptions on H^{n-1}, all of $C^1, \ldots, C^n$ and all the Z_i then are cycles in $H^n := H^{n-1} \cup H'$, as well as in G.

This completes the inductive definition of the cycles C^n. Since each C^n is a subgraph of Z^{n-1} and $Z^n = Z^{n-1} - C^n$, no edge of C^n is left in Z^n, and so the C^n are indeed edge-disjoint. By the choice of the edges $e = vw$, every edge of $Z = Z^0$ lies in some C^n, and the theorem follows. □

As mentioned before, the cycle space of a graph is not obviously closed under taking infinite sums. Indeed, let $(Z_i)_{i \in I}$ be a thin family of elements of $\mathcal{C}(G)$ (so that $Z := \sum_{i \in I} Z_i$ is well defined), and for each i let $Z_i = \sum_{j \in J_i} C_i^j$ where all the C_i^j are cycles. Then the canonical way to establish Z as an element of $\mathcal{C}(G)$ would be to write it as the 'sum' $Z = \sum_{i \in I, j \in J_i} C_i^j$. But this 'sum' may be ill-defined, since the family of all the cycles C_i^j need not be thin even though $(Z_i)_{i \in I}$ is a thin family. For example, if a vertex v lies on exactly two cycles C_i^j for each i, and if both these cycles contain the same two edges at v, then v is not a vertex of Z_i (since we suppress isolated vertices in our definition of sum) and hence does not contradict the thinness of the family $(Z_i)_{i \in I}$; but it does prevent the family of all the C_i^j from being thin.

This phenomenon does not occur, however, when the cycles C_i^j in each of the sums $Z_i = \sum_{j \in J_i} C_i^j$ are edge-disjoint: then $V(Z_i) = \bigcup_{j \in J_i} V(C_i^j)$, and hence if both $(Z_i)_{i \in I}$ and all the $(C_i^j)_{j \in J_i}$ are thin families then so is $(C_i^j)_{i \in I, j \in J_i}$. Theorem 1.32 therefore implies that $\mathcal{C}(G)$ is indeed closed under infinite as well as finite sums:

Corollary 1.34 *In* TOP *the cycle space of an infinite graph is closed under taking sums.*

1.11 An open problem

The subgraphs C of a finite graph G that are cycles or other elements of the cycle space of G are easily characterized without any reference to a notion of cyclicity (such as cyclic sequences of vertices etc.). For example, C is a cycle if and only if it is 2-regular and connected, and C is an element of $\mathcal{C}(G)$ if and only if all its vertices have even degree. Similarly, $C \in \mathcal{C}(G)$ if and only if C is orthogonal to every cut of G, ie. meets every cut in an even number of edges.

Since our definition of an infinite cycle appeals to an external notion of cyclicity in an even stronger sense by making reference to the topology of the unit circle, it seems all the more desirable to have similar characterizations for infinite cycles:

Problem *Characterize the cycles and the elements of the cycle space in an infinite graph in purely combinatorial terms.*

Theorem 1.31 offers such a characterization in terms of cuts. Alternatively, one might try to extend the finite "even degree" characterization of the cycle space to infinite graphs. Clearly, any such characterization will have to refer to ends, but the idea is that such reference should not explicitly appeal to the topology on $\overline{G}$. For example, one might try to define the 'degree' of an end of G in such a way that a subgraph C of G lies in $\mathcal{C}(G)$ if and only if all its vertices have even degree and all its ends have even or infinite degree. One of the problems with such an approach will be in which subgraph to measure the 'degrees' of these ends: probably not in G itself (since an end ω of G that lies on C can contain further rays that have little to do with C), and certainly not in C (where ω will typically split up into many unrelated ends).

Chapter 2

Induced subdivisions in $K_{s,s}$-free graphs of large average degree

2.1 Introduction

A classical theorem of Mader states that for every graph H there exists $d = d(H)$ such that every graph G of average degree at least d contains a subdivision of H. Obviously, the result becomes false if we ask for an *induced* subdivision of H. Here we prove that this stronger assertion holds if G is 'locally sparse' in the sense that it fails to contain some complete bipartite graph $K_{s,s}$:

Theorem 2.1 *For every graph H and every $s \in \mathbb{N}$ there exists $d = d(H, s)$ such that every graph G of average degree at least d contains either a $K_{s,s}$ as a subgraph or an induced subdivision of H.*

Of course, one cannot replace 'subdivision' by 'subgraph', as for example there exist graphs which have both arbitrarily large average degree and arbitrarily large girth. On the other hand, Kierstead and Penrice [21] proved that if H is a tree then one can indeed find it as an induced subgraph in any $K_{s,s}$-free graph of sufficiently large average degree:

Theorem 2.2 *For every tree T and every $s \in \mathbb{N}$ there exists $d = d(T, s)$ such that every graph of average degree at least d contains either $K_{s,s}$ as a subgraph or an induced copy of T.*

They used this result to prove a special case of the conjecture of Gyárfás and Sumner that given a tree T and $s \in \mathbb{N}$, every K_s-free graph of sufficiently large chromatic number contains an induced copy of T. Scott [43] proved that this conjecture becomes true if we only require an induced subdivision of T. Motivated by this result, he proposed a conjecture which is analogous to Theorem 2.1—replacing 'average degree' by 'chromatic number' and $K_{s,s}$ by K_s.

We now briefly outline the organization of this chapter and the strategy of our proof of Theorem 2.1. Consider a $K_{s,s}$-free graph G of large average degree. In Section 2.2 we prepare the ground for the proof by collecting some tools which we will need later on. In particular, it turns out that in order to find an induced subdivision of H in G, it suffices to find an induced subdivision of any graph H' with large enough average degree so that every edge of H' is subdivided exactly once. We will call this a *1-subdivision* of H'. So both the set B of branch vertices and the set S of subdividing vertices have to be independent in G.

The first step towards finding such a 1-subdivision of H' is to find a large independent set I in G (Section 2.3). Ideally, we would like to find another independent set B' such that the bipartite subgraph between I and B' has large average degree. In this case, one can find B in the smaller of B' and I and S in the larger of the two. Unfortunately, we cannot guarantee that such a set B' always exists. However, in Section 2.4 we will show that one can come fairly close: we will find sets $I' \subseteq I$ and B' such that the bipartite subgraph between I' and B' has large average degree and $G[B']$ has small chromatic number. In Section 2.5, which constitutes the core of our proof, we then show how to find our induced 1-subdivision of H' within $G[I' \cup B']$. In Section 2.6 we put everything together to complete the proof of Theorem 2.1.

The special case of Theorem 2.1 when H is a cycle admits a much simpler elementary proof, which we include in Section 2.7. The main tool used in this proof is the case of Theorem 2.2 where T is a path. In Section 2.8 we offer an alternative proof of the entire Theorem 2.2.

2.2 Notation and tools

All graphs considered in this chapter are finite, and all logarithms are base two. We write $d(G) := 2|E(G)|/|V(G)|$ for the *average degree* of a graph G, $\delta(G)$ for its minimum degree and $\chi(G)$ for its chromatic number. Given a vertex x of G, we denote by $d(x)$ or $d_G(x)$ the degree of x and by $N(x)$ or $N_G(x)$ the set of neighbours of x. Given graphs G and H we say that G is *H-free* if G does not contain H as a subgraph. A *subdivision* of a graph H is a graph G obtained from H by replacing the edges of H with internally disjoint paths between their endvertices. We view $V(H)$ as a subset of $V(G)$ and call these vertices the *branch vertices* of G. A *1-subdivision* of a graph H is the graph obtained from H by replacing the edges of H with internally disjoint paths of length two.

For disjoint sets $A, B \subseteq V(G)$ we write $e(A, B)$ for the number of A–B edges in G and $(A, B)_G$ for the bipartite subgraph of G whose vertex classes are A, B and whose edges are the A–B edges in G. If we say that a bipartite graph (A', B') is a subgraph of (A, B) then we tacitly assume that $A' \subseteq A$ and $B' \subseteq B$. We shall frequently consider the following class of graphs.

Definition. Given non-negative numbers d, i and $k \leq d/4$, we say that a bipartite graph (A, B) is a *(d, i, k)-graph* if $|A| \geq d^{12i}|B|$ and $d/4-k \leq d(a) \leq 4d$ for all vertices $a \in A$. (So the order of A and B matters here.)

We now list some results which we need later on in the proof of Theorem 2.1. We shall frequently use the following simple observations. Proofs are for example included in [6, Prop. 1.2.2 resp. Cor. 5.2.3].

Proposition 2.3 *Every graph G contains an induced subgraph of average degree at least $d(G)$ and minimum degree at least $d(G)/2$.*

Proposition 2.4 *Every graph G contains an induced subgraph of minimum degree at least $\chi(G) - 1$.*

Clearly, it suffices to prove Theorem 2.1 for graphs G which do not have subgraphs of average degree $> d(G)$. So the propositions enable us to assume that $\delta(G) \geq d(G)/2$ and $\chi(G) \leq d(G) + 1$.

The following theorem of Mader (for a proof see e.g. [6, Thm. 3.6.1]) implies that it suffices to show that G contains an induced 1-subdivision of any graph H' of large enough average degree. Indeed, from the theorem it follows that H' contains a subdivision of H; and it is easily checked that the corresponding subdivision of H in G is induced.

Theorem 2.5 *For every $r \in \mathbb{N}$ there exists $d = d(r)$ such that every graph of average degree at least d contains a subdivision of K_r.*

Bollobás and Thomason as well as Komlós and Szemerédi independently showed that the order of magnitude of $d(r)$ is r^2 (see e.g. [6, Thm. 8.1.1]). We shall also use the following well known upper bound for the average degree of $K_{s,s}$-free graphs (see e.g. [5, p. 74]).

Theorem 2.6 *If G is a $K_{s,s}$-free graph then $d(G) \leq c_s|G|^{1-1/s}$ where c_s is some constant depending on s.*

The next lemma is a special case of Chernoff's inequality (see for example [2, Thm. A.1.12 and A.1.13]).

Lemma 2.7 *Let $X_1, \dots, X_n$ be independent 0-1 random variables with $\mathbb{P}(X_i = 1) = p$ for all $i \leq n$, and let $X := \sum_{i=1}^n X_i$. Then $\mathbb{P}(X \geq 2\mathbb{E}X) \leq (4/\mathrm{e})^{-\mathbb{E}X}$ and $\mathbb{P}(X \leq \mathbb{E}X/2) \leq \mathrm{e}^{-\mathbb{E}X/8}$.*

One case which arises in our proof of Theorem 2.1 is that we first find an induced bipartite subgraph (A, B) of large average degree in G and then find an induced subdivision of H in (A, B). To carry out this second step, it will turn out to be useful if the vertices in A have almost the same degree and $|B|$ is much smaller than $|A|$. The following lemma shows that by replacing (A, B) with an induced subgraph we can always satisfy these two additional conditions. The lemma is a slight extension of [37, Lemma 2.4]. Although the proof is almost the same, we include it here for completeness.

Lemma 2.8 *Let $r \geq 2^6$, $s \geq 1$ and $d \geq 8r^{12s+1}$. Then every bipartite graph of average degree d contains an induced copy of an $(r, s, 0)$-graph.*

Proof. Clearly, we may assume that our given bipartite graph has no subgraph of average degree $> d$. So by Proposition 2.3 this graph contains an induced subgraph $G = (A, B)$ such that $\delta(G) \geq d/2$, $d(G) = d$ and $|A| \geq |B|$. Thus at most half of the vertices of A have degree at least $2d$ in G. So, writing A' for the set of all vertices in A of degree at most $2d$, we have $|A'| \geq |A|/2 \geq |B|/2$.

Let us now consider a random subset B_p of B which is obtained by including each vertex of B independently with probability $p := r/d$. For every $a \in A'$ let $X_a := |N_G(a) \cap B_p|$. Then $r/2 \leq \mathbb{E}X_a \leq 2r$. Given B_p, let us call $a \in A'$ *useful* if $r/4 \leq X_a \leq 4r$. Lemma 2.7 implies that

$$\mathbb{P}(a \text{ is not useful}) \leq \mathbb{P}(X_a \geq 2\mathbb{E}X_a) + \mathbb{P}(X_a \leq \mathbb{E}X_a/2) \leq (4/\mathrm{e})^{-r/2} + \mathrm{e}^{-r/16} \leq \frac{1}{4}.$$

Hence the expected number of vertices in A' which are not useful is at most $|A'|/4$. So Markov's inequality (which states that $\mathbb{P}(X \geq c\,\mathbb{E}X) \leq 1/c$ for every $c \geq 1$) implies that

$$\mathbb{P}(\text{at least half of the vertices in } A' \text{ are not useful}) \leq \frac{1}{2}.$$

Moreover, using Lemma 2.7 again,

$$\mathbb{P}(|B_p| \geq 2p|B|) = \mathbb{P}(|B_p| \geq 2\mathbb{E}|B_p|) \leq (4/\mathrm{e})^{-p|B|} \leq \frac{1}{4}.$$

So the probability that both $|B_p| \leq 2p|B|$ and that at least half of the vertices in A' are useful is at least $1/2 - 1/4 > 0$. Hence there exists a choice B^* for B_p which has these two properties. Let A^* be the set of useful vertices in A'. Then $r/4 \leq d_{(A^*,B^*)_G}(a) \leq 4r$ for every vertex $a \in A^*$ and

$$|A^*| \geq \frac{|A'|}{2} \geq \frac{|B|}{4} \geq \frac{|B^*|}{8p} = \frac{d|B^*|}{8r} \geq r^{12s}|B^*|.$$

Thus $(A^*, B^*)_G$ is an induced $(r, s, 0)$-subgraph of G. $\square$

2.3 Independent sets

Clearly, every graph G of maximum degree Δ has an independent set of size at least $|G|/\chi(G) \geq |G|/(\Delta + 1)$. Lemma 2.9 shows that we obtain a small but significant improvement if G is $K_{s,s}$-free. The proof is based on Alon's elegant proof of the result that any triangle-free graph H of maximum degree Δ contains an independent set of size $c|H|\log\Delta/\Delta$ (see e.g. [2], the result itself is due to Ajtai, Komlós and Szemerédi [1]).

Alternatively, we could have applied another result in [1]: for all ε there exists a constant c_0 so that every graph with maximum degree at most Δ which contains at most $|G|\Delta^{2-\varepsilon}$ triangles has an independent set of size at least $c_0|G|\log\Delta/\Delta$. But Theorem 2.6 implies that in a $K_{s,s}$-free graph G the neighbourhood of any vertex x can span at most $c_s d(x)^{2-1/s} \leq c_s\Delta^{2-1/s}$ edges and thus G contains at most $c_s|G|\Delta^{2-1/s}$ triangles. Although proof of Lemma 2.9 given below yields a weaker bound, it is simpler and has the advantage of being self-contained.

Lemma 2.9 *Let $\Delta > 2$. For every $s \in \mathbb{N}$ there exists $c' = c'(s)$ such that every $K_{s,s}$-free graph G of maximum degree at most Δ has an independent set of size at least*

$$f := c'|G|\frac{(\log \Delta)^{1/s}}{\Delta \log\log \Delta}.$$

Proof. Let $n := |G|$. Let I be an independent set chosen uniformly at random from all independent sets of G. For every vertex $x \in G$ define

$$Z_x := \begin{cases} \Delta & \text{if } x \in I; \\ |N(x) \cap I| & \text{otherwise.} \end{cases}$$

Then

$$\sum_{x \in G} Z_x = \sum_{x \in I} Z_x + \sum_{x \notin I} Z_x \leq \Delta|I| + e(I, V(G) \setminus I) \leq 2\Delta|I|.$$

So it suffices to show that $\mathbb{E}(\sum_{x\in G} Z_x) \geq 2\Delta f$. Given any vertex $x \in G$, let $I_x := I \setminus (N(x) \cup \{x\})$. Rather than directly showing that $\mathbb{E}(\sum_{x\in G} Z_x)$ is large, we will show that $\mathbb{E}(Z_x | I_x)$ is large for every vertex x and every I_x.

Let N_x be the set of all neighbours of x which are not adjacent to a vertex in I_x. We will now show that if N_x is large then the average size of an independent subset of N_x is large as well. So suppose first that $|N_x| \geq 2$. Since $G[N_x]$ is $K_{s,s}$-free, it follows from Theorem 2.6 that every subgraph H of $G[N_x]$ has average degree at most $c_s|H|^{1-1/s} \leq c_s|N_x|^{1-1/s}$. Thus by Proposition 2.4 we have that $\chi(G[N_x]) \leq c_s|N_x|^{1-1/s} + 1 \leq 2c_s|N_x|^{1-1/s}$. So $G[N_x]$ has an independent set of size at least $|N_x|^{1/s}/(2c_s) =: \alpha$. Hence $G[N_x]$ contains at least $2^\alpha/2$ independent sets of size at least $\alpha/2$. Put $\beta := \alpha/(4 \log |N_x|)$. Then the number of independent subsets of N_x of size at most β is at most

$$\binom{|N_x|}{0} + \cdots + \binom{|N_x|}{\lfloor \beta \rfloor} \leq |N_x|^{2\beta} = 2^{2\beta \log |N_x|} = 2^{\alpha/2}.$$

If $|N_x| \geq (4c_s)^s$ then $2^\alpha/2 \geq 2^{\alpha/2}$ and $\alpha/2 \geq 2\beta$; and so in this case the average size ℓ_x of an independent subset of N_x is at least β.

Now note that, writing k_x for the number of independent sets in N_x, for every $|N_x| \geq 0$ we have

$$\mathbb{E}(Z_x | I_x) \geq \frac{\Delta + k_x \ell_x}{1 + k_x} \geq \frac{\Delta}{2k_x} + \frac{\ell_x}{2}.$$

Thus, if $|N_x| \geq (\log \Delta)/2$ and if c' is sufficiently small, then

$$\mathbb{E}(Z_x | I_x) \geq \frac{\ell_x}{2} \geq \frac{\beta}{2} \geq \frac{|N_x|^{1/s}}{16c_s \log |N_x|} \geq \frac{2c'(\log \Delta)^{1/s}}{\log\log \Delta},$$

while if $0 \leq |N_x| \leq (\log \Delta)/2$ then

$$\mathbb{E}(Z_x | I_x) \geq \frac{\Delta}{2 \cdot 2^{|N_x|}} \geq \frac{\Delta}{2 \cdot 2^{(\log \Delta)/2}} = \frac{\sqrt{\Delta}}{2} \geq \frac{2c'(\log \Delta)^{1/s}}{\log\log \Delta}.$$

Hence we have $\mathbb{E}(Z_x) \geq 2\Delta f/n$ and so $\mathbb{E}(\sum_{x\in G} Z_x) = \sum_{x\in G} \mathbb{E}(Z_x) \geq 2\Delta f$, which completes the proof. $\square$

Corollary 2.10 *For every $s \in \mathbb{N}$ there exists $d_0 = d_0(s)$ such that every $K_{s,s}$-free graph G of average degree $d \geq d_0$ contains an independent set of size at least $|G|(\log d)^{1/(s+1)}/d$.*

Proof. Let G' be the subgraph of G induced by the vertices of degree at most $2d$. Clearly, $|G'| \geq |G|/2$. If d is sufficiently large, then by Lemma 2.9, G' (and thus G) has an independent set of size at least $|G|(\log d)^{1/(s+1)}/d$. $\square$

2.4 Finding a 'nearly' induced bipartite subgraph of large average degree

As remarked in the introduction, we would like to find an induced bipartite subgraph of large average degree in our original graph G. The aim of this section is to prove that if G does not contain such a subgraph, we can still come close to it: By Corollary 2.10 we may assume that G contains a large independent set I. We will use this to find a subgraph (A, B) of large average degree so that $A \subseteq I$ (so A is independent) and B has small chromatic number and is much smaller than A. The following lemma shows how to construct one colour class of B.

Lemma 2.11 *Let I be an independent set in a graph G such that $d(x) \geq d/2$ for every $x \in I$ and $|I| = 2c|G|/d$ for some $c \geq 2$. Suppose that $\chi(G) \leq 3d$. Then G has one of the following properties.*

(i) *G contains an induced bipartite subgraph whose average degree is at least $(\log c)/24$.*

(ii) *There are a set $I' \subseteq I$ and an independent set J in $G - I$ such that in G every vertex of I' has exactly one neighbour in J, $|J| \leq |I| \log c/c$ and $|I'| \geq |I|/4(\log c)^2$.*

Proof. Put $n := |G|$, $\overline{I} := V(G) \setminus I$ and let Y be the set of all vertices in $\overline{I}$ which have at least $c/2$ neighbours in I. Then $e(I, \overline{I} \setminus Y) \leq c|\overline{I} \setminus Y|/2 \leq cn/2$. On the other hand the degree of every vertex in I is at least $d/2$, and so we have that $e(I, \overline{I}) \geq cn$. Thus $e(I, Y) \geq cn/2$. As $\chi(G) \leq 3d$, there exists an independent set $A \subseteq Y$ such that

$$e(I, A) \geq \frac{e(I, Y)}{3d} \geq \frac{cn}{6d} = \frac{|I|}{12}. \tag{2.1}$$

Note also that

$$\frac{c}{2} \cdot |A| \leq e(I, A). \tag{2.2}$$

We may assume that the average degree of $(I, A)_G$ is at most $(\log c)/2$ (otherwise $(I, A)_G$ would be as desired in (i)). Since every vertex in A has at least $c/2 \geq (\log c)/2$ neighbours in I, this implies that $|I| \geq |A|$. Therefore

$$\frac{c}{2} \cdot |A| \leq e(I, A) = \frac{1}{2} \cdot d((I, A)_G)(|I| + |A|) \leq \frac{\log c}{2} \cdot |I|,$$

and hence
$$|A| \leq \frac{|I| \log c}{c}. \tag{2.3}$$
Using a probabilistic argument, we will show that there exist sets $J \subseteq A$ and $I' \subseteq I$ as desired in (ii). To make this work, we first need to replace I with the set $I_1 \subseteq I$ of all vertices which have at least one and at most $\log c$ neighbours in A. So let us first estimate the size of I_1. Denote by I_2 the set of all vertices in I which have no neighbours in A and put $I_3 := I \setminus (I_1 \cup I_2)$. We will show that we may assume that both $e(I_1, A) \geq e(I, A)/2$ and $|I_1| \geq |I|/\log c$. Suppose to the contrary that $e(I_1, A) \leq e(I, A)/2$. Then $e(I_3, A) \geq e(I, A)/2$ and so (2.2) implies that $e(I_3, A) \geq c|A|/4$. Thus on average, a vertex in A has at least $c/4$ neighbours in I_3. As every vertex in I_3 has at least $\log c$ neighbours in A, it follows that $(I_3, A)_G$ is as desired in (i). Hence we may assume that $e(I_1, A) \geq e(I, A)/2$. Next suppose that $|I_1| \leq |I|/\log c$. Then
$$e(I_1, A) \geq e(I, A)/2 \overset{(2.1)}{\geq} |I|/24 \geq |I_1|(\log c)/24$$
and
$$e(I_1, A) \geq e(I, A)/2 \overset{(2.2)}{\geq} c|A|/4.$$
Thus $(I_1, A)_G$ is as desired in (i). Therefore we may also assume that $|I_1| \geq |I|/\log c$.

Let us now consider a random subset A_p of A which is obtained by including each $a \in A$ independently with probability $p := 1/(2 \log c)$. Call a vertex $x \in I_1$ *useful* if it has exactly one neighbour in A_p. Using the definition of I_1 it follows that for every $x \in I_1$
$$\begin{aligned}\mathbb{P}(x \text{ is useful}) &= |N(x) \cap A| \cdot p \cdot (1-p)^{|N(x) \cap A| - 1} \geq 1 \cdot p \cdot (1-p)^{\lfloor \log c \rfloor} \\ &\geq p(1 - p\lfloor \log c \rfloor) \geq p/2.\end{aligned}$$
(The second inequality can be easily proved by induction.) Hence the expected number of useful vertices in I_1 is at least $p|I_1|/2$. So there exists a choice J for A_p such that at least $p|I_1|/2$ vertices in I_1 are useful. Let I' be the set of these useful vertices. Then
$$|I'| \geq \frac{p|I_1|}{2} = \frac{|I_1|}{4 \log c} \geq \frac{|I|}{4(\log c)^2}$$
and
$$|J| \leq |A| \overset{(2.3)}{\leq} \frac{|I| \log c}{c}.$$
So I' and J are as desired in (ii). $\square$

By repeated applications of Lemma 2.11 we obtain the following result.

Lemma 2.12 *Let $c \geq 2^{512}$, $d > 2c$ and let G be a graph of minimum degree at least $d/2$. Suppose that $\chi(G) \leq d+1$ and that G has an independent set I of size $2c|G|/d$. Put $r := \lfloor \log \log c \rfloor$. Then G has one of the following properties.*

(i) *G contains an induced bipartite subgraph whose average degree is at least $(\log c)/48$.*

(ii) *There are a set $I^* \subseteq I$ and disjoint independent subsets $J_1, \dots, J_r$ of $G - I^*$ such that every vertex of I^* has exactly one neighbour in each J_k, $|I^*| \geq |I|/4^r(\log c)^{2r}$ and $|J_k| \leq 4|I| \log c/c$ for every $k \leq r$.*

Proof. The proof follows from r applications of Lemma 2.11. Indeed, let $I_0 := I$ and suppose inductively that for some $0 \leq \ell < r$ we already have obtained a set $I_\ell \subseteq I$ and disjoint independent sets $J_1, \dots, J_\ell$ in $G - I_\ell$ such that every vertex of I_ℓ has exactly one neighbour in each J_k, $|I_\ell| = \lceil |I|/4^\ell(\log c)^{2\ell} \rceil$ and $|J_k| \leq 4|I| \log c/c$ for every $1 \leq k \leq \ell$. Put $n := |G|$, $G' := G - (J_1 \cup \dots \cup J_\ell)$, $n' := |G'|$ and $d' := d/2$. Thus $d_{G'}(x) \geq d/2 - \ell \geq d/4 = d'/2$ for every $x \in I_\ell$. Moreover, since $|J_k| \leq 4n \log c/c$, we have that $n' \geq n/2$. Let c' be defined by $|I_\ell| = 2c'n'/d'$. Using $|I_\ell| \leq |I|$ it follows that $c' \leq c$. On the other hand

$$\frac{|I|}{4^\ell(\log c)^{2\ell}} \leq |I_\ell| = \frac{2c'n'}{d'} \leq \frac{4c'n}{d},$$

and so

$$c' \geq \frac{c}{2 \cdot 4^\ell(\log c)^{2\ell}} = \frac{c}{2(2\log c)^{2\ell}}.$$

In particular, $c' \geq 2$. Since also $\chi(G_\ell) \leq d+1 \leq 3d'$, we may apply Lemma 2.11 to the graph G' and the independent set I_ℓ. As

$$\frac{\log c'}{24} \geq \frac{\log c - 1 - \log((2\log c)^{2\ell})}{24} \geq \frac{\log c - 1 - 2r\log(2\log c)}{24} \geq \frac{\log c}{48}$$

we may assume that we have $I_{\ell+1} \subseteq I_\ell$ and $J_{\ell+1}$ satisfying condition (ii) of Lemma 2.11. Hence

$$|I_{\ell+1}| \geq \frac{|I_\ell|}{4(\log c')^2} \geq \frac{|I_\ell|}{4(\log c)^2} \geq \frac{|I|}{4^{\ell+1}(\log c)^{2(\ell+1)}},$$

and

$$|J_{\ell+1}| \leq \frac{|I_\ell| \log c'}{c'} \leq \frac{2 \cdot 4^\ell \cdot |I_\ell|(\log c)^{2\ell+1}}{c} \leq \frac{4|I| \log c}{c}.$$

Note that we may assume that $|I_{\ell+1}| = \lceil |I|/4^{\ell+1}(\log c)^{2(\ell+1)} \rceil$ by making $I_{\ell+1}$ smaller if necessary. This completes the induction step. □

Corollary 2.13 *For every $s \in \mathbb{N}$ there exists $c(s)$ such that the following holds. Let $c \geq c(s)$, $d > 2c$ and let G be a graph of minimum degree at least $d/2$. Suppose that G has an independent set I of size $2c|G|/d$ and that $\chi(G) \leq d+1$. Put $r := \lfloor \log\log c \rfloor$. Then G has one of the following properties.*

(i) *G contains an induced bipartite subgraph whose average degree is at least $(\log c)/48$.*

(ii) *There are disjoint vertex sets $A, B \subseteq V(G)$ such that A is independent, $\chi(G[B]) \leq r$ and $(A, B)_G$ is an $(r, s, 0)$-graph.*

Proof. Applying Lemma 2.12 we may assume that G contains independent sets I^* and $J_1, \dots J_r$ satisfying condition (ii) of Lemma 2.12. Let $A := I^*$ and $B := J_1 \cup \dots \cup J_r$. Clearly, every vertex of A has degree r in the bipartite graph $(A, B)_G$ and $\chi(G[B]) \le r$. Thus it remains to show that $|A| \ge r^{12s}|B|$. But

$$\frac{|A|}{|B|} \ge \frac{c}{4^{r+1} r (\log c)^{2r+1}} \ge r^{12s},$$

if c is sufficiently large. □

2.5 Finding an induced 1-subdivision of a graph of large average degree

In the previous section we showed that we may assume that our original graph G contains a bipartite subgraph (A, B) of large average degree such that A is independent in G and $G[B]$ has small chromatic number (or is possibly independent as well). In this section we will show that this (A, B) contains a 1-subdivision of some graph of large average degree such that this 1-subdivision is induced in G. (Our theorem will then follow immediately from Theorem 2.5.)

To accomplish this, we first find a 1-subdivision of some graph H' of large average degree in (A, B) (Corollary 2.15). The branch vertices of this 1-subdivision are vertices in B, its subdivided edges are paths of length two in (A, B) and so the midpoints of the subdivided edges are vertices in A. In Lemma 2.16 we then show how to find a subgraph H'' of H' for which every midpoint of a subdivided edge is joined in G only to the two endpoints of this edge and to no other branch vertex. As A is independent, it follows that every edge of G which prevents the 1-subdivision of H'' from being induced must join two branch vertices, i.e. two vertices in B. So if B is also independent then this 1-subdivision is induced in G, as desired. The case when B is not independent is more difficult and dealt with in Lemma 2.18.

Let us now introduce some notation. A path P of length two in a bipartite graph (A, B) is called a *hat* of G if it begins and ends in B. A set $\mathcal{H}$ of hats of (A, B) is *uncrowded* if any two hats in $\mathcal{H}$ join distinct pairs of vertices and have distinct midpoints. (So the sets of subdivided edges of the 1-subdivisions of the graphs H' and H'' described above are both uncrowded; and conversely, an uncrowded set of hats can serve as the set of subdivided edges of a 1-subdivision whose set of branch vertices is B.)

Lemma 2.14 *Let $r, i \ge 1$ and $0 \le k \le r/8$. Let $G = (A, B)$ be an (r, i, k)-graph. Then either G has an uncrowded set of at least $r^{11}|B|/2^8$ hats or there are a vertex $b' \in B$ and an induced copy (A', B') of an $(r, i-1, k+1)$-graph in $G - b'$ such that $\emptyset \ne A' \subseteq N_G(b')$.*

Proof. Let us first suppose that every vertex $b \in B$ satisfies

$$|N^2(b)| \ge d(b)/r^{12(i-1)},$$

where $N^2(b)$ is the set of all vertices with distance two from b. In other words, for each $b \in B$ there is a set $\mathcal{H}_b$ of at least $d(b)/r^{12(i-1)}$ hats in G which have b as one endvertex, but whose other endvertices are distinct. Note that every pair of vertices in B belongs to at most two hats in $\bigcup_{b\in B}\mathcal{H}_b$. Hence there are at least $|E(G)|/2r^{12(i-1)}$ hats with distinct pairs of endpoints. Since the degree of every vertex $a \in A$ is at most $4r$, at most $(4r)^2$ of these hats have a as their midpoint. Thus G has a uncrowded set of at least

$$\frac{|E(G)|}{2 \cdot 16r^{12(i-1)+2}} \geq \frac{(r/4-k)|A|}{2^5 r^{12(i-1)+2}} \geq \frac{(r/4-k)r^{12i}|B|}{2^5 r^{12(i-1)+2}} \geq \frac{r^{11}|B|}{2^8}$$

hats, as required.

So we may assume that there is a vertex $b' \in B$ with

$$|N^2(b')| < d(b')/r^{12(i-1)}.$$

Let $A' := N(b')$ and $B' := N^2(b')$. Then $(A', B')_G$ has the required properties.

□

The proof of the preceding lemma shows that in the case where we failed to find a large set of uncrowded hats (i.e. a 1-subdivision of some graph of large average degree), there must be a vertex b' so that the set of vertices with distance two from b' is much smaller that the neighbourhood of b'. However, if this happens we can reapply the lemma to the bipartite graph induced by these sets. In case of renewed failure, we can iterate the process—but if we encounter i successive failures, then this means that G contains contains a $K_{i,i}$:

Corollary 2.15 *Let $s \in \mathbb{N}$ and let $r \geq 8s$. Let $G = (A, B)$ be a $K_{s,s}$-free $(r, s, 0)$-graph. Then there exists $0 \leq i \leq s$ such that G contains an induced copy (A', B') of an $(r, s-i, i)$ graph which has an uncrowded set of at least $r^{11}|B'|/2^8$ hats.*

Proof. Applying Lemma 2.14 repeatedly, assume that there are sequences $(A, B) = (A_0, B_0) \supseteq (A_1, B_1) \supseteq \cdots \supseteq (A_s, B_s)$ of induced subgraphs of G and $b_1, b_2, \ldots, b_s$ of distinct vertices in B such that, for each $0 < i \leq s$, (A_i, B_i) is an $(r, s-i, i)$-graph and $\emptyset \neq A_i \subseteq N_G(b_i)$. Note that every vertex in A_s has degree at least $r/4 - s \geq r/8$ and so

$$s \leq \frac{r}{8} \leq |B_s| = r^{12(s-s)}|B_s| \leq |A_s|.$$

Thus together with any s vertices from A_s the vertices $b_1, \ldots, b_s$ induce a $K_{s,s}$ in G, a contradiction.

□

We say that an uncrowded set $\mathcal{H}$ of hats of a bipartite graph (A, B) is *induced* if $\bigcup \mathcal{H}$ an induced subgraph of (A, B), i.e. if every midpoint of a hat in $\mathcal{H}$ has degree two in (A, B).

Lemma 2.16 *Let $r \geq 1$ and let $G = (A, B)$ be a bipartite graph with $d(a) \leq 4r$ for every vertex $a \in A$. Suppose that G has an uncrowded set $\mathcal{H}$ of at least $r^{11}|B|/2^8$ hats. Then there is an induced subgraph $G' = (A', B')$ of G which has an induced uncrowded set $\mathcal{H}'$ of at least $r^9|B'|/2^{15}$ hats.*

Proof. We may assume that A consists only of midpoints of hats in $\mathcal{H}$. Since $\mathcal{H}$ is uncrowded, every vertex $a \in A$ is the midpoint of exactly one hat in $\mathcal{H}$, and we say that a *owns* the endvertices of these hat. So every vertex in A owns exactly two vertices in B and

$$|A| = |\mathcal{H}| \geq \frac{r^{11}|B|}{2^8}.$$

Let us consider a random subset B_p of B which is obtained by including each vertex of B independently with probability $p := 1/(8r)$. Given B_p, let us call a vertex $a \in A$ *useful* if $N(a) \cap B_p$ consists precisely of the two vertices owned by a. Thus

$$\mathbb{P}(a \text{ is useful}) = p^2(1-p)^{d(a)-2} \geq p^2(1-p)^{\lfloor 4r \rfloor} \geq p^2(1-\lfloor 4r \rfloor p) \geq p^2/2,$$

and so the expected number of useful vertices is at least $p^2|A|/2$. Hence there exists a choice B' for B_p such that at least $p^2|A|/2$ vertices in A are useful. Let A' denote the set of these vertices, and let $\mathcal{H}'$ be the set consisting of all hats in $\mathcal{H}$ whose midpoints lie in A'. Then

$$|\mathcal{H}| = |A'| \geq \frac{|A|}{2^7 r^2} \geq \frac{r^9|B|}{2^{15}} \geq \frac{r^9|B'|}{2^{15}},$$

and so $(A', B')_G$ and $\mathcal{H}'$ have the required properties. $\square$

Corollary 2.17 *Let $s \in \mathbb{N}$ and $r \geq 8s$. Let $G = (A, B)$ be an $(r, s, 0)$ graph. Then either G contains a $K_{s,s}$ or an induced 1-subdivision of some graph H with $d(H) \geq r^9/2^{14}$.*

Proof. We may apply Corollary 2.15 and Lemma 2.16 to obtain an induced bipartite graph $G' = (A', B') \subseteq G$ and a set $\mathcal{H}'$ of hats as in Lemma 2.16. Let H be the graph whose vertex set is B' and in which $b, b' \in B'$ are joined by an edge if there is a hat in $\mathcal{H}'$ whose endvertices are b and b'. So every edge of H corresponds to a hat in $\mathcal{H}'$. As $\mathcal{H}'$ is induced, the 1-subdivision of H is induced in G' (and thus in G). Moreover $|E(H)| = |\mathcal{H}'| \geq r^9|B'|/2^{15}$, as desired. $\square$

Lemma 2.18 *Let $r \geq 2^{25}$. Let A, B be a vertex partition of a graph G such that A is independent, $\chi(G[B]) \leq r$ and $d(G') \leq r^3$ for every $G' \subseteq G[B]$. Suppose that $(A, B)_G$ has an induced uncrowded set $\mathcal{H}$ of at least $r^9|B|/2^{15}$ hats. Then G contains an induced 1-subdivision of some graph H with $d(H) \geq r$.*

Proof. Let H_0 be the graph whose vertex set is B and in which $b, b' \in B$ are joined by an edge if they are the endpoints of a hat in $\mathcal{H}$. Hence G contains a 1-subdivision of H_0. Note that $|E(H_0)| = |\mathcal{H}|$ and so $d(H_0) \geq r^9/2^{14}$. Let H_1 be a subgraph of H_0 with

$$\delta(H_1) \geq \frac{r^9}{2^{15}}, \tag{2.4}$$

and put $B_1 := V(H_1)$ (where B_1 is thought of as a subset of B). Let G^* be the 1-subdivision of H_1 contained in G. Note that every edge which prevents

G^* from being induced must join two branch vertices of G^*, i.e. vertices in B_1. Using a probabilistic argument, we will show that H_1 contains a subgraph H_2 of average degree at least r whose 1-subdivision in G is induced.

Let $F := G[B_1]$ and let B_1' denote the set of all vertices $b \in B_1$ with $d_F(b) \leq 2r^3$. Then

$$2r^3|B_1 \setminus B_1'| \leq 2|E(F)| = d(F)|F| \leq r^3|B_1|$$

and thus

$$|B_1'| \geq \frac{|B_1|}{2}. \tag{2.5}$$

Consider a random subset B_p of B_1 which is obtained by including each vertex of B_1 independently with probability $p = 1/(4r^3)$. Given B_p, call a vertex $b \in B_1'$ *useful* if

(a) $b \in B_p$,

(b) $N_F(b) \cap B_p = \emptyset$,

(c) $|(N_{H_1}(b) \setminus N_F(b)) \cap B_p| \geq pr^9/2^{17}$.

Thus every useful vertex is isolated in $G[B_p]$ and in the graph H_1 it has many neighbours which are contained in B_p. The aim now is to show that with non-zero probability the set I_0 of useful vertices is large. As the chromatic number of $G[B_p]$ is small compared to $|N_{H_1}(b) \cap B_p|$ for any useful vertex b, there will be an independent set in $B_p \setminus I_0$ which together with I_0 induces a subgraph H_2 of H_1 with large average degree. Observe that the 1-subdivision of H_2 in G will be induced.

To prove that with non-zero probability B_1' contains many useful vertices, first note that for every $b \in B_1'$ the random variable $X := |(N_{H_1}(b) \setminus N_F(b)) \cap B_p|$ is binomially distributed with

$$\mathbb{E}X = p|N_{H_1}(b) \setminus N_F(b)| \geq p|\delta(H_1) - d_F(b)| \overset{(2.4)}{\geq} \frac{pr^9}{2^{16}} \geq 8.$$

So Lemma 2.7 implies that

$$\mathbb{P}(X \leq \frac{pr^9}{2^{17}}) \leq \mathbb{P}(X \leq \frac{\mathbb{E}X}{2}) \leq \mathrm{e}^{-\mathbb{E}X/8} \leq \frac{1}{2}.$$

Moreover, note that the events (a), (b) and (c) are mutually independent. Thus for every vertex $b \in B_1'$ we have that

$$\mathbb{P}(b \text{ is useful}) \geq p \cdot (1-p)^{d_F(b)} \cdot \frac{1}{2} \geq p \cdot (1-p)^{\lfloor 2r^3 \rfloor} \cdot \frac{1}{2} \geq \frac{p(1 - \lfloor 2r^3 \rfloor p)}{2} \geq \frac{p}{4}.$$

Hence by (2.5) the expected number of useful vertices is at least $p|B_1'|/4 \geq p|B_1|/8$. So there exists a choice B_2 for B_p such that at least $p|B_1|/8$ vertices in B_1' are useful. Let I_0 denote the set of these vertices. Every useful vertex is

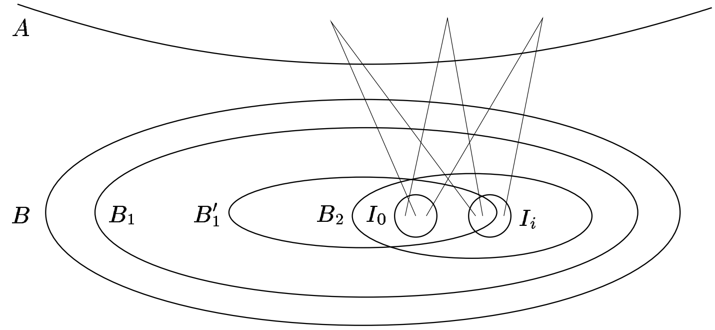

Figure 2.1: Finding an independent set of vertices in F which induces many hats

contained in B_2 and has at least $pr^9/2^{17}$ neighbours in H_1 which are contained in B_2. Thus there are at least

$$\frac{1}{2} \cdot \frac{pr^9}{2^{17}} \cdot \frac{p|B_1|}{8} = \frac{r^3|B_1|}{2^{25}}$$

edges of H_1 which emanate from vertices contained in I_0. Since $\chi(G[B]) \leq r$, we may partition $G[B_2 \setminus I_0]$ into r independent sets, $I_1, \dots, I_r$ say. Then there exists $0 \leq i \leq r$ such that at least a $1/(r+1)$th of the edges of H_1 emanating from I_0 ends in I_i. But then the subgraph H_2 of H_1 induced by $I_0 \cup I_i$ has at least

$$\frac{1}{r+1} \cdot \frac{r^3|B_1|}{2^{25}} \geq \frac{r|B_1|}{2}$$

edges and so it has average degree at least r. Moreover, since in F both I_0 and I_i are independent and no vertex in $B_2 \supseteq I_i$ is joined to a vertex in I_0, it follows that $I_0 \cup I_i$ is independent in G. As mentioned above, this implies that the 1-subdivision of H_2 is induced in G. □

By successively applying Corollary 2.15 and Lemmas 2.16 and 2.18 we obtain the following result.

Corollary 2.19 *Let $s \in \mathbb{N}$ and $r \geq \max\{8s, 2^{25}\}$. Let G be a $K_{s,s}$-free graph and let $A, B \subseteq V(G)$ be disjoint sets of vertices such that A is independent, $\chi(G[B]) \leq r$, $d(G') \leq r^3$ for every $G' \subseteq G[B]$ and so that $(A, B)_G$ is an $(r, s, 0)$ graph. Then G contains an induced 1-subdivision of some graph H with $d(H) \geq r$.* □

2.6 Proof of Theorem 2.1

We can now put everything together.

Proof of Theorem 2.1. Suppose that G is a $K_{s,s}$-free graph with $d(G) = d \geq d_0$ where d_0 is sufficiently large compared to s and $|H|$. Put $n := |G|$. Clearly, we may assume that G has no subgraph of average degree $> d$. So Propositions 2.3

and 2.4 enable us to assume that $\delta(G) \geq d/2$ and $\chi(G) \leq d+1$. Note also that it suffices to show that G contains an induced 1-subdivision of some graph H' with large enough average degree. Indeed, Theorem 2.5 implies that H' contains a subdivision of H and it is easy to show that the corresponding subdivision of H in G is induced. So Lemma 2.8 and Corollary 2.17 imply that Theorem 2.1 holds if G contains an induced bipartite subgraph of large average degree—we will make use of this fact twice in what follows.

Turning to the proof itself, we first apply Corollary 2.10 to G, which gives us an independent set I of size $2cn/d$ where $c \geq (\log d)^{1/(s+1)}/2$. We then apply Corollary 2.13 to obtain (without loss of generality) disjoint sets $A, B \subseteq V(G)$ as in condition (ii) of the corollary. In other words, A is independent, $\chi(G[B]) \leq r$ and $(A, B)_G$ is an $(r, s, 0)$-graph, where $r = \lfloor \log \log c \rfloor$. Now if $G[B]$ has an (induced) subgraph G' whose average degree is at least r^3 then, as $\chi(G') \leq r$, there must be two disjoint independent sets B_1, B_2 of G' such that

$$e((B_1, B_2)_{G'}) \geq \frac{|E(G')|}{\binom{r}{2}} \geq \frac{d(G')|G'|}{r^2} \geq r|G'| \geq r(|B_1| + |B_2|).$$

Hence $(B_1, B_2)_G$ is an induced bipartite subgraph of average degree at least $2r$. So we may assume that $d(G') \leq r^3$ for every $G' \subseteq G[B]$. But then Corollary 2.19 implies that G contains an induced 1-subdivision of some graph H' of large average degree, as desired. □

A result of Thomassen [53] states that for every $k, \ell \in \mathbb{N}$ there exists $f = f(k, \ell)$ such that every graph of minimum degree at least f contains a subdivision of some graph H with minimum degree at least k in which every edge is subdivided exactly ℓ times. Note that, using this result, our proof of Theorem 2.1 gives the following analogue for odd integers ℓ: For every $k, s \in \mathbb{N}$ and every odd integer ℓ there exists $g = g(k, \ell, s)$ such that every $K_{s,s}$-free graph of minimum degree at least g contains an induced subdivision of some graph H with minimum degree at least k in which every edge is subdivided exactly ℓ times.

2.7 Induced cycles

In this section we give a short proof of the special case of Theorem 2.1 when H is a cycle:

Theorem 2.20 *For every $\ell, s \in \mathbb{N}$ there exists $d = d(\ell, s)$ such that every graph of average degree at least d contains either $K_{s,s}$ as a subgraph or an induced cycle of length at least ℓ.*

To prove this result we shall need the case of Theorem 2.2 when T is a path. This special case immediately follows from a theorem of Galvin, Rival and Sands [17, Thm. 4] that every sufficiently large $K_{s,s}$-free graph which has a Hamiltonian path contains a long induced path. The proof of the latter result—an elegant application of Ramsey's theorem—is much simpler than that of the general case of Theorem 2.2. The result itself is also used as a tool in [21] to prove Theorem 2.2.

Proof of Theorem 2.20. We will prove the theorem for $d(\ell, s) = 4 \cdot \lceil \max\{d(P_{s\ell^2}, s), 4s^{s+1}\ell^{2s}, 4(2c_s\ell^2)^s\}\rceil$ where c_s is the constant in Theorem 2.6 and $d(P_{s\ell^2}, s)$ is as defined in Theorem 2.2. By a theorem of Mader (see e.g. [6, Thm. 1.4.2]), every graph G has a $\lfloor d(G)/4 \rfloor$-connected subgraph. So we may assume that we are given an k-connected graph G where $k \geq d(\ell, s)/4$. Thus in particular $d(G) \geq k$, and hence by Theorem 2.2 we may assume that G contains an induced path $P = x \dots y$ of length $s\ell^2$. Menger's theorem now implies that G contains at least k internally disjoint paths joining x to y. Let $\mathcal{Q}$ be the set consisting of all the paths not meeting $\mathring{P}$, where $\mathring{P}$ denotes the interior of P. So $|\mathcal{Q}| \geq k - s\ell^2 \geq k/2$. By short-cutting the paths in $\mathcal{Q}$ if necessary, we may assume that they are induced. Let $\mathcal{Q}'$ be the set consisting of all paths in $\mathcal{Q}$ of length at most ℓ.

Case 1. $|\mathcal{Q}'| \geq k/4$.

Suppose first that there exists a $Q \in \mathcal{Q}'$ such that G has less than $s\ell$ edges joining $\mathring{Q}$ to $\mathring{P}$. In this case, P has a segment S of length at least ℓ such that no edge in G joins $\mathring{S}$ to $\mathring{Q}$. Let S be a maximal such segment. It is easy to show that then G contains an induced cycle consisting of S together with two suitable edges e_1, e_2 joining the endpoints of S to $\mathring{Q}$ and the segment of Q between e_1 and e_2. Since this cycle has length at least ℓ, we may assume that for every $Q \in \mathcal{Q}'$ there are at least $s\ell$ edges joining $\mathring{Q}$ to $\mathring{P}$. Thus in the interior of every $Q \in \mathcal{Q}'$ there must be a vertex x_Q which has at least s neighbours on $\mathring{P}$. Since the number of s-element subsets of $V(\mathring{P})$ is at most

$$\binom{s\ell^2}{s} \leq (s\ell^2)^s \leq \frac{k}{4s} \leq \frac{|\mathcal{Q}'|}{s},$$

there are s paths in $\mathcal{Q}'$, $Q_1, \dots, Q_s$ say, and a set X of s vertices in $\mathring{P}$ such that each x_{Q_i} is joined to all of X. Thus G contains a $K_{s,s}$.

Case 2. $|\mathcal{Q} \setminus \mathcal{Q}'| \geq k/4$.

Let $\mathcal{Q}'' := \mathcal{Q} \setminus \mathcal{Q}'$. For every $Q \in \mathcal{Q}''$ let Q' denote the segment of $\mathring{Q}$ (directed from x to y) consisting of the first ℓ vertices. If there are distinct paths $Q_1, Q_2 \in \mathcal{Q}''$ such that no edge of G joins Q_1' to Q_2' then one can easily show that G has an induced cycle containing Q_1' or Q_2' (or both). Since the length of this cycle is at least ℓ, we may assume that for every distinct $Q_1, Q_2 \in \mathcal{Q}''$ there is an edge joining Q_1' to Q_2'. Consider $G' := G[\bigcup_{Q \in \mathcal{Q}''} V(Q')]$. Then

$$|G'| = \ell |\mathcal{Q}''|$$

and

$$|E(G')| \geq \binom{|\mathcal{Q}''|}{2} \geq \frac{|\mathcal{Q}''|^2}{4}.$$

Together with $|\mathcal{Q}''| \geq k/4 \geq (2c_s\ell^2)^s$ (where c_s is the constant of Theorem 2.6) this implies that

$$d(G') \geq |\mathcal{Q}''|/2\ell \geq c_s \ell |\mathcal{Q}''|^{1-1/s} > c_s |G'|^{1-1/s}.$$

From Theorem 2.6 it now follows that G' contains a $K_{s,s}$. $\square$

2.8 Induced trees

For our alternative proof of Theorem 2.2 we first need some notation. If T is a rooted tree and $i \in \mathbb{N}$ then the *ith level* of T is the set of all vertices of T whose distance to the root of T is exactly i. Given vertices $v, w \in T$ we say that w lies *above* v and v lies *below* w if v lies on the path in T which joins w to the root of T. A vertex w lying above v is said to be a *successor* of v if it is adjacent to v in T. We denote by $\mathrm{br}(v)$ the subtree of T which is induced by all vertices above v and whose root is v. Two vertices of T are *incomparable* if none of them lies above the other. We denote by T_a^b the rooted tree in which every vertex which is not a leaf has exactly a successors and in which every leaf has distance b from the root. Given a graph G and a rooted subtree T of G, we say that T is *pseudo-induced* in G if the endvertices of every edge in $E(G[V(T)]) \setminus E(T)$ are incomparable in T.

Proof of Theorem 2.2. Clearly, it suffices to prove the theorem for all trees of the form T_a^b. So let G be a $K_{s,s}$-free graph of average degree d where d is sufficiently large compared to a, b and s. It is easy to see that G contains every tree T with $|T| \le d/2$ as a (not necessarily induced) subgraph. Thus the following assertion implies that G also contains large pseudo-induced trees:

> *For all $i, j \in \mathbb{N}$ there exist $h = h(i,j)$ and $r' = r'(i,j)$ so that for any $T_{r'}^h \subseteq G$, $G[V(T_{r'}^h)]$ contains a pseudo-induced T_i^j whose leaves are leaves of the $T_{r'}^h$ and in which a vertex w lies above v if and only if w lies above v in the $T_{r'}^h$ (for all $v, w \in V(T_i^j)$).* $\quad (*)$

The proof of $(*)$ proceeds by induction on j. The case $j = 1$ is trivial since every star is pseudo-induced. So suppose that $j > 1$ and that $(*)$ holds for $j-1$. Let $h^* := h(i, j-1)$ and $r^* := r'(i, j-1)$; let r' and h be sufficiently large compared to r^*, h^* and s, and set r' to be much larger than h. Suppose that T is a copy of $T_{r'}^h$ in G.

It is easily seen that T contains a copy T' of $T_{\lceil r'/2 \rceil}^h$ such that in G either all or none of the leaves of T' are adjacent to the root x of T. Indeed, colour a leaf of T blue if it is adjacent to x and green otherwise. Now keep moving towards the root of T, colouring a vertex blue if at least half of its successors are blue and green otherwise, until all the vertices of T are coloured. T must contain a monochromatic copy T' of $T_{\lceil r'/2 \rceil}^h$, which then has the desired property. By repeating this argument, it follows that T contains a copy U of T_r^h such that every vertex $u \in U$ is adjacent in G to either all or none of the leaves above u in U and where $r \ge r'/2^h$.

Let us now assign to each leaf v of U a list of all those integers ℓ for which v is joined in G to the unique vertex below v that lies in the ℓth level of U. These lists never contain more than $s-1$ entries. Indeed, suppose that the list of v has (at least) s entries, $\ell_1 > \cdots > \ell_s$ say. Then for $k = 1, \ldots, s$ the unique vertex v_k lying below v in the ℓ_kth level is joined to v and so, by the choice of U, v_k must be joined to all leaves of U lying above it. In particular, v is a successor of v_1 and each v_k is joined to all the $r \ge s$ successors of v_1. Thus G contains a $K_{s,s}$, contradicting our assumption.

So the above assignment of lists to the leaves of U can be viewed as a colouring of these leaves with at most h^s colours. Since U has exactly r^h leaves, there must be a set R of at least r^h/h^s leaves which all have the same colour, L say. As L has less than s entries, there must be an entry $p \geq h/s$ of L such that $q \notin L$ for all q with $p - h/s < q < p$; and so if h is sufficiently large, this is true for all q with $p - h^* \leq q < p$. Colour a vertex u in the pth level of U red if u is adjacent in G to at least r^* leaves in R lying above u. Moving towards the root of U, colour a vertex red if at least r^* of its successors are red and let R_ℓ be the set of red vertices on the ℓth level ($\ell \leq p$). Since every vertex in the pth level lies below r^{h-p} leaves of U, and since there are r^p vertices in the pth level of U, it follows that

$$|R_p|r^{h-p} + r^p r^* \geq |R| \geq \frac{r^h}{h^s}.$$

This in turn implies that $|R_p| \geq \frac{r^p}{2h^s}$, provided that r is sufficiently large. Similarly, for $q < p$ we have

$$|R_q|r + r^q r^* \geq |R_{q+1}|.$$

If r is sufficiently large, it follows inductively that $|R_q| \geq r^q/(2^{p-q+1}h^s)$ for all $q \leq p$. In particular, we have $|R_{p-h^*}| \geq 1$. Now pick any $u \in R_{p-h^*}$ and apply the induction hypothesis to a red copy of $T_{r^*}^{h^*}$ in $\mathrm{br}(u) \subseteq U$ to find a pseudo-induced T_i^{j-1} as in $(*)$. Since every leaf v of this T_i^{j-1} is a red vertex on the pth level of U, it follows that v is joined in G to $r^* \geq i$ red leaves above v. Let U' denote the extension of the T_i^{j-1} to the copy of T_i^j thus obtained. Recall that for each $\ell = p - h^*, \ldots, p-1$ no leaf $w \in R$ of U is joined in G to a vertex lying below w in the ℓth level of U. Thus U' must be as required in $(*)$, which completes the induction step.

We have shown that if $d = d(G)$ is sufficiently large, then G contains a pseudo-induced copy of $T_{a'}^b$. Thus the theorem will follow if we can show that for every $a, b \in \mathbb{N}$ there exists an a' such that any pseudo-induced copy of $T_{a'}^b$ in G contains a copy of T_a^b that is induced in G. But as already observed by Kierstead and Penrice [21], this can easily be verified by applying Ramsey's theorem and induction on b. Indeed, the case $b = 1$ follows immediately by applying Ramsey's theorem to the subgraph of G spanned by the leaves of a sufficiently large star. Now suppose that $b > 1$ and that the induction hypothesis holds for $b - 1$. Let $a' \in \mathbb{N}$ be sufficiently large and let $v_1, \ldots, v_{a'}$ be an enumeration of the successors of the root of $T_{a'}^b \subseteq G$. The induction hypothesis implies that $\mathrm{br}(v_i)$ contains a copy U_i of T_a^{b-1} that is induced in G ($i = 1, \ldots, a'$). Label the vertices in each U_i, and for $i < j$ let B_{ij} denote the labelled ordered bipartite graph whose first vertex class is $V(U_i)$, whose second vertex class is $V(U_j)$ and whose edges are the U_i–U_j edges in G. Let us now consider the complete graph $K_{a'}$ in which the edge ij is coloured with the graph B_{ij}. So the number of colours only depends on a and b but not on a'. Thus if a' is sufficiently large, then by Ramsey's theorem there exists a monochromatic K_t where $t := \max\{a, 2s\}$. We may assume that the vertices of this K_t are $1, \ldots, t$. Note that G cannot contain a U_i–U_j edge for any $1 \leq i < j \leq t$. For if there exists an edge joining the ℓth vertex of U_i to the kth vertex of U_j, then G contains a $K_{s,s}$ whose vertex classes are formed by the ℓth vertices of $U_1, \ldots, U_s$ and the kth vertices of $U_{s+1}, \ldots, U_{2s}$. Thus the copy of T_a^b formed by the root of the $T_{a'}^b$ and $U_1, \ldots, U_a$ must be induced in G. $\square$

Chapter 3

Subgraphs of large average degree containing no cycle of length less than six

3.1 Introduction

Thomassen [52] conjectured that for all integers k, g there exists an integer $f(k,g)$ such that every graph G of average degree at least $f(k,g)$ contains a subgraph of average degree at least k and girth at least g (where the *girth* of G is the length of the shortest cycle in G). Erdős and Hajnal [16] made a conjecture analogous to that of Thomassen with both occurrences of average degree replaced by chromatic number. The case $g = 4$ of the conjecture of Erdős and Hajnal was proved by Rödl [42], while the general case is still open.

The existence of graphs of both arbitrarily high average degree and high girth follows for example from the result of Erdős that there exist graphs of high girth and high chromatic number. The case $g = 4$ of Thomassen's conjecture (which corresponds to forbidding triangles) is trivial since every graph can be made bipartite by deleting at most half of its edges. Thus $f(k,4) \leq 2k$. The purpose of this chapter is to prove the case $g = 6$ of the conjecture.

Theorem 3.1 *For every k there exists $d = d(k)$ such that every graph of average degree at least d contains a subgraph of average degree at least k whose girth is at least six.*

A straightforward probabilistic argument shows that Thomassen's conjecture is true for graphs G which are almost regular in the sense that their maximum degree is not much larger than their average degree (see Lemma 3.3 for the C_4-case). Indeed, such graphs G do not contain too many short cycles. Thus if we consider the graph G_p obtained by selecting each edge of G with probability p (for a suitable p), it is easy to show that with nonzero probability G_p contains far fewer short cycles than edges. Deleting one edge on every short cycle then yields a subgraph of G with the desired properties.

Thus the conjecture would hold in general if every graph of sufficiently large average degree would contain an almost regular subgraph of large average degree. However, this is not the case: Pyber, Rödl and Szemerédi [37] showed that there are graphs with $cn \log \log n$ edges which do not contain a k-regular subgraph (for all $k \geq 3$). These graphs cannot even contain an almost regular subgraph of large average degree, since e.g. another result in [37] states that every graph with at least $c_k n \log(\Delta(G))$ edges contains a k-regular subgraph. On the other hand, the latter result implies that every graph G with at least $c_k n \log n$ edges contains a k-regular subgraph (which was already proved by Pyber [36]), and thus, if k is sufficiently large, G contains also a subgraph of both high average degree and high girth.

3.2 Proof of Theorem 3.1

All graphs considered in this chapter are finite. We write $e(G)$ for the number of edges of a graph G and $d(G) := 2e(G)/|G|$ for the average degree of G. We say that a graph is C_4*-free* if it does not contain a C_4 as a subgraph. We prove the following quantitative version of Theorem 3.1. (It implies Theorem 3.1 since every graph can be made bipartite by deleting at most half of its edges.) We remark that we have made no attempt to optimize the bounds given in the theorem.

Theorem 3.2 *Let $k \geq 2^{16}$ be an integer. Then every graph of average degree at least $64k^{3+2\cdot 11^{64k^3}}$ contains a C_4-free subgraph of average degree at least k.*

We now give a sketch of the proof of Theorem 3.2. As a preliminary step we find a bipartite subgraph (A, B) of the given graph G which has large average degree and where the vertices in A all have the same degree. We then inductively construct a C_4-free subgraph of (A, B) in the following way. Let $a_1, a_2, \dots$ be an enumeration of the vertices in A. At stage i we will have found a C_4-free subgraph G_i of (A, B) whose vertex classes are contained in $\{a_1, \dots, a_i\}$ and B, and such that the vertices in $V(G_i) \cap A$ all have the same degree in G_i. We then ask whether the subgraph of (A, B) consisting of G_i together with all the edges of G incident with a_{i+1} (and their endvertices) contains many C_4's. If this is the case, the vertex a_{i+1} is 'useless' for our purposes. We then let $G_{i+1} := G_i$ and consider the next vertex a_{i+2}. But if a_{i+1} is not 'useless', we add a_{i+1} together with suitable edges to G_i to obtain a new C_4-free graph G_{i+1}. We then show that either the C_4-free graph G^* consisting of the union of all the G_i has large average degree or else that there is a vertex $x \in B$ and a subgraph (A', B') of $(A, B) - x$ which has similar properties as (A, B) and such that $A' \subseteq N(x)$ (Lemma 3.5). In the latter case, we apply the above procedure to this new graph (A', B'). If this again does not yield a C_4-free subgraph with large average degree, there will be a vertex $x' \in B'$ and a subgraph (A'', B'') of $(A', B') - x'$ as before. So both x and x' are joined in G to all vertices in A''. Continuing this process, we will either find a C_4-free subgraph with large average degree or else a large $K_{s,s}$. But $K_{s,s}$ is regular and so, as was already mentioned in Section 3.1, it contains a C_4-free subgraph as required (Lemma 3.3).

The following lemma implies that Theorem 3.2 holds for the class of all graphs whose maximum degree is not much larger than their average degree. It can easily be generalized to longer cycles.

Lemma 3.3 *If G is a graph of average degree d and maximum degree αd, then G contains a C_4-free subgraph of average degree at least $d^{1/3}/(4\alpha)$.*

Proof. Let $n := |G|$ and put $k := d^{1/3}/(4\alpha)$. Let G_p denote the (random) spanning subgraph of G obtained by including each edge of G in G_p with probability $p := 2k/d$. Let X_4 denote the number of labelled C_4's in G_p and let X_e denote the number of edges in G_p. Then $\mathbb{E}[X_e] = pdn/2$. Since the number of C_4's contained in G is at most $\frac{dn}{2}(\alpha d)^2$ (indeed, every C_4 is determined by first choosing an edge $xy \in G$ and then choosing a neighbour of x and a neighbour of y so that these neighbours are joined by an edge in G), it follows that

$$\mathbb{E}[X_4] \le \frac{dn}{2}(\alpha d)^2 p^4 \le \frac{8\alpha^2 k^3}{d} \cdot p \cdot \frac{dn}{2} \le \mathbb{E}[X_e]/2.$$

Let $X := X_e - X_4$. Then by the above, $\mathbb{E}[X] \ge \mathbb{E}[X_e]/2 = pdn/4 = kn/2$. Thus $\mathbb{P}[X \ge kn/2] > 0$, and so G contains a subgraph H with the property that if we delete an edge from each C_4 in H, the remaining graph H' still has at least $kn/2$ edges. Thus H' is as desired. □

Proposition 3.4 *Let $D > 0$, $0 \le c_0 < 1$ and $c_1 \ge 1$. Let $G = (A, B)$ be a bipartite graph with at least $D|A|$ edges and such that $d(a) \le c_1 D$ for every vertex $a \in A$. Then there are at least $(1 - c_0)/(c_1 - c_0)|A|$ vertices $a \in A$ with $d(a) \ge c_0 D$.*

Proof. Let t denote the number of vertices $a \in A$ with $d(a) \ge c_0 D$. Then $c_1 Dt + c_0 D(|A| - t) \ge e(G) \ge D|A|$, which implies that $t(c_1 D - c_0 D) \ge |A|(D - c_0 D)$. □

Given $c, d \ge 0$, we say that a bipartite graph (A, B) is a *(d, c)-graph* if A is non-empty, $|B| \le c|A|$ and $d(a) = \lceil d \rceil$ for every vertex $a \in A$. Given a graph G and disjoint sets $A, B \subseteq V(G)$, we write $(A, B)_G$ for the induced bipartite subgraph of G with vertex classes A and B.

Lemma 3.5 *Let $c, d \in \mathbb{N}$ be such that d is divisible by c, $c \ge 2^{16}$ and $d \ge 4c^3$. Let $G = (A, B)$ be a $(d/c, c)$-graph. Then G contains either a C_4-free subgraph of average degree at least c or there exists a vertex $x \in B$ and a $(d/c^{11}, c^{11})$-graph $(A', B') \subseteq G$ such that $A' \subseteq N(x)$ and $B' \subseteq B \setminus \{x\}$.*

Proof. Given a bipartite graph (X, Y) and a set $Y' \subseteq Y$, we say that a path P of length two whose endvertices both lie in Y' is a *hat* of Y', and that the endvertices of P *span* this hat.

Let $a_1, a_2, \ldots$ be an enumeration of the vertices in A. Let us define a sequence $A_0 \subseteq A_1 \subseteq \ldots$ of subsets of A and a sequence $G_0 \subseteq G_1 \subseteq \ldots$ of subgraphs of G such that the following holds for all $i = 0, 1, \ldots$:

G_i is C_4-free and has vertex classes $A_i \subseteq \{a_1, \dots, a_i\}$ and B, and $d_{G_i}(a) = 2c^2$ for every $a \in A_i$.

To do this, we begin with $A_0 := \emptyset$ and the graph G_0 consisting of all vertices in B (and no edges). For every $i \geq 1$ in turn, we call the vertex a_i *useless* if $N_G(a_i)$ spans at least $d^2/(8c^4)$ hats contained in G_{i-1}. If a_i is useless, we put $A_i := A_{i-1}$ and $G_i := G_{i-1}$. If a_i is not useless, let us consider the auxiliary graph H on $N_G(a_i)$ in which two vertices $x, y \in N_G(a_i)$ are joined if they span a hat contained in G_{i-1}. Since a_i is not useless, we have that

$$
\begin{aligned}
e(\overline{H}) = \binom{d_G(a_i)}{2} - e(H) &\geq \left(\frac{d/c - 1}{d/c} - \frac{1}{4c^2}\right) \frac{d_G(a_i)^2}{2} \\
&\geq \left(1 - \frac{1}{2c^2}\right) \frac{d_G(a_i)^2}{2},
\end{aligned}
$$

where the last inequality holds since $d \geq 4c^3$. Turán's theorem (see e.g. [6, Thm. 7.1.1.]) applied to $\overline{H}$ now shows that H contains an independent set of size at least $2c^2$. Hence there are $2c^2$ edges of G incident with a_i such that the graph consisting of G_{i-1} together with a_i and these edges does not contain a C_4. We then let G_i be this graph and put $A_i := A_{i-1} \cup \{a_i\}$.

Let $A^* := \bigcup_i A_i$ and $G^* := \bigcup_i G_i$. Thus the accepted graph G^* is C_4-free. Let $A^1 := A \setminus A^*$, and let $G^1 := (A^1, B)_G$. We show that either G^* has average degree at least c (which correponds to Case 1 below) or else that there are $x \in B$ and (A', B') as in the statement of the lemma (Case 2). We will distinguish these two cases according to the properties of the neighbourhoods and the second neighbourhoods of the vertices in B. For this, we need some definitions.

For every $a \in A^1$ consider the auxiliary graph H_a on $N_{G^1}(a) = N_G(a)$ in which two vertices are joined by an edge if they span a hat contained in the accepted graph G^*. Since a is useless, this graph has at least $d^2/(8c^4)$ edges (and d/c vertices), and so it has average degree at least $d/(4c^3)$. By Proposition 2.3, H_a contains a subgraph H'_a with minimum degree at least $d/(8c^3)$, and so with at least $1 + d/(8c^3)$ vertices. Let $B^2 := \bigcup_{a \in A^1} V(H'_a)$, and let G^2 be the subgraph of G^1 whose vertex set is $A^1 \cup B^2$ and in which every $a \in A^1$ is joined to all of $V(H'_a)$. Thus the following holds.

For every $a \in A^1$ we have that $d_{G^2}(a) \geq 1 + d/(8c^3)$, and every vertex in $N_{G^2}(a)$ spans a hat contained in G^* with at least $d/(8c^3)$ other vertices in $N_{G^2}(a)$. $(*)$

Given any vertex $x \in B^2$, let G^2_x denote the subgraph of G^2 induced by the vertices in $A^2_x := N_{G^2}(x)$ and $B^2_x := N_{G^2}(N_{G^2}(x)) \setminus \{x\}$. Let

$$
u := \frac{d}{2^8 c^7},
$$

and say that a vertex $b \in B^2_x$ is *x-rich* if $d_{G^2_x}(b) \geq u$.

Case 1. For every vertex $x \in B^2$ we have that

$$\sum_{b \in B_x^2,\ b \text{ is } x\text{-rich}} d_{G_x^2}(b) \leq \frac{e(G_x^2)}{16c^2}. \tag{3.1}$$

We will show that in this case, every vertex $x \in B^2$ is incident with at least $8c^2 d_{G^2}(x)$ edges of the accepted graph G^* and thus that $e(G^*) \geq 8c^2 e(G^2)$. Before doing this, let us first show that the latter implies that the average degree of G^* is at least c. Indeed, since $e(G^1) = d|A^1|/c$, we have

$$e(G^2) \overset{(*)}{\geq} \frac{d}{8c^3}|A^1| = \frac{1}{8c^2} e(G^1).$$

Thus $e(G^*) \geq e(G^1)$. Also $d_{G^*}(a) = 2c^2$ for every $a \in A^*$ while $d_{G^1}(a) = d/c \geq 2c^2$ for every $a \in A^1$, and so

$$d(G^* \cup G^1) \geq \frac{2 \cdot 2c^2|A|}{|A| + |B|} \geq \frac{4c^2|A|}{(1+c)|A|} \geq 2c.$$

Recalling that $e(G^*) \geq e(G^1)$, this now shows that $d(G^*) \geq d((G^* \cup G^1) - E(G^1)) \geq d(G^* \cup G^1)/2 \geq c$.

Thus it suffices to show that $d_{G^*}(x) \geq 8c^2 d_{G^2}(x)$ for every vertex $x \in B^2$. So let $x \in B^2$, and put $t := d_{G^2}(x) = |A_x^2|$. Let B_x^3 be the subset of B_x^2 obtained by deleting all x-rich vertices, and let $G_x^3 := (A_x^2, B_x^3)_{G_x^2}$. Let $y_1, \ldots, y_t$ be an enumeration of the vertices in A_x^2. For all $i = 1, \ldots, t$, let N_i denote the set of all vertices in $N_{G_x^2}(y_i) = N_{G^2}(y_i) \setminus \{x\}$ spanning a hat with x which is contained in G^*. Hence by $(*)$

$$|N_i| \geq \frac{d}{8c^3}. \tag{3.2}$$

We now use the existence of these hats to show that x is incident with at least $8c^2 t$ edges of G^* (namely edges contained in these hats). Let $N_i' := N_i \cap B_x^3$ and $n_i := |N_i \setminus N_i'|$. Thus $n_i \leq d_{G_x^2}(y_i) - d_{G_x^3}(y_i)$, and so

$$\sum_{i=1}^{t} n_i \leq e(G_x^2) - e(G_x^3) \overset{(3.1)}{\leq} \frac{e(G_x^2)}{16c^2} \leq \frac{dt}{16c^3}.$$

Hence

$$\sum_{i=1}^{t} |N_i'| = \sum_{i=1}^{t} (|N_i| - n_i) \overset{(3.2)}{\geq} \frac{dt}{8c^3} - \frac{dt}{16c^3} = \frac{dt}{16c^3}.$$

But every vertex of G_x^3 lies in at most u of the sets $N_1', \ldots, N_t'$, since $d_{G_x^3}(b) \leq u$ for every $b \in B_x^3$. Thus

$$\left| \bigcup_{i=1}^{t} N_i' \right| \geq \frac{1}{u} \sum_{i=1}^{t} |N_i'| \geq 16c^4 t.$$

That means that x spans hats contained in G^* with at least $16c^4 t$ other vertices in B_x^3. But as every vertex in A^* has degree $2c^2$ in G^*, this implies that x

has at least $16c^4t/(2c^2) \geq 8c^2t$ neighbours in G^*. So we have shown that $d_{G^*}(x) \geq 8c^2 d_{G^2}(x)$ for every $x \in B^2$, as desired.

Case 2. There exists a vertex $x \in B^2$ not satisfying (3.1).

Let B_x^4 be the set of all x-rich vertices in B_x^2, let $G_x^4 := (A_x^2, B_x^4)_{G_x^2}$ and put $t := d_{G^2}(x) = |A_x^2|$. Then the choice of x implies that $t > 0$ and

$$e(G_x^4) \geq \frac{e(G_x^2)}{16c^2} \overset{(*)}{\geq} \frac{1}{16c^2} \cdot \frac{dt}{8c^3} = \frac{dt}{2^7c^5}.$$

Hence the average degree in G_x^4 of the vertices in A_x^2 is at least $D' := d/(2^7c^5)$. Proposition 3.4, applied with $D = D'$, $c_0 = 1/2$ and $c_1 = d/(cD') = 2^7c^4$, now implies that there are at least

$$\frac{1-c_0}{c_1-c_0} \cdot t = \frac{t}{2(c_1 - \frac{1}{2})} \geq \frac{t}{2c_1} = \frac{t}{2^8c^4}$$

vertices $a \in A_x^2$ with $d_{G_x^4}(a) \geq D'/2 \geq d/c^{11}$. Let A_x^4 be the set of these vertices. Thus $|A_x^4| \geq t/(2^8c^4)$. But then the subgraph of $(A_x^4, B_x^4)_{G_x^4}$ obtained by deleting edges so that every vertex in A_x^4 has degree $\lceil d/c^{11} \rceil$ is a $(d/c^{11}, c^{11})$-graph. Indeed, the only thing that remains to be checked is that $|B_x^4| \leq c^{11}|A_x^4|$. But since

$$u|B_x^4| = \frac{d}{2^8c^7}|B_x^4| \leq e(G_x^4) \leq \frac{td}{c} \leq 2^8c^3d|A_x^4|,$$

this follows by recalling that $c \geq 2^{16}$. □

We can now put everything together.

Proof of Theorem 3.2. We may assume (by deleting edges if necessary) that the given graph G has average degree $d := 64k^{3+2\cdot 11^{64k^3}}$. Now we pick a bipartite subgraph G' of G which has average degree at least $d/2$. By Proposition 2.3, there is a subgraph G'' of G' which has minimum degree at least $d/4$. By definition, G'' has a vertex bipartition into A and B so that $|A| \geq |B|$. Let G_0 be the subgraph of G'' obtained by deleting sufficiently many edges to ensure that all vertices in A have degree exactly $d/2^{16}$. We then apply Lemma 3.5 to this G_0. If this fails to produce a C_4-free subgraph of average degree at least k, we obtain a vertex $x_1 \in B_0$ and a $(d/k^{11}, k^{11})$-graph $G_1 = (A_1, B_1)$ with $A_1 \subseteq N_{G_0}(x_1)$ and $B_1 \subseteq B_0 \setminus \{x_1\}$ to which we can apply Lemma 3.5 again. Continuing in this way, after $s := 64k^3$ applications of Lemma 3.5, we either found a C_4-free subgraph of average degree at least k, or sequences $x_1, \dots, x_s$ and $G_1 = (A_1, B_1), \dots, G_s = (A_s, B_s)$, where G_s is a $(d/k^{11^s}, k^{11^s})$-graph. But then each x_i is joined in G to every vertex in A_s. Since A_s is non-empty, we have $|B_s| \geq d/k^{11^s}$ and so in fact

$$|A_s| \geq |B_s|/k^{11^s} \geq d/k^{2\cdot 11^s} = s.$$

Thus G contains the complete bipartite graph $K_{s,s}$. The result now follows by applying Lemma 3.3 to this $K_{s,s}$. □

Chapter 4

Partitions of graphs with high minimum degree or connectivity

4.1 Introduction

It is well-known that the vertex set of every graph G of minimum degree at least $2\ell - 1$ can be partitioned into sets S and T such that the bipartite subgraph $(S, T)_G$ of G between S and T has minimum degree at least ℓ. Moreover, Hajnal [19] and Thomassen [50] proved that for every ℓ there exists $k = k(\ell)$ such that the vertex set of every graph of minimum degree at least k can be partitioned into sets S and T such that the graphs $G[S]$ and $G[T]$ induced by these sets both have minimum degree at least ℓ. They also proved an analogue where minimum degree is replaced by connectivity.

On the other hand, it is easily seen that we cannot simultaneously require that $G[S]$, $G[T]$ and $(S, T)_G$ have large minimum degree or connectivity if G has large minimum degree or connectivity (see Proposition 4.6). In Section 4.2 we show that we can nevertheless strengthen the results of Hajnal and Thomassen by requiring that each vertex in S has many neighbours in T:

Theorem 4.1 *For every $\ell \in \mathbb{N}$ there exists $k = k(\ell) \in \mathbb{N}$ such that the vertex set of every graph G of minimum degree at least k can be partitioned into non-empty sets S and T such that both $G[S]$ and $G[T]$ have minimum degree at least ℓ and every vertex in S has at least ℓ neighbours in T.*

Theorem 4.2 *For every $\ell \in \mathbb{N}$ there exists $k = k(\ell) \in \mathbb{N}$ such that the vertex set of every k-connected graph G can be partitioned into non-empty sets S and T such that both $G[S]$ and $G[T]$ are ℓ-connected and every vertex in S has at least ℓ neighbours in T.*

Theorem 4.2 can be applied to show the existence of non-separating structures in highly connected graphs. A well-known result in this area is the theorem of Thomassen [49] which states that every $(\ell + 3)$-connected graph G contains an induced cycle C such that $G - V(C)$ is ℓ-connected. Applying the result of

Mader that every graph of sufficiently large average degree contains a subdivision of a given graph H to the graph $G[S]$ obtained by Theorem 4.2, it immediately follows that every highly connected graph G contains a non-separating subdivision of H:

Corollary 4.3 *For every $\ell \in \mathbb{N}$ and every graph H there exists $k = k(\ell, H) \in \mathbb{N}$ such that every k-connected graph G contains a subdivision TH of H such that $G - V(TH)$ is ℓ-connected.*

Using Theorem 2.1 of Chapter 2 we obtain the following analogue for induced subdivisions:

Corollary 4.4 *For all $\ell, s \in \mathbb{N}$ and every graph H there exists $k = k(\ell, s, H) \in \mathbb{N}$ such that every k-connected $K_{s,s}$-free graph G contains an induced subdivision TH of H such that $G - V(TH)$ is ℓ-connected.*

Strengthening the theorem of Mader, Larman and Mani [29] showed that every sufficiently highly connected graph G contains a subdivision of a given graph H with prescribed branch vertices. Thus the following conjecture of Thomassen, which greatly strengthens Theorem 4.2, would even imply the existence of a non-separating such subdivision.

Conjecture (Thomassen [51]) *For every $\ell \in \mathbb{N}$ there exists $k = k(\ell) \in \mathbb{N}$ such that if G is a k-connected graph and $X \subseteq V(G)$ consists of at most ℓ vertices then the vertex set of G can be partitioned into non-empty sets S and T such that $X \subseteq S$, each vertex in S has at least ℓ neighbours in T and both $G[S]$ and $G[T]$ are ℓ-connected.*

If true, this would also imply a conjecture of Lovász (see [51]) that there exists a function $f(\ell)$ such that, for any two vertices x and y in an $f(\ell)$-connected graph G, there exists an induced x–y path P such that $G - V(P)$ is ℓ-connected. In Section 4.3 we observe that $f(1) = 3$.

Finally, for graphs of high minimum degree we can strengthen the result of Hajnal and Thomassen in a different direction. It turns out that instead of asking for high minimum degree on one side of the bipartite graph $(S, T)_G$ as in Theorem 4.1, one can require $(S, T)_G$ to have high average degree:

Theorem 4.5 *The vertex set of every graph G of minimum degree at least $2^{32}\ell$ can be partitioned into non-empty sets S and T such that both $G[S]$ and $G[T]$ have minimum degree at least ℓ, $(S, T)_G$ has average degree at least ℓ and $|S|, |T| \geq |G|/2^{18}$.*

It is not hard to obtain similar results about partitions of graphs of high average degree or chromatic number. In Section 4.5 we will present a simple probabilistic argument which shows that for every $\ell \geq 192$ the vertex set of every graph G of average degree at least 40ℓ can be partitioned into S and T such that each of $G[S]$, $G[T]$ and $(S, T)_G$ has average degree at least ℓ. We also observe that the vertex set of every graph G with large chromatic number and large minimum degree can be partitioned into S and T such that both $G[S]$ and $G[T]$ still have large chromatic number and $(S,T)_G$ has large minimum degree.

Let us now recall some notation. Given a graph G, we write $e(G)$ for the number of its edges, $d(G) := 2e(G)/|G|$ for its average degree, $\delta(G)$ for its minimum degree and $\chi(G)$ for its chromatic number. If S and T are disjoint sets of vertices of G, then we denote by $G[S]$ the subgraph of G induced by S and write $(S,T)_G$ for the bipartite subgraph of G whose vertex classes are S and T and whose edges are all the S–T edges in G. We denote by $e_G(S,T)$ or, if this is unambiguous, by $e(S,T)$ the number of these edges and call $(S,T)_G$ the *bipartite subgraph of G between S and T*. We write $N(x)$ for the neighbourhood of a vertex $x \in G$. All graphs considered in this chapter are finite.

4.2 Proof of Theorems 4.1 and 4.2

Before proving Theorems 4.1 and 4.2, let us first observe that these results are best possible in the sense that one cannot additionally require that each vertex in T has many neighbours in S, i.e. that $(S,T)_G$ has large minimum degree.

Proposition 4.6 *For every $k \in \mathbb{N}$ there is a k-connected graph G whose vertex set cannot be partitioned into non-empty sets S and T such that each vertex of G has a neighbour in both S and T.*

Proof. Let G^n_k be the bipartite graph whose vertex classes are $X := \{1,\dots,n\}$ and the set $X^{(k)}$ consisting of all k-element subsets of X, and in which $x \in X$ is joined to $Y \in X^{(k)}$ if $x \in Y$. It is easy to see that G^n_k is k-connected if n is sufficiently large compared to k. Consider any partition of $V(G^n_k)$ into non-empty sets S and T. As we may assume that $n \geq 2k-1$, one of S, T contains at least k vertices from X. Suppose that this is true for S and let $Y \in X^{(k)}$ be a vertex of G^n_k corresponding to a set of k vertices in $X \cap S$. Then Y has no neighbour in T, which proves the proposition. $\square$

In the proof of Theorems 4.1 and 4.2 we will apply the following quantitative versions (see [19, Remark 3 and Thm. 4.3]) of the results of Hajnal and Thomassen mentioned in Section 4.1.

Theorem 4.7 *The vertex set of every graph G of minimum degree at least 3ℓ can be partitioned into non-empty sets S and T such that both $G[S]$ and $G[T]$ have minimum degree at least ℓ.*

Theorem 4.8 *The vertex set of every 8ℓ-connected graph G can be partitioned into non-empty sets S and T such that both $G[S]$ and $G[T]$ are ℓ-connected.*

In the proof of Theorem 4.1 we will also use a lemma which is an easy consequence of Theorem 4.7.

Lemma 4.9 *The vertex set of every graph G of minimum degree at least 3ℓ can be partitioned into non-empty sets S and T such that both $G[S]$ and $G[T]$ have minimum degree at least ℓ and every subgraph of $G[S]$ has average degree less than 6ℓ.*

Proof. Theorem 4.7 implies that $V(G)$ can be partitioned into non-empty sets S and T such that both $\delta(G[S]) \geq \ell$ and $\delta(G[T]) \geq \ell$. Choose S and T satisfying these properties such that S is minimal. Then S, T are as desired. Indeed, if $G[S]$ had a subgraph of average degree at least 6ℓ, then by Proposition 2.3 it would also contain a subgraph H of minimum degree at least 3ℓ. Apply Theorem 4.7 to obtain a partition S', T' of $V(H)$. Let X be a maximal subset of $S \setminus T'$ such that $S' \subseteq X$ and $\delta(G[X]) \geq \ell$. Then each vertex in $S \setminus (T' \cup X)$ has at least $3\ell - \ell = 2\ell$ neighbours in G outside X. Hence X, $V(G) \setminus X$ is a partition contradicting the minimality of S. □

The following lemma shows that if S, T is a partition as provided by Lemma 4.9, then we can alter it by successively moving vertices from S to T to obtain a partition S', T' satisfying Theorem 4.1 except that $G[S']$ is only required to have large average degree. Theorem 4.1 will then immediately follow since $G[S']$, having large average degree, contains a subgraph of large minimum degree. The 'moreover' part of Lemma 4.10 will only be used in the proof of Theorem 4.2.

Lemma 4.10 *Let $c, k, \ell, r \in \mathbb{N}$ be such that $\ell \geq r$ and $k \geq 2^6 c\ell r$. Let G be a graph of minimum degree at least k and let S, T be a partition of $V(G)$ such that $d(G[S]) \geq \ell$, $\delta(G[T]) \geq \ell$ and every subgraph of $G[S]$ has average degree less than $c\ell$. Then there exists $S' \subseteq S$ such that, writing $T' := V(G) \setminus S'$, every vertex in S' has at least r neighbours in T', $d(G[S']) \geq \ell/8$ and $\delta(G[T']) \geq r$. Moreover, T' can be obtained from T by successively adding vertices having at least r neighbours in the superset of T already constructed.*

Proof. Let $A \subseteq S$ be the set of all vertices in S having at least $4c\ell$ neighbours in S, and let $C := S \setminus A$. Thus $4c\ell|A| \leq 2e(G[S])$, and so, as every subgraph of $G[S]$ has average degree less than $c\ell$,

$$e(G[A]) \leq \frac{c\ell|A|}{2} \leq \frac{e(G[S])}{4}. \tag{4.1}$$

Define B to be the set of all those vertices in S which have less than r neighbours in T. (So if $B = \emptyset$, then $S' := S$ and $T' := T$ would be a partition as required in the lemma.) Since $\delta(G) \geq k \geq 4c\ell + r$, we have $B \subseteq A$. Also

$$\frac{k|B|}{2} \leq (k-r)|B| \leq 2e(G[S]). \tag{4.2}$$

Let B' be the set of all those vertices in B which have at least r neighbours in C. For every vertex $x \in B'$ choose a set N_x of r neighbours in C. Let $e_S(\bigcup_{x \in B'} N_x)$ denote the number of all edges in $G[S]$ which are incident to some vertex in $\bigcup_{x \in B'} N_x$. (So here we also count those edges with both endvertices in $\bigcup_{x \in B'} N_x$.) Then

$$e_S(\bigcup_{x \in B'} N_x) \leq |B'| 4c\ell r \overset{(4.2)}{\leq} \frac{16c\ell r \cdot e(G[S])}{k} \leq \frac{e(G[S])}{4}. \tag{4.3}$$

Thus if we add $\bigcup_{x \in B'} N_x$ to T then we can ensure that every vertex in B' has at least r neighbours in the resulting superset of T while maintaining high average degree in the resulting subgraph of $G[S]$. We now have to deal with the vertices in $B \setminus B'$. As some of these vertices may have (nearly) all their neighbours in $B \setminus B'$, but we nevertheless want to ensure that every vertex in S' has many neighbours in T', we will move all of $B \setminus B'$ into T'. For each vertex of $B \setminus B'$ which has most of its neighbours outside $B \setminus B'$ (and thus in $A \setminus (B \setminus B')$), we will also move some of these neighbours into T' in order to ensure that $G[T']$ has large minimum degree. The main difficulty is that the removal of these neighbours should not decrease the average degree of the resulting subgraph of $G[S]$ too much.

So let $A_1 := B \setminus B'$, $A_2 := A \setminus A_1$ and let $A_1' \subseteq A_1$ be the set of all those vertices which have at least $k/2$ neighbours in A_2. Put $C' := C \setminus \bigcup_{x \in B'} N_x$ (Fig. 4.1). We will show that we can add $A_1 \cup \bigcup_{x \in B'} N_x$ to T together with a set $A_2' \subseteq A_2$ such that

(i) each vertex in A_1' has at least r neighbours in A_2',

(ii) $e(A_2', C') \leq e(A_2, C')/2$.

As we shall see, the partition thus obtained is as required in the lemma. However, let us first show that there exists a set A_2' satisfying (i) and (ii).

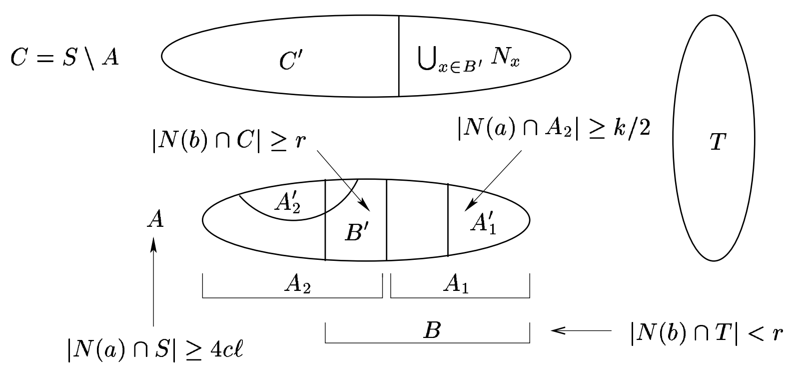

Figure 4.1: The set-up of the proof of Lemma 4.10.

Let $a_1, a_2, \ldots$ be any enumeration of the vertices in A_2 such that their degrees in $(A_2, C')_G$ form a non-increasing sequence. For every $x \in A_1'$ let R_x be the set consisting of the r rightmost neighbours of x in $a_1, a_2, \ldots$. Note that, writing $[a_i]$ for the set $\{a_1, \ldots, a_i\}$,

$$\textit{If } a \in R_x \textit{ then } x \textit{ has at least } k/2 - r \textit{ neighbours in } [a]. \qquad (*)$$

Let $A_2' := \bigcup_{x \in A_1'} R_x$. Then clearly A_2' satisfies condition (i). To show that A_2' also satisfies (ii), we first prove the following

Claim. *For all $a \in A_2$ at most half of the vertices in $[a]$ are contained in A_2'.*

Suppose not and let a be a vertex in A_2 such that more than half of the vertices in $[a]$ are contained in A_2'. Consider the bipartite subgraph H of $G[S]$ between $[a]$ and the set A_1^* consisting of all those vertices $x \in A_1'$ for which $[a] \cap R_x \neq \emptyset$.

Since $[a] \cap A_2'$ contains at most r vertices lying in a common set R_x, it follows that $|A_1^*| \geq |[a]|/(2r)$. Moreover, $(*)$ implies that each vertex in A_1^* has degree at least $k/2 - r \geq k/4$ in H. So if $|A_1^*| \geq |[a]|$, then

$$e(H) \geq \frac{k|A_1^*|}{4} \geq \frac{k|H|}{8} \geq \frac{c\ell|H|}{2};$$

while if $|A_1^*| \leq |[a]|$, we have

$$e(H) \geq \frac{k|A_1^*|}{4} \geq \frac{k|[a]|}{8r} \geq \frac{k|H|}{16r} \geq \frac{c\ell|H|}{2}.$$

Thus in each case H is a subgraph of $G[S]$ of average degree at least $c\ell$, contradicting our assumption. This proves the claim.

By the claim there exists an injection which sends every vertex $a \in A_2'$ to a vertex in $A_2 \setminus A_2'$ preceding a in the sequence $a_1, a_2, \ldots$. The choice of the enumeration $a_1, a_2, \ldots$ now implies (ii).

Put $T' := T \cup A_1 \cup A_2' \cup \bigcup_{x \in B'} N_x$ and $S' := V(G) \setminus T'$. Let us first show that $G[S']$ has large average degree. Clearly,

$$e(G[S']) \geq e(A_2 \setminus A_2', C') \overset{\text{(ii)}}{\geq} e(A_2, C')/2.$$

Thus if $e(A_2, C') \geq e(G[S])/4$ then $e(G[S']) \geq e(G[S])/8$, as required. To verify the other case, note that

$$\begin{aligned} e(G[S']) &\geq e(G[S]) - e(G[A]) - e_S(\bigcup_{x \in B'} N_x) - e(A_1, C') - e(A_2', C') \\ &\overset{(4.1),(4.3)}{\geq} e(G[S]) - \frac{e(G[S])}{4} - \frac{e(G[S])}{4} - |B \setminus B'|r - e(A_2', C') \\ &\overset{(4.2),\text{(ii)}}{\geq} \frac{e(G[S])}{4} - \frac{e(A_2, C')}{2}. \end{aligned}$$

Thus $e(G[S']) \geq e(G[S])/8$ also holds in the case when $e(A_2, C') \leq e(G[S])/4$. This proves that $d(G[S']) \geq \ell/8$. Furthermore, every vertex in $S' \subseteq S \setminus A_1 = S \setminus (B \setminus B')$ has at least r neighbours in $(T \cup \bigcup_{x \in B'} N_x) \subseteq T'$.

Let us now verify that T' can be obtained from T by successively adding vertices having at least r neighbours in the superset of T already constructed. Clearly, as $B \cap \bigcup_{x \in B'} N_x = \emptyset$, $T_1 := T \cup \bigcup_{x \in B'} N_x$ can be constructed in this way. As $A_2' \subseteq (A \setminus B) \cup B'$, every vertex in A_2' has at least r neighbours in T_1, and thus also $T_1 \cup A_2'$ can be constructed. Since by (i) every vertex in A_1' has at least r neighbours in A_2', we can now construct $T_1 \cup A_2' \cup A_1'$. Finally, to construct T, we now successively move vertices from $A_1 \setminus A_1'$ to the other side of the partition as long as they have at least r neighbours there. Suppose that we are not able to exhaust all of $A_1 \setminus A_1'$ in this way, but are left with some non-empty set X. As every vertex in $X \subseteq A_1 \setminus A_1' \subseteq B \setminus B'$ has less than r neighbours in C, less than $k/2$ neighbours in A_2, less than r neighbours in $T' \setminus X$ and as $X \cup C \cup A_2 \cup (T' \setminus X) = V(G)$, it follows that the graph $G[X]$ has minimum degree at least $k - k/2 - 2r \geq c\ell$, contradicting our assumption on $G[S]$. Thus T' can be constructed in the required way. In particular as $\delta(G[T]) \geq \ell \geq r$ this implies that $\delta(G[T']) \geq r$, so S' and T' are as desired.

□

Proof of Theorem 4.1. We will prove the theorem for $k \geq 2^{10} \cdot 6\ell^2 = 2^{11} \cdot 3\ell^2$. Apply Lemma 4.9 to obtain a partition S'', T'' with $\delta(G[S'']) \geq 16\ell$, $\delta(G[T'']) \geq 16\ell$ and such that every subgraph of $G[S'']$ has average degree less than $6 \cdot 16\ell$. Then apply Lemma 4.10 to find a partition S', T' such that $d(G[S']) \geq 2\ell$, $\delta(G[T']) \geq \ell$ and each vertex in S' has at least ℓ neighbours in T'. By Proposition 2.3, the graph $G[S']$ contains a subgraph H of minimum degree at least ℓ. Then $S := V(H)$, $T := V(G) \setminus S$ is a partition as required. $\square$

To prove Theorem 4.2, we need the following theorem of Mader (see e.g. [6, Thm. 1.4.2]) and an analogue to Lemma 4.9.

Theorem 4.11 *Every graph G has a $\lfloor d(G)/4 \rfloor$-connected subgraph.*

Lemma 4.12 *The vertex set of every 9ℓ-connected graph G can be partitioned into non-empty sets S and T such that $G[T]$ is ℓ-connected, the average degree of $G[S]$ is at least ℓ and every subgraph of $G[S]$ has average degree at most 32ℓ.*

Proof. Theorem 4.8 implies that $V(G)$ can be partitioned into non-empty sets S and T such that $d(G[S]) \geq \ell$, $G[T]$ is ℓ-connected and $|T| \geq 2\ell$. Choose S and T satisfying these properties such that S is minimal. We will show that S and T are as desired. First note that $G[S]$ cannot be 8ℓ-connected. For otherwise, by Theorem 4.8, S can be partitioned into two non-empty sets S_1 and S_2 such that each $G[S_i]$ is ℓ-connected. As G was 2ℓ-connected and $|S|, |T| \geq 2\ell$, there are at least 2ℓ independent S–T edges in G. So some S_i, S_2 say, must contain endvertices of at least ℓ of these edges. But then $G[T \cup S_2]$ is ℓ-connected, and so S_1, $T \cup S_2$ is a partition contradicting the minimality of S.

Let us now show that $G[S]$ does not even contain an 8ℓ-connected subgraph (and hence, by Theorem 4.11, no subgraph of average degree at least 32ℓ). Suppose that H is an 8ℓ-connected subgraph of $G[S]$. As $G[S]$ is not 8ℓ-connected, it has a cut-set X with $|X| < 8\ell$. Let C be the unique component of $G[S] - X$ such that $V(H) \subseteq V(C) \cup X$. Let $S' := V(C) \cup X$ and $T' := V(G) \setminus S'$. Then $G[T']$ is ℓ-connected. Indeed, if it has a cut-set Y with $|Y| < \ell$, then, as $G[T]$ is ℓ-connected, there is a component D of $G[T'] - Y$ such that $T \subseteq V(D) \cup Y$. Since there is no path from $F := T' \setminus (V(D) \cup Y)$ to C within $G[S]$ which avoids X and no path in G from F to T which avoids $S' \cup Y = V(C) \cup X \cup Y$, it follows that $X \cup Y$ separates C from F in G, contradicting the fact that G is 9ℓ-connected (Fig. 4.2). We may now successively move vertices from S' to T'

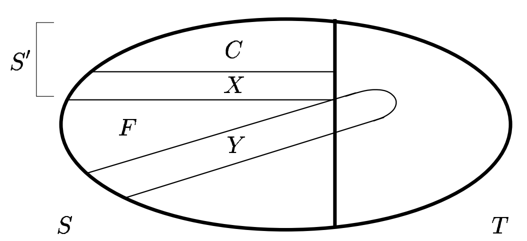

Figure 4.2: $X \cup Y$ separates C from F in G.

if they have less than 8ℓ neighbours in the subset of S' already constructed.

(Thus each of these vertices has at least ℓ neighbours in the superset of T' and therefore adding it preserves the ℓ-connectedness.) This process must terminate as H is a subgraph of $G[S']$ with minimum degree at least 8ℓ. The partition obtained in this way contradicts the minimality of S. $\square$

Proof of Theorem 4.2. We will prove the theorem for $k \geq 2^{16}\ell^2$. First apply Lemma 4.12 to obtain a partition S'', T'' such that $d(G[S'']) \geq 2^5\ell$, $G[T'']$ is $2^5\ell$-connected and such that every subgraph of $G[S'']$ has average degree less than $32 \cdot 2^5\ell$. Now apply Lemma 4.10 to find a partition S' and T' such that each vertex in S' has at least ℓ neighbours in T', $d(G[S']) \geq 4\ell$, $\delta(G[T']) \geq \ell$ and T' can be obtained from T'' by successively adding vertices having at least ℓ neighbours in the superset of T'' already constructed. Thus $G[T']$ is ℓ-connected, since $G[T'']$ is. By Theorem 4.11, $G[S']$ contains an ℓ-connected subgraph H. As each vertex in $S' \setminus V(H)$ has ℓ neighbours in T', $G[T' \cup (S' \setminus V(H))] = G - V(H)$ is ℓ-connected. Thus $S := V(H)$, $T := V(G) \setminus S$ is a partition as required. $\square$

4.3 Non-separating substructures

If H is a structure whose existence is guaranteed by high connectivity, then Theorem 4.2 implies that in every highly connected graph G there exists a copy of H such that $G - V(H)$ is still highly connected. Indeed, if S, T is a partition as in Theorem 4.2, then any copy of H in $G[S]$ will do. Corollary 4.3, which follows by combining Theorem 4.2 with Mader's theorem on subdivisions (Theorem 2.5), was one such example. Similarly, Corollary 4.4 follows by combining Theorem 4.2 with Theorem 2.1. Analogous results for graphs of large minimum degree can be deduced from Theorem 4.1.

The conjecture of Lovász mentioned in Section 4.1 also concerns non-separating structures—in this case an induced path between two prescribed vertices. The following simple argument shows that at least a special case of this conjecture holds.

Proposition 4.13 *For any two vertices x, y of a 3-connected graph G there exists an induced x–y path P in G such that $G - V(P)$ is connected.*

Note that we cannot replace '3-connected' by '2-connected'. Indeed, if x, y denote the vertices of degree three in $G = K_{2,3}$, then the removal of any x–y path from G makes the graph disconnected. In the proof of the proposition we will use the following well-known result of Tutte (see e.g. [6, Thm. 3.2.2]).

Theorem 4.14 *For every 3-connected graph G there is a sequence $G_0, \dots, G_k$ of 3-connected graphs such that, for all $i = 1, \dots, k$, G_{i-1} can be obtained from G_i by contracting one edge of G_i, $G_0 = K_4$ and $G_k = G$.*

Given an x–y path P, we write $\mathring{P}$ for its subpath $P - \{x, y\}$.

Proof of Proposition 4.13. By induction on the length of the sequence in Theorem 4.14. Clearly, the proposition holds if xy is an edge of G and therefore in particular if $G = K_4$. So suppose that $xy \notin E(G)$ and that $e = ab$ is an edge of G such that the graph G' obtained by contracting e is 3-connected and satisfies the proposition. Given a vertex $z \in G$ let v_z denote its image in G', and let $v_e := v_a = v_b$. The 3-connectedness of G' implies that we may assume that a (say) either has more than three neighbours in G or is joined to a vertex of G outside $N(b) \cup \{b\}$. Let P' be an induced v_x–v_y path in G' such that $G' - V(P')$ is connected. Let P be the unique induced x–y path in G whose image in G' is P' and which avoids a if possible (so if $v_e \in \mathring{P}'$ and both a and b are joined to the two neighbours of v_e on P' then P is obtained from P' by replacing v_e with b and not with a). Suppose that $G - V(P)$ is not connected. As $G' - V(P')$ is connected, $V(P)$ clearly cannot be a cut-set of G if $e \in P$ or $e \in G - V(P)$. Thus $v_e \in P'$ but $e \notin P$. Similarly it is easy to see that the endvertex z of e which does not lie on P must form a component of $G - V(P)$. Hence z has all its neighbours on P. Let z' denote the other endvertex of e. Thus P was obtained from P' by replacing v_e with z'. Since P' is induced, v_e has at most two neighbours on P'. As the image in G' of every neighbour of z in $V(P) \setminus \{z'\}$ must be a neighbour of v_e on P' and as $d_G(z) \geq 3$, it follows that $z' \in \mathring{P}$ and $N(z)$ consists of z' and the two neighbours of z' on P. Thus our assumption on a implies that $z = b$, contradicting the choice of P. □

4.4 Proof of Theorem 4.5

Similarly to Section 4.2, before proving the theorem, let us first show that it is best possible in the sense that we cannot always partition the vertex set of a graph of large minimum degree into non-empty sets S and T satisfying Theorems 4.1 and 4.5 simultaneously, i.e. we cannot always find S, T such that both $G[S]$ and $G[T]$ have large minimum degree, $(S,T)_G$ has large average degree and every vertex in S has many neighbours in T.

Proposition 4.15 *For every $k \in \mathbb{N}$ there is a k-connected graph G whose vertex set cannot be partitioned into non-empty sets S and T such that $\delta(G[S]) \geq 1$, $\delta(G[T]) \geq 1$, $d((S,T)_G)) \geq 1$ and each vertex in S has at least one neighbour in T.*

Note that Proposition 4.6 is a special case of Proposition 4.15, but has the advantage that its proof is less technical. Moreover, we will refer to the proof of Proposition 4.6 again in Section 4.5 to deduce a result about partitions of graphs of large chromatic number and large minimum degree.

Proof of Proposition 4.15. Let G_k^n, X and $X^{(k)}$ be as defined in the proof of Proposition 4.6 and suppose that S, T is a partition of $V(G_k^n)$ contradicting Proposition 4.15. Then $0 < |X \cap S| < k$. (Indeed, if $|X \cap S| \geq k$ and Y is a set consisting of any k vertices in $X \cap S$, then all neighbours of the vertex $Y \in X^{(k)}$ are contained in S, contradicting our assumptions on S and T.) Moreover, as any vertex Y in $X^{(k)} \cap S$ must contain at least one of the at most $k-1$ points in $X \cap S$ (when Y is viewed as a k-element subset of X), we have

$$|X^{(k)} \cap S| \le \sum_{i=1}^{k-1} \binom{k-1}{i} \binom{n}{k-i} \le 2^k n^{k-1}.$$

If n is sufficiently large, this implies

$$\begin{aligned} e(S,T) &\le |X \cap S| \cdot \binom{n}{k-1} + |X^{(k)} \cap S| \cdot k \\ &\le k(1+2^k)n^{k-1} < \frac{1}{2} \cdot \binom{n}{k} \le \frac{|G^n_k|}{2}. \end{aligned}$$

Thus $d((S,T)_{G^n_k}) < 1$, a contradiction. $\square$

By Proposition 2.3, every graph G of large average degree contains a subgraph H of large minimum degree. The following lemma implies that if G has many vertices of large degree, then H can be chosen to contain a constant fraction of the vertices in G.

Lemma 4.16 *Let $\ell \in \mathbb{N}$ and $0 < \alpha \le 1$. Let G be a graph and let $X \subseteq V(G)$ be a set of at least $\alpha|G|$ vertices having degree at least ℓ in G. Then G has a subgraph of minimum degree at least $\alpha\ell/4$ which contains at least $|X|/4$ of the vertices in X.*

Proof. By Proposition 2.3, G contains a subgraph of minimum degree at least $\alpha\ell/2$. Let $H \subseteq G$ be maximal with minimum degree at least $\alpha\ell/4$ and suppose that it contains less than $|X|/4$ of the vertices in X. By the choice of H, every vertex of $G - V(H)$ has less than $\alpha\ell/4$ neighbours in H. Thus $G - V(H)$ still has at least $3|X|/4 \ge 3\alpha|G|/4$ vertices having degree at least $\ell - \alpha\ell/4 = (1-\alpha/4)\ell$ in $G - V(H)$. Thus

$$d(G - V(H)) \ge \frac{(3\alpha|G|/4) \cdot (1-\alpha/4)\ell}{|G|} \ge \frac{3\alpha}{4} \cdot \frac{3\ell}{4} \ge \frac{\alpha\ell}{2}.$$

By Proposition 2.3, $G - V(H)$ contains a subgraph H' of minimum degree at least $\alpha\ell/4$. But $H \cup H'$ contradicts the maximality of H. $\square$

As in Chapter 2 we will need a special case of Chernoff's inequality, this time in a slightly different version (see [2, Thm. A.11 and A.13]).

Lemma 4.17 *Let $X_1, \dots, X_n$ be independent 0-1 random variables with $\mathbb{P}(X_i = 1) = p$ for all $i \le n$, and let $X := \sum_{i=1}^n X_i$. Then $\mathbb{P}(|X - \mathbb{E}X| \ge \mathbb{E}X/2) \le 2\mathrm{e}^{-\mathbb{E}X/16}$ and $\mathbb{P}(|X - \mathbb{E}X| \ge \mathbb{E}X/4) \le 2\mathrm{e}^{-\mathbb{E}X/64}$.*

Proof of Theorem 4.5. Let $k := 2^{32}\ell$ and let $(A,B)_G$ be a spanning bipartite subgraph of G of minimum degree at least $k/2$. Suppose that $|A| \ge |B|$. Delete edges if necessary to obtain a bipartite graph H in which each vertex in A has precisely $k/2$ neighbours in B. (Then B may contain some vertices of degree less than $k/2$.) Consider a random partition of B into sets B_1, B_2 where each vertex of B is included in B_1 with probability $1/2$ independently of all other vertices in B. Call a vertex $a \in A$ *good* if it has at least $k/8$ neighbours in each B_i, and *bad* otherwise. Given $a \in A$, write N_a for the number of those neighbours of a in H that lie in B_1. Lemma 4.17 together with the fact that $\mathbb{E}(N_a) = k/4$ implies that

$$\mathbb{P}(a \text{ is bad}) \leq \mathbb{P}(|N_a - \mathbb{E}(N_a)| \geq \mathbb{E}(N_a)/2) \leq 2\mathrm{e}^{-\mathbb{E}(N_a)/16} \leq 1/4.$$

Hence the expected number of bad vertices in A is at most $|A|/4$. Thus there exists a partition B_1, B_2 of B such that the set A' of good vertices in A satisfies

$$|A'| \geq \frac{3|A|}{4}. \tag{4.4}$$

Let $H' := H - (A \setminus A')$. Now consider a random partition of A' into A'_1 and A'_2 where again each vertex of A' is included in A'_1 with probability $1/2$ independently of all other vertices in A'. Then, by Lemma 4.17,

$$\begin{aligned}\mathbb{P}\left(\left||A'_1| - |A'|/2\right| \geq |A'|/8\right) &= \mathbb{P}\left(\left||A'_1| - \mathbb{E}(|A'_1|)\right| \geq \mathbb{E}(|A'_1|)/4\right) \\ &\leq 2\mathrm{e}^{-\mathbb{E}(|A'_1|)/64} \leq 1/3. \end{aligned} \tag{4.5}$$

Let us say that a vertex $b \in B$ is *good* if $d_{H'}(b) \geq 3k/2^{17}$ and if each A'_i contains at least one quarter of the neighbours of b in H', otherwise call b *bad*. Given a vertex $b \in B$, let N'_b denote the number of those neighbours of b in H' that lie in A'_1. Using Lemma 4.17 again, if $d_{H'}(b) \geq 3k/2^{17}$ we have

$$\mathbb{P}(b \text{ is bad}) \leq \mathbb{P}(|N'_b - \mathbb{E}(N'_b)| \geq \mathbb{E}(N'_b)/2) \leq 2\mathrm{e}^{-\mathbb{E}(N'_b)/16} \leq 1/2^{14}.$$

Call an edge of H' *good* if it is incident to a good vertex of B, and *bad* otherwise. Let Y denote the number of bad edges of H'. Let E' be the set of all those edges in H' whose endvertex in B has degree less than $3k/2^{17}$ in H'. As

$$|E'| \leq \frac{3k|B|}{2^{17}} \leq \frac{3k|A|}{2^{17}} \overset{(4.4)}{\leq} \frac{k|A'|}{2^{15}} = \frac{e(H')}{2^{14}},$$

it follows that

$$\begin{aligned}\mathbb{E}(Y) &= \sum_{e \in E(H')} \mathbb{P}(\text{the endvertex of } e \text{ in } B \text{ is bad }) \\ &\leq |E'| + \sum_{e \in E(H') \setminus E'} \mathbb{P}(\text{the endvertex of } e \text{ in } B \text{ is bad}) \\ &\leq |E'| + e(H')/2^{14} \leq e(H')/2^{13}.\end{aligned}$$

Now Markov's inequality implies that

$$\mathbb{P}(Y \geq e(H')/2^{12}) \leq \mathbb{P}(Y \geq 2\mathbb{E}Y) \leq 1/2.$$

Together with (4.5) this implies that there exists a partition A'_1, A'_2 of A' such that at most $e(H')/2^{12}$ of the edges of H' are bad and such that $\left||A'_1| - |A'|/2\right| \leq |A'|/8$. Thus

$$|A'_i| \geq \frac{3|A'|}{8} \tag{4.6}$$

for $i=1,2$. We will now show that there exist sets $A_i^*\subseteq A_i'$ and $B_i^*\subseteq B_i$ such that each $(A_i^*,B_i^*)_{H'}$ has large minimum degree and $(A_2^*,B_1^*)_{H'}$ contains a constant fraction of the edges of H'. Once we have obtained these sets, we will suitably extend $A_1^*\cup B_1^*$ to S and $A_2^*\cup B_2^*$ to T. To find A_1^* and B_1^*, consider the graph $H_1:=(A_1',B_1)_{H'}$. As every vertex in A_1' has at least $k/8$ neighbours in B_1 and

$$|A_1'| \overset{(4.6)}{\geq} \frac{3|A'|}{8} \overset{(4.4)}{\geq} \frac{|A|}{4} \geq \frac{|G|}{8} \geq \frac{|H_1|}{8}, \tag{4.7}$$

we may apply Lemma 4.16 with $X:=A_1'$ and $\alpha:=1/8$ to obtain sets $A_1^*\subseteq A_1'$ and $B_1^*\subseteq B_1$ such that the minimum degree of $(A_1^*,B_1^*)_{H'}$ is at least $k/2^8$ and

$$|A_1^*| \geq \frac{|A_1'|}{4} \overset{(4.7)}{\geq} \frac{|G|}{2^5}. \tag{4.8}$$

Thus

$$e_{H'}(A_1^*,B_1^*) \geq \frac{k|A_1^*|}{2^8} \overset{(4.8)}{\geq} \frac{k|A_1'|}{2^{10}} \overset{(4.6)}{\geq} \frac{3k|A'|}{2^{13}}.$$

But as at most $e(H')/2^{12}=k|A'|/2^{13}$ edges in $H'\supseteq(A_1^*,B_1^*)_{H'}$ are bad, it follows that at least $k|A'|/2^{12}$ of the edges from $(A_1^*,B_1^*)_{H'}$ are good in H'. In other words,

$$\sum_{b\in B_1^*,\, b \text{ is good}} d_{H'}(b) \geq k|A'|/2^{12},$$

and therefore there are also at least $k|A'|/2^{14}$ edges in H' joining (the good vertices in) B_1^* to A_2'.

Now consider the graph $(A_2',B_1^*)_{H'}$ and let $A_2^\diamond$ be the set of all vertices in A_2' which have at least $k/2^{15}$ neighbours in B_1^*. Then

$$\frac{k|A_2^\diamond|}{2}+\frac{k|A'|}{2^{15}} \geq e_{H'}(A_2',B_1^*) \geq \frac{k|A'|}{2^{14}},$$

and therefore

$$|A_2^\diamond| \geq \frac{|A'|}{2^{14}} \overset{(4.4)}{\geq} \frac{|A|}{2^{15}} \geq \frac{|G|}{2^{16}}. \tag{4.9}$$

Now consider the graph $(A_2^\diamond,B_2)_{H'}$. As in this graph the degree of every vertex in $A_2^\diamond$ is at least $k/8$, we may apply Lemma 4.16 with $X:=A_2^\diamond$ and $\alpha:=1/2^{16}$ to obtain sets $A_2^*\subseteq A_2^\diamond$ and $B_2^*\subseteq B_2$ such that $(A_2^*,B_2^*)_{H'}$ has minimum degree at least $k/2^{21}$ and

$$|A_2^*| \geq \frac{|A_2^\diamond|}{4} \overset{(4.9)}{\geq} \frac{|G|}{2^{18}}.$$

Since every vertex in $A_2^\diamond\supseteq A_2^*$ has at least $k/2^{15}$ neighbours in B_1^*, it follows that

$$e_{H'}(A_2^*,B_1^*) \geq \frac{k|A_2^*|}{2^{15}} \geq \frac{k|G|}{2^{33}}.$$

Let $S':=A_1^*\cup B_1^*$ and $T':=A_2^*\cup B_2^*$. Then $|S'| \overset{(4.8)}{\geq} |G|/2^5$, $|T'|\geq|G|/2^{18}$, $\delta(G[S'])\geq k/2^8\geq\ell$, $\delta(G[T'])\geq k/2^{21}\geq\ell$ and $(S',T')_G$ has at least $k|G|/2^{33}$ edges. Choose $S\supseteq S'$ and $T\supseteq T'$ such that S and T are disjoint, both $G[S]$ and $G[T]$ have minimum degree at least $k/2^{21}$ and so that $S\cup T$ is maximal.

Then $S \cup T = V(G)$. Indeed, otherwise every vertex in $G - (S \cup T)$ has at least $k - 2k/2^{21}$ neighbours outside $S \cup T$, i.e. $\delta(G - (S \cup T)) \geq k/2$. But then $S, V(G) \setminus S$ is a partition contradicting the choice of S, T. Note that $e_G(S,T) \geq e_G(S',T') \geq e_{H'}(A_2^*, B_1^*)$ and so $d((S,T)_G) \geq k/2^{32} = \ell$, as desired. □

4.5 Partitions with constraints on the average degree or the chromatic number

For completeness we also present a partition result for graphs of large average degree: the vertex set of every such graph can be partitioned into S and T such that each of $G[S]$, $G[T]$ and $(S,T)_G$ still has large average degree. The straightforward probabilistic argument used to prove this proposition was already one ingredient in the proof of Theorem 4.5.

Proposition 4.18 *For any $d \geq 96$, the vertex set of every graph G with at least $40d|G|$ edges can be partitioned into two non-empty sets S and T such that $G[S]$ and $G[T]$ as well as $(S,T)_G$ have at least $d|G|$ edges.*

Proof. Let $n := |G|$ and let A, B be a partition of $V(G)$ so that $e(A,B) \geq e(G)/2$. Let $A' \subseteq A$ be the set of vertices in A which have at least d neighbours in B. Now consider a random partition of B into sets B_1, B_2 where a vertex $b \in B$ is included in B_1 with probability $1/2$ independently of all the other vertices in B. We call a vertex $a \in A'$ *good* if each B_i contains at least one quarter of its neighbours in B, and *bad* otherwise. Given a vertex $a \in A'$, let N_a be the number of neighbours of a in B_1. Lemma 4.17 now implies that

$$\mathbb{P}(a \text{ is bad}) = \mathbb{P}(|N_a - \mathbb{E}(N_a)| \geq \mathbb{E}(N_a)/2) \leq 2\mathrm{e}^{-\mathbb{E}N_a/16} \leq 1/4.$$

We call an edge between A and B *good* if it is incident to a good vertex in A', and *bad* otherwise. Let Y be the number of bad edges. As the number of edges between $A \setminus A'$ and B is at most $dn \leq e(A,B)/4$, it follows that

$$\begin{aligned}\mathbb{E}(Y) &\leq e(A \setminus A', B) + \sum_{e \in E((A',B)_G)} \mathbb{P}(\text{the endvertex of } e \text{ in } A' \text{ is bad}) \\ &\leq e(A,B)/4 + e(A',B)/4 \leq e(A,B)/2.\end{aligned}$$

Thus there exists a partition B_1, B_2 of B such that the number of bad edges is at most $e(A,B)/2$. Let A'' be the set of good vertices in A' with respect to this partition. Thus

$$e(A'', B) \geq e(A,B)/2. \tag{4.10}$$

Let $a_1, a_2, \ldots$ be an enumeration of the vertices in A'' such that their degrees in $(A'', B)_G$ form a non-increasing sequence. Put $A_1 := \{a_1, a_3, \ldots\}$ and $A_2 := \{a_2, a_4, \ldots\}$. Then

$$e(A_1, B) \geq e(A'', B)/2 \overset{(4.10)}{\geq} 5dn,$$

and therefore

$$e(A_2, B) \geq e(A_1 \setminus \{a_1\}, B) \geq e(A_1, B) - n \geq 4dn.$$

Since all the vertices in A_i are good, this implies that for each pair i, j we have $e(A_i, B_j) \geq dn$. Thus setting $S := A_1 \cup B_1$ and $T = (A \setminus A_1) \cup B_2$ completes the proof. □

We remark that the same proof, carried out with more care regarding the constants, also shows that for every $\varepsilon > 0$ the vertex set of every bipartite graph with m edges and sufficiently large average degree can be partitioned into S and T such that both $G[S]$ and $G[T]$ contain at least $(1/4 - \varepsilon)m$ edges and $(S,T)_G$ contains at least $(1/2 - \varepsilon)m$ edges (as one would expect for a random partition).

To conclude this section, let us prove a simple result about partitions of graphs of large chromatic number and large minimum degree: the vertex set of every such graph can be partitioned into sets S and T such that both $G[S]$ and $G[T]$ have large chromatic number and $(S,T)_G$ has large minimum degree. Note that we cannot additionally require $G[S]$ and $G[T]$ to have large minimum degree. In fact, there need not even exist a partition S, T such that each of $G[S]$, $G[T]$ and $(S,T)_G$ has minimum degree at least one. (Let H_k^n be the graph obtained from the graph G_k^n defined in the proof of Proposition 4.6 by making $G_k^n[X]$ complete. Then $\chi(H_k^n) \geq n$ and $\delta(H_k^n) = k$, but the proof of Proposition 4.6 shows that there is no partition S, T having the properties in question.)

Proposition 4.19 *The vertex set of every graph G can be partitioned into non-empty sets S and T such that $\chi(G[S]) = \lceil \chi(G)/2 \rceil$, $\chi(G[T]) = \lfloor \chi(G)/2 \rfloor$ and every vertex $x \in G$ has at least $d^*(x) := \min\{\lfloor \chi(G)/2 \rfloor, d(x)/2\}$ neighbours in the partition set not containing x.*

Proof. Let S, T be a partition with $\chi(G[S]) = \lceil \chi(G)/2 \rceil$ and $\chi(G[T]) = \lfloor \chi(G)/2 \rfloor$ and such that $e(S,T)$ is maximal under these conditions. (By partitioning $V(G)$ into the vertices in the first $\lceil \chi(G)/2 \rceil$ colour classes and the remaining $\lfloor \chi(G)/2 \rfloor$ colour classes, it is clear that such a partition exists.) Let x be a vertex in S and suppose that x has less than $d^*(x) \leq \chi(G[T])$ neighbours in T. Then the chromatic number of $G[T \cup \{x\}]$ is still $\chi(G[T])$. As $\chi(G) \leq \chi(G[S \setminus \{x\}]) + \chi(G[T \cup \{x\}])$, this implies that $\chi(G[S \setminus \{x\}]) = \chi(G[S])$. But as $d^*(x) \leq d(x)/2$, we have $e(S \setminus \{x\}, T \cup \{x\}) > e(S,T)$, contradicting the choice of S, T. Similarly it can be shown that each vertex $x \in T$ has at least $d^*(x)$ neighbours in S. Thus S, T is a partition as required. □

It is not hard to show that the vertex set of every graph G of large chromatic number and large average degree can be partitioned into S, T such that both $G[S]$ and $G[T]$ have large chromatic number and each of $G[S]$, $G[T]$ and $(S,T)_G$ has large average degree. The idea is to use an averaging argument to find a set $X \subseteq V(G)$ (consisting of some of the colour classes of G) such that $G[X]$ still has large chromatic number but only a small fraction of the edges of G is incident to vertices in X. Then one can apply Proposition 4.18 to $G - X$ to obtain a partition S', T'. Adding one half of the colour classes of $G[X]$ to each of S' and T' gives a partition of $V(G)$ as desired.

Chapter 5

Forcing complete minors by high external connectivity

5.1 Introduction

A fundamental result of Robertson and Seymour [40] states that a graph has large tree-width if and only if it contains a large grid minor. In their short proof of this theorem, Diestel et al. [9] introduced the concept of externally connected sets. A set $X \subseteq V(G)$ is *externally k-connected* in G if $|X| \geq k$ and, for all subsets $Y, Z \subseteq X$ with $|Y| = |Z| \leq k$, there are $|Y|$ disjoint Y-Z paths in G without inner vertices or edges in $G[X]$. We say that a subgraph H of G is *externally k-connected* in G if $V(H)$ is externally k-connected in G. For example the bottom row of a $k \times k$ grid G is externally k-connected in G. A result from [9] is that a graph G has large tree-width if and only if it contains a large externally highly connected set X. Thus such a set X forces a large grid minor in G, even if $G[X]$ consists of isolated vertices.

One of the central tools in the proof of the Graph Minor Theorem of Robertson and Seymour is the observation that every large externally highly connected grid forces a large complete minor (and thus so do the graphs with sufficiently large tree-width). Indeed, if we take a large grid H and add $\binom{r}{2}$ independent edges in such a way that any endvertex of such an edge has horizontal distance at least r from every other such endvertex as well as distance at least r from the boundary of the grid, then in the resulting graph we can contract suitable (zig-zag) paths in H to vertices of a K_r whose edges are the edges added to H. Now if H is externally highly connected in some other graph G, these additional edges can be found (subdivided) as paths through G. In this chapter, we address the question raised by Grohe [18] whether thinner structures than grids can still force large complete minors in the same way.

Given a graph property $\mathcal{P}$, let us say that a class $\mathcal{H}$ of graphs *forces large minors from $\mathcal{P}$* if for every $r \in \mathbb{N}$ there is a $k \in \mathbb{N}$ such that whenever a graph $H \in \mathcal{H}$ is externally k-connected in another graph G, then G has a minor of order at least r in $\mathcal{P}$. In this terminology, our observation above says that grids force large complete minors, while the result from [9] cited earlier says that the class of all edgeless finite graphs forces large grid minors. And Grohe's question is whether a class $\mathcal{H}$ of graphs 'substantially thinner' than grids can still force large complete minors. For instance:

Problem A *Is there a class $\mathcal{H}$ of bounded tree-width that forces large complete minors?*

We shall see that the answer to this question is yes. So the next problem will be to find, if possible, the 'thinnest' such class $\mathcal{H}$. To make this precise, let us write $\mathcal{H} \preccurlyeq \mathcal{H}'$ for classes $\mathcal{H}$ and $\mathcal{H}'$ of finite graphs if for every $H \in \mathcal{H}$ there exists an $H' \in \mathcal{H}'$ such that H is a minor of H'. If both $\mathcal{H} \preccurlyeq \mathcal{H}'$ and $\mathcal{H}' \preccurlyeq \mathcal{H}$ then $\mathcal{H}$ and $\mathcal{H}'$ are *equivalent.* (For example, the class of grids is the unique least element, up to equivalence, among the classes of unbounded tree-width.) The Graph Minor Theorem implies that there are no infinite strictly descending chains of graph properties $\mathcal{H}$ with respect to $\preccurlyeq$ (see e.g. [12, Lemma 1.5]). Thus even if there is no least class forcing large complete minors, we can still try to find the minimal ones:

Problem B *Determine the $\preccurlyeq$-minimal classes of graphs that force large complete minors.*

In this chapter, we shall settle Problem A in the affirmative by constructing four inequivalent classes $\mathcal{H}'$, $\mathcal{H}^4$, $\mathcal{H}^{2,3}$ and $\mathcal{H}^{3,2}$ of graphs such that each of them has bounded tree-width but forces large complete minors (Theorem 5.1). All these classes are $\preccurlyeq$-minimal with respect to the property of forcing large complete minors (Theorem 5.2). Indeed, I conjecture that, up to equivalence, $\mathcal{H}'$, $\mathcal{H}^4$, $\mathcal{H}^{2,3}$ and $\mathcal{H}^{3,2}$ are the only such $\preccurlyeq$-minimal classes, which would also settle Problem B. But this conjecture remains open.

5.2 Preliminary observations and statement of results

First recall that, as was observed in Section 5.1, every class $\mathcal{H}$ of finite graphs that has unbounded tree-width forces large complete minors. So we can restrict our attention to classes $\mathcal{H}$ whose tree-width is bounded. If even the path-width of the graphs in $\mathcal{H}$ is bounded, say $\mathrm{pw}(H) < \ell$ for all $H \in \mathcal{H}$, it turns out that we can join the vertices of any $H \in \mathcal{H}$ bijectively to the bottom row of a grid to obtain a graph that has no K_r minor for any $r > 3\ell + 4$ but in which H is externally $|H|$-connected. (See [15, Lemma 2.3] for a proof.)

Similarly, the vertices of every outerplanar graph H can be joined bijectively to the bottom row of a grid to obtain a planar graph in which H is externally $|H|$-connected. Thus outerplanar graphs do not even force a K_5 minor. (A graph is *outerplanar* if it can be drawn in the plane so that all its vertices lie on the boundary of the outer face.) Hence if all graphs in $\mathcal{H}$ can be made outerplanar by deleting ℓ vertices in each of them, then the graphs in $\mathcal{H}$ do not force a $K_{\ell+5}$ minor.

Now the graphs of unbounded path-width are precisely those that contain arbitrarily large binary trees as minors (see e.g. [6, p. 260]), while the outerplanar graphs are precisely those that contain neither K_4 nor $K_{2,3}$ as a minor (see e.g. [48]). So if the graphs in $\mathcal{H}$ are to force arbitrarily large complete minors, they must contain unboundedly large binary trees and at the same time unboundedly many copies of K_4 or $K_{2,3}$ as minors.

Our main result in this chapter says that a natural combination of these conditions is also sufficient. Let T_2^n denote the binary tree of height n. Let H'_n be the disjoint union of n graphs each of which is obtained from T_2^n by adding a new vertex and joining it to the leaves of T_2^n. Let H_n^4 be the graph obtained from T_2^n by identifying each of its leaves with a vertex of a K_4 (where the K_4's glued to different leaves of T_2^n are disjoint from each other and from the rest of T_2^n). Let $H_n^{2,3}$ be the graph obtained from T_2^n by identifying each of its leaves with a vertex of a $K_{2,3}$ having degree two, and let $H_n^{3,2}$ be the graph obtained from T_2^n by identifying each of its leaves with a vertex of a $K_{2,3}$ having degree three (where the $K_{2,3}$'s glued to different leaves of T_2^n are disjoint from each other and from the rest of T_2^n). Let $\mathcal{H}'$ be the class consisting of all H'_n, and define $\mathcal{H}^4$, $\mathcal{H}^{2,3}$ and $\mathcal{H}^{3,2}$ similarly. It is easy to show that these classes are incomparable with respect to $\preccurlyeq$. The following theorem states that each of them forces large complete minors.

Theorem 5.1 *Given an integer r, there are integers $k = k(r)$ and $n = n(r)$ with the following property. Whenever a graph G contains an externally k-connected set X such that $G[X]$ has a minor isomorphic to any of H'_n, H_n^4, $H_n^{2,3}$ or $H_n^{3,2}$, there is a K_r minor in G.*

Moreover, each of the four classes $\mathcal{H}'$, $\mathcal{H}^4$, $\mathcal{H}^{2,3}$ and $\mathcal{H}^{3,2}$ is $\preccurlyeq$-minimal with the property of forcing large complete minors:

Theorem 5.2 *If $\mathcal{H}$ is a class of finite graphs which forces large complete minors and if $\mathcal{H} \preccurlyeq \mathcal{H}^*$, where $\mathcal{H}^*$ is one of the classes $\mathcal{H}'$, $\mathcal{H}^4$, $\mathcal{H}^{2,3}$ and $\mathcal{H}^{3,2}$, then $\mathcal{H}$ is equivalent to $\mathcal{H}^*$.*

Our proof of Theorem 5.1 uses methods as developed in [9]. An alternative approach would have been to show that any graph containing H'_n, H_n^4, $H_n^{2,3}$ or $H_n^{3,2}$ for sufficiently large n as an externally highly connected subgraph cannot be nearly embedded in a given surface (see e.g. [15] for definitions). Theorem 5.1 would then follow from the theorem of Robertson and Seymour [41], that characterizes the structure of graphs without a K_r minor (r fixed).

This chapter is organized as follows. In Section 5.3 we show that the graphs in $\mathcal{H}'$ force arbitrarily large complete minors, while in Section 5.4 we prove the same for the graphs in $\mathcal{H}^4$, $\mathcal{H}^{2,3}$ and $\mathcal{H}^{3,2}$. Theorem 5.2 is proved in the final section of this chapter.

5.3 Trees attached to stars

All graphs considered in this chapter are finite, and all trees will have a root. The *binary tree of height* $n \geq 1$ is the tree in which the root has degree two, all leaves have distance n from the root, and all other vertices have degree three.

Let $H_n^\sharp$ be the graph obtained from the binary tree T of height n by adding a new vertex x and joining it to all leaves of T. Thus H_n' is the disjoint union of n copies of $H_n^\sharp$. We call T the *binary tree in* $H_n^\sharp$. The leaves of T will be called *leaves of* $H_n^\sharp$, and x will be its *new vertex.*

Theorem 5.3 *Given an integer r, there exist integers d, f, n with the following property. Whenever a graph G contains an externally d-connected set X such that $G[X]$ has the graph consisting of f disjoint copies of $H_n^\sharp$ as a minor, there is a K_r minor in G.*

For the proof of the theorem we will need the following definitions and lemmas. Given two vertices x and y of a tree T, we say that x is *above* y if y lies on the path from x to the root of T. A vertex x is called *successor* of y, if x is a neighbour of y and lies above y. Two vertices of T are *incomparable* if none of them lies above the other. The *branch above* x is the subtree of T induced by the set of all vertices above x (including x itself). If e is an edge of T, then the *branch above* e is the branch above the heighest endvertex of e. A *branch strictly above* x is a branch above a successor of x. If S is a subtree of a tree T, we take the unique lowest vertex of S in T as the root of S. In what follows, we assume that for any given tree T we have chosen a linear ordering σ_T of its vertices in such a way that for every incomparable pair x, y of vertices in T the vertices of the branch above x either all precede or all succeed those of the branch above y; and if x is above y, then x succeeds y. Thus such an ordering may be obtained by considering a drawing of T. For a subtree of T or a subdivision of T we choose the ordering induced by σ_T.

If $P = x_1 \dots x_n$ is a (directed) path and $1 \le i \le n$, we write $Px_i := x_1 \dots x_i$, $x_iP := x_i \dots x_n$, $P\mathring{x}_i := x_1 \dots x_{i-1}$ and $\mathring{x}_iP := x_{i+1} \dots x_n$ for the appropriate subpaths of P.

Lemma 5.4 *Let T be the binary tree of height $n \ge 2$ and A a set of leaves of T. Let $h \le n$ be a positive integer. If $|A| \ge n^h$, then T contains a subdivision S of a binary tree of height h such that all leaves of S are contained in A and the root of S (when S is viewed as the subdivision of a binary tree) can be taken as the lowest vertex of S in T.*

Proof. Induction on h. If $h = 1$ the assertion holds. Assume that $h > 1$ and the statement is true for smaller values of h. Since $2 + (n-1)|A|/n \le |A|$, there is a vertex x in T such that each of the branches strictly above x contains $\ge \lfloor |A|/n \rfloor \ge n^{h-1}$ leaves in A. The result follows by taking x as the root of S and applying the induction hypothesis to each of the two branches strictly above x. □

An $r \times t$ *pseudogrid* is a graph consisting of r disjoint directed paths $W_1, \dots, W_r$ and t disjoint directed paths $V_1, \dots, V_t$ such that each W_i consists of t consecutive (vertex disjoint) segments, every V_j meets every W_i exactly in its jth segment, and V_j meets W_i before it meets W_{i+1} (for all $1 \le i < r$). The W_i are the *horizontal* and the V_j the *vertical paths* of the pseudogrid. A tree T is *s-attached to a pseudogrid G* if there is a set A of s leaves of T such that in each of them there begins a vertical path of G, and G meets T only in A. A is the set of *attached leaves* of T. T is *tidily s-attached to a pseudogrid* G if T is s-attached to G and the order of the leaves in A (in the restriction of the ordering σ_T on A) corresponds to the order of the vertical paths in G starting in these leaves. Let $T_1, \dots, T_k$ be disjoint trees. We say that $T_1, \dots, T_k$ are [*tidily*] *s-attached to a pseudogrid* G if each T_i is [tidily] s-attached to G and the vertical paths of G starting in T_i either all precede or all succeed those starting in T_{i+1} (for all $1 \leq i < k$).

A family $\mathcal{P} = \{P_1, \dots, P_k\}$ of directed paths in a tree T is *nested* if the P_i are disjoint, each of them joins two leaves of T, and for all $1 \leq i, j \leq k$, the first vertex of P_i precedes the last vertex of P_j in the ordering σ_T.

Lemma 5.5 *Let G be a graph that contains $\binom{r}{2}$ disjoint subgraphs $G_1, \dots, G_{\binom{r}{2}}$ such that each G_i contains a subdivision T_i of the binary tree of height $2r-1$ and $G_i - T_i$ has a component C_i that is joined (by edges) to every leaf of T_i. Let C be the union of the C_i. If $T_1, \dots, T_{\binom{r}{2}}$ are tidily 2^{2r-1}-attached to an $r \times \binom{r}{2}2^{2r-1}$ pseudogrid in $G - C$, then G contains a K_r minor.*

Proof. Note that every T_i has a set of r nested paths. We may join the nested paths of all T_i using suitable paths in the pseudogrid to obtain a set $\mathcal{P}$ of r disjoint paths such that each of them meets every T_i (Fig. 5.1).

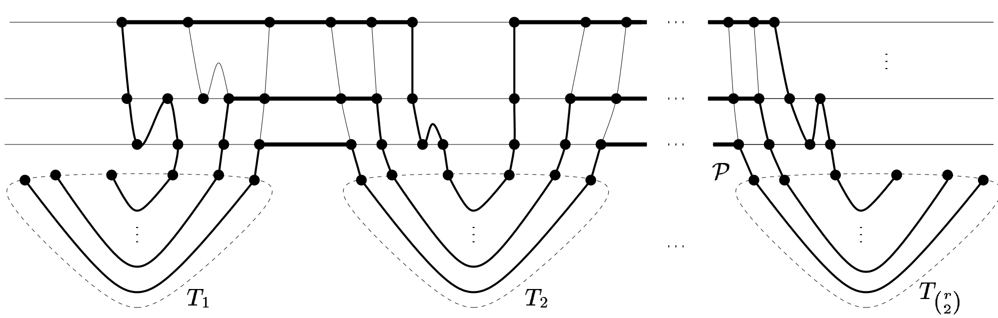

Figure 5.1: Finding a set $\mathcal{P}$ of disjoint paths

The paths in $\mathcal{P}$ are the branch sets of a (subdivided) K_r minor, as any two of them may be joined by a path through one of the C_i. □

Various proofs of the following result of Erdős and Szekeres can be found in [45].

Lemma 5.6 *Every sequence of n distinct integers contains a monotone subsequence of length at least $\sqrt{n}$.*

Lemma 5.7 *Let r, n, s be positive integers such that $n \geq 2$ and $\sqrt{s} \geq n^{2r-1}$, and let $k := \binom{r}{2}$. Let G be a graph that contains k disjoint copies $H_1, \dots, H_k$ of $H_n^\sharp$. For all $1 \leq i \leq k$ let T_i be the binary tree in H_i and x_i its new vertex. Suppose that $T_1, \dots, T_k$ are s-attached to an $r \times sk$ pseudogrid G' in $G - \{x_1, \dots, x_k\}$. Then G contains a K_r minor.*

Proof. Lemma 5.6 applied to the sets of attached leaves of every T_i shows that $T_1, \ldots, T_k$ are tidily $\lceil\sqrt{s}\,\rceil$-attached to an $r \times \lceil\sqrt{s}\,\rceil k$ subpseudogrid of G'. Now Lemma 5.7 follows immediately by first applying Lemma 5.4 and then Lemma 5.5. □

The next lemma is proved in [9, Lemma 6].

Lemma 5.8 *Let $G = (A, B)$ be a bipartite graph, $|A| = a$, $|B| = b$, and let $c \le a$ and $d \le b$ be positive integers. Suppose that G has at most $(a-c)(b-d)/d$ edges. Then there exist $C \subseteq A$ and $D \subseteq B$ such that $|C| = c$ and $|D| = d$ and $C \cup D$ is independent in G.*

For a set I of vertices in a graph G let $N(I)$ denote its neighbourhood. We will also make use of the following easy consequence of Hall's theorem.

Lemma 5.9 *Suppose that $G = (A, B)$ is a bipartite graph such that $s|A| = |B|$ for some positive integer s and $|N(I)| \ge s|I|$ for all subsets $I \subseteq A$. Then G contains $|A|$ disjoint stars, each of them having s edges and their centre in A.*

Proof. Form a new bipartite graph $G' = (A', B)$ by replacing each vertex $a \in A$ by s new vertices and joining each of them to all the neighbours of a. Then G' satisfies Hall's condition, and a matching in G' yields the required disjoint stars. □

Proof of Theorem 5.3. It suffices to show the following assertion.

> *Let $c := 2^{r^5} r^{r+4}$, $s := c^{\lfloor r^2/2 \rfloor - 2}$ and $n := r^2 c^{r^4}$. Let H be the disjoint union of r copies of $H^\sharp_{\lceil \log_2 n + \log_2 s \rceil}$ and sr copies of $H^\sharp_{\lceil \log_2 n \rceil}$. Let G be a graph containing H as an externally nrs-connected subgraph. Then G contains a K_r minor.* (∗)

We may assume that $r \ge 4$. Let A be the set consisting of the r copies of $H^\sharp_{\lceil \log_2 n + \log_2 s \rceil}$ in H, and let B be the set consisting of the sr copies of $H^\sharp_{\lceil \log_2 n \rceil}$ in H. Choose ns leaves of every graph in A, and let Z denote the set consisting of all these leaves. Similarly, choose n leaves of every graph in B, and let Z' denote the set consisting of all these leaves. As H is externally nrs-connected in G, there is a set $\mathcal{Q}$ of $|Z| = nrs$ disjoint Z-Z' paths having no inner vertices in H. Lemma 5.9 implies that we may label the graphs in A by $H_1, \ldots, H_r$ and the graphs in B by H_{ik}, where $1 \le i \le r$ and $1 \le k \le s$, such that for all i, k the number of H_i-H_{ik} paths in $\mathcal{Q}$ is $\ge n/r^2 s$. Indeed, consider the bipartite graph (A, B) in which $S \in A$ is joined to $T \in B$ if there are $\ge n/r^2 s$ S-T paths in $\mathcal{Q}$. We have to check that the assumption of Lemma 5.9 holds. Suppose not, and choose $I \subseteq A$ such that $|N(I)| < s|I|$. Then there are $\ge |I|ns - |N(I)|n \ge n$ paths in $\mathcal{Q}$ which join a graph in I to a graph in $B \backslash N(I)$. Then $\ge n/rs$ of these paths all have one endvertex in the same graph of $B \backslash N(I)$, T say, and $\ge n/r^2 s$ of those paths all have the other endvertex in the same graph of I, S say. But then T is a neighbour of S in (A, B), a contradiction.

For all $i = 1, \ldots, r$ let $\mathcal{Q}_i$ be a set of n/r^2 paths from $\mathcal{Q}$ such that for all $k = 1, \ldots, s$ exactly $n/r^2 s$ of them join H_i to H_{ik}. Choose $n/r^2 s c^{\lfloor r^2/2 \rfloor}$ leaves of H_{ik} that are disjoint from the endvertices of paths from $\mathcal{Q}_i$, and let Y_i be the union (over $k = 1, \ldots, s$) of all these leaves. Since H is externally nrs-connected in G, for all pairs ij with $1 \leq i < j \leq r$ there is a set $\mathcal{P}_{ij}$ of disjoint Y_i-Y_j paths in G having no inner vertices in H. To show $(*)$, we will try to find single paths $P_{ij} \in \mathcal{P}_{ij}$ that are both disjoint for different pairs ij and disjoint from the paths in any $\mathcal{Q}_i$, and thus link up the graphs consisting of H_i together with H_{ik} for all $k = 1, \ldots, s$ and the paths in $\mathcal{Q}_i$ to form a K_r minor in G. If that is not possible, there will be either two sets $\mathcal{P}_{pq}$ and $\mathcal{P}_{ij}$ such that many paths of $\mathcal{P}_{pq}$ meet many paths of $\mathcal{P}_{ij}$, and we shall then use this 'intersection property' to find a K_r minor within the graph consisting of the H_{pk} $(1 \leq k \leq s)$ together with the paths in $\mathcal{P}_{pq} \cup \mathcal{P}_{ij}$, or there will be a $\mathcal{P}_{pq}$ and a $\mathcal{Q}_i$ such that many paths of $\mathcal{P}_{pq}$ meet many paths of $\mathcal{Q}_i$, and in this case we will find a K_r minor within the graph consisting of the H_{ik} $(1 \leq k \leq s)$ together with the paths in $\mathcal{P}_{pq} \cup \mathcal{Q}_i$.

Let $\sigma : \{ij \mid 1 \leq i < j \leq r\} \to \{0, 1, \ldots, \binom{r}{2} - 1\}$ be any bijection. Starting with $\ell = 0$, for successive ℓ and $pq := \sigma^{-1}(\ell)$, we will try to find a path $P \in \mathcal{P}_{pq}$ that is disjoint from the previous selected paths and replace both the later sets $\mathcal{P}_{ij}$ and all sets $\mathcal{Q}_i$ by smaller sets of paths disjoint from P. More precisely, let $\ell^* \leq \binom{r}{2}$ be maximal such that for all $0 \leq \ell < \ell^*$, $1 \leq i \leq r$, $i < j \leq r$ (if $i < r$) there exist sets $\mathcal{P}^\ell_{ij}$ and $\mathcal{Q}^\ell_i$ satisfying the following conditions.

(i) $\mathcal{P}^\ell_{ij}$ is a non-empty set of disjoint Y_i-Y_j paths having no inner vertices in H.

(ii) $\mathcal{Q}^\ell_i$ is a subset of $\mathcal{Q}_i$ of size $|\mathcal{Q}^\ell_i| = n/r^2c^{2\ell}$, and each H_{ik} $(1 \leq k \leq s)$ contains endvertices of $\leq n/r^2sc^\ell$ paths from $\mathcal{Q}^\ell_i$.

As soon as $\mathcal{P}^\ell_{ij}$ and $\mathcal{Q}^\ell_i$ are defined, let H^ℓ_{ij} be the graph consisting of all paths in $\mathcal{P}^\ell_{ij}$, and F^ℓ the graph consisting of all paths contained in some $\mathcal{Q}^\ell_i$. Furthermore, let Y^ℓ_{ij} be the set of all endvertices of paths from $\mathcal{P}^\ell_{ij}$ in Y_i and Y^ℓ_{ji} the set of those in Y_j.

(iii) If $\sigma(ij) < \ell$, then $\mathcal{P}^\ell_{ij}$ has exactly one element P^ℓ_{ij}, and P^ℓ_{ij} is disjoint from any path belonging to a set $\mathcal{P}^\ell_{ab}$ with $ab \neq ij$ and any path belonging to a set $\mathcal{Q}^\ell_a$ (for all a).

(iv) If $\sigma(ij) = \ell$, then $|\mathcal{P}^\ell_{ij}| = n/r^2c^{\ell r^2 + \lfloor r^2/2 \rfloor}$ and each H_{ik} $(1 \leq k \leq s)$ contains endvertices of $\leq n/r^2sc^{\ell r^2 + \lfloor r^2/2 \rfloor - \ell}$ paths from $\mathcal{P}^\ell_{ij}$. Moreover, for every edge $e \in E(H^\ell_{ij}) \backslash E(F^\ell)$ there are no $|\mathcal{P}^\ell_{ij}|$ disjoint Y^ℓ_{ij}-Y^ℓ_{ji} paths in the graph $(H^\ell_{ij} \cup F^\ell) - e$.

(v) If $\sigma(ij) > \ell$, then $|\mathcal{P}^\ell_{ij}| = n/r^2c^{(\ell+1)r^2}$ and each H_{ik} $(1 \leq k \leq s)$ contains endvertices of $\leq n/r^2sc^{(\ell+1)r^2 - \ell}$ paths from $\mathcal{P}^\ell_{ij}$.

(vi) If $\ell = \sigma(pq) < \sigma(ij)$, then for every edge $e \in E(H^\ell_{ij}) \backslash E(H^\ell_{pq})$ there are no $|\mathcal{P}^\ell_{ij}|$ disjoint Y^ℓ_{ij}-Y^ℓ_{ji} paths in the graph $(H^\ell_{ij} \cup H^\ell_{pq}) - e$.

If $\ell^* = \binom{r}{2}$, then by (i)–(iii) we have a (subdivided) K_r minor (the branch sets are the graphs consisting of H_i together with H_{ik} for all $k = 1, \ldots, s$ and all

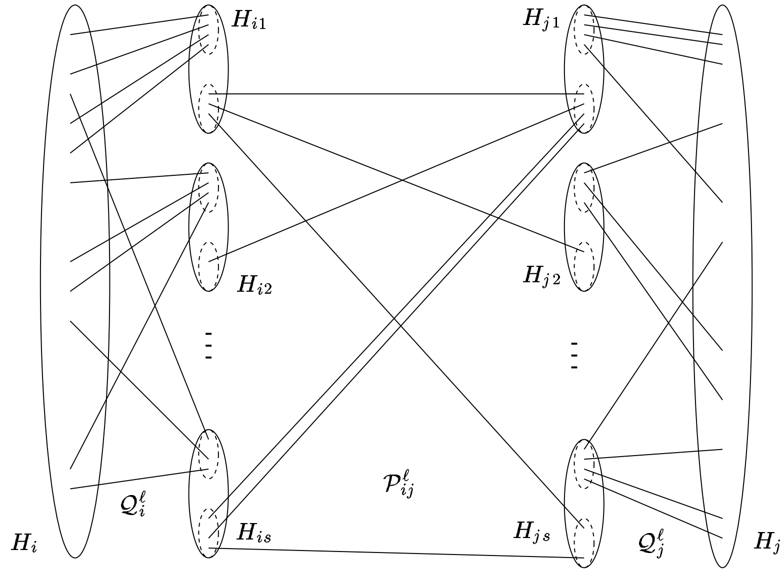

Figure 5.2: The set-up of the proof of Theorem 5.3

paths in $\mathcal{Q}_i^{\binom{r}{2}-1}$). Thus we may assume that $\ell^* < \binom{r}{2}$. To see that $\ell^* > 0$, first note that condition (ii) holds with $\mathcal{Q}_i^0 := \mathcal{Q}_i$. Let $pq := \sigma^{-1}(0)$. Then setting $\mathcal{P}_{pq}^0 := \mathcal{P}_{pq}$ would satisfy condition (i) and the first half of (iv), but may not satisfy its second half. If so, let H_{pq} be the graph consisting of all paths in $\mathcal{P}_{pq}$, and choose $I \subseteq E(H_{pq})\backslash E(F^0)$ maximal such that there are $|\mathcal{P}_{pq}|$ disjoint Y_p-Y_q paths in the graph $(H_{pq} \cup F^0) - I$; then let $\mathcal{P}_{pq}^0$ be such a set of paths. It is easily checked that this choice of $\mathcal{P}_{pq}^0$ satisfies (i) and (iv). For all ij with $\sigma(ij) > 0$, choose a subset of $\mathcal{P}_{ij}$ containing $n/r^2c^{r^2}$ paths such that from every H_{ik} there start exactly $n/r^2sc^{r^2}$ of them. Then the obtained set of paths satisfies condition (v) and may be modified similarly as before to obtain a good choice for $\mathcal{P}_{ij}^0$.

Let $\ell := \ell^* - 1$. Hence conditions (i)–(vi) are satisfied for ℓ, but cannot be satisfied for $\ell + 1$. Let $pq := \sigma^{-1}(\ell)$. We first show that there is no path P in $\mathcal{P}_{pq}^\ell$ that avoids $\geq |\mathcal{P}_{ij}^\ell|/c$ of the paths in $\mathcal{P}_{ij}^\ell$ for all ij with $\sigma(ij) > \ell$ as well as $|\mathcal{Q}_i^\ell|/c$ of the paths in $\mathcal{Q}_i^\ell$ for all i. Suppose there is such a path P. We will show that then we can satisfy conditions (i)–(vi) for $\ell + 1$. Indeed, condition (ii) implies that for all i we may discard paths from $\mathcal{Q}_i^\ell$ avoiding P to find a set $\mathcal{Q}_i^{\ell+1}$ of $|\mathcal{Q}_i^\ell|/c^2$ paths in $\mathcal{Q}_i^\ell$ such that each of them avoids P and each H_{ik} contains endvertices of $\leq n/r^2sc^{\ell+1}$ paths from $\mathcal{Q}_i^{\ell+1}$. Thus $\mathcal{Q}_i^{\ell+1}$ satisfies condition (ii). Similarly, for every ij with $\sigma(ij) > \ell$ we can find a set $\mathcal{P}'_{ij}$ of $|\mathcal{P}_{ij}^\ell|/c^2 = n/r^2c^{(\ell+1)r^2+2}$ paths from $\mathcal{P}_{ij}^\ell$ such that each H_{ik} contains endvertices of $\leq n/r^2sc^{(\ell+1)r^2-\ell+1}$ of these paths and P avoids all of them. If $\sigma(ij) = \ell + 1$, we choose a set $\mathcal{P}''_{ij}$ of $|\mathcal{P}'_{ij}|/c^{\lfloor r^2/2\rfloor-2} = n/r^2c^{(\ell+1)r^2+\lfloor r^2/2\rfloor}$ paths in $\mathcal{P}'_{ij}$ such that each H_{ik} contains endvertices of $\leq n/r^2sc^{(\ell+1)r^2+\lfloor r^2/2\rfloor-(\ell+1)}$ of them. $\mathcal{P}''_{ij}$ can be modified as before to yield a good choice for $\mathcal{P}_{ij}^{\ell+1}$. In a similar way, one can now define $\mathcal{P}_{ij}^{\ell+1}$ for all ij with $\sigma(ij) > \ell+1$, contradicting the maximality of ℓ^*.

Thus for every path P in $\mathcal{P}_{pq}^{\ell}$ there is either a pair ij with $\sigma(ij) > \ell$ such that P avoids $< |\mathcal{P}_{ij}^{\ell}|/c$ paths in $\mathcal{P}_{ij}^{\ell}$, or there is an i such that P avoids $< |\mathcal{Q}_i^{\ell}|/c$ paths in $\mathcal{Q}_i^{\ell}$. Hence there is a set $\mathcal{P}$ of $\geq |\mathcal{P}_{pq}^{\ell}|/r^2$ paths from $\mathcal{P}_{pq}^{\ell}$ together with either a pair ij such that each path in $\mathcal{P}$ avoids $< |\mathcal{P}_{ij}^{\ell}|/c$ paths in $\mathcal{P}_{ij}^{\ell}$ or an i such that each path in $\mathcal{P}$ avoids $< |\mathcal{Q}_i^{\ell}|/c$ paths in $\mathcal{Q}_i^{\ell}$.

Case 1. Each path in $\mathcal{P}$ avoids less than $|\mathcal{Q}_i^{\ell}|/c$ paths in $\mathcal{Q}_i^{\ell}$.

We first use Lemma 5.8 to find a set $\mathcal{V}$ of $\geq \lceil |\mathcal{Q}_i^{\ell}|/2 \rceil = |\mathcal{Q}_i^{\ell}|/2$ paths from $\mathcal{Q}_i^{\ell}$ and a set $\mathcal{W}$ of r paths from $\mathcal{P}$ such that every path in $\mathcal{V}$ meets every path in $\mathcal{W}$. Indeed, to show the existence of such sets we have to check that the bipartite graph $(\mathcal{Q}_i^{\ell}, \mathcal{P})$, in which $Q \in \mathcal{Q}_i^{\ell}$ is joined to $P \in \mathcal{P}$ if P avoids Q, has not too many edges. But this follows, since the number of edges of this bipartite graph is at most

$$|\mathcal{Q}_i^{\ell}||\mathcal{P}|/c \;\leq\; |\mathcal{Q}_i^{\ell}||\mathcal{P}|/4r \;\leq\; (|\mathcal{Q}_i^{\ell}| - |\mathcal{Q}_i^{\ell}|/2)(|\mathcal{P}| - r)/r.$$

We now find a set $\mathcal{V}'$ of $\geq |\mathcal{V}|/r^r$ paths from $\mathcal{V}$ and a labelling of the paths from $\mathcal{W}$ by $W_1, \dots, W_r$ such that on its way from $\bigcup_{k=1}^{s} H_{ik}$ to H_i every path from $\mathcal{V}'$ meets W_a before it meets W_{a+1} (for all $1 \leq a < r$). Recall that, by condition (ii), each H_{ik} ($k = 1, \dots, s$) contains endvertices of $\leq n/r^2 s c^{\ell}$ paths from $\mathcal{Q}_i^{\ell}$, and, note that $r \geq 4$ implies that $\binom{r}{2} \leq \lfloor r^2/2 \rfloor - 2$. Together with

$$|\mathcal{V}'|/2 \;\geq\; |\mathcal{Q}_i^{\ell}|/4r^r \;=\; \frac{n}{4r^{r+2}c^{2\ell}} \;\geq\; \frac{n}{c^{\ell+\binom{r}{2}}} \;\geq\; r^2 \frac{n}{r^2 s c^{\ell}}$$

and

$$|\mathcal{V}'|/2 \;\geq\; \frac{n}{c^{\ell r^2+1}} \;\geq\; (s - r^2)c\frac{n}{r^2 c^{\ell r^2 + \lfloor r^2/2 \rfloor}} \;=\; (s - r^2)c|\mathcal{P}_{pq}^{\ell}|,$$

this implies that there are r^2 of the H_{ik}, without loss of generality $H_{i1}, \dots, H_{ir^2}$, such that each of them contains endvertices of $\geq c|\mathcal{P}_{pq}^{\ell}|$ paths from $\mathcal{V}'$. For all $k = 1, \dots, r^2$, let $\mathcal{V}_k$ be the set of all paths from $\mathcal{V}'$ begining at H_{ik}. We shall now prove that the paths from $\mathcal{W}$ together with many paths from each of $\mathcal{V}_k$ form a pseudogrid (the paths from $\mathcal{W}$ will be its horizontal paths and those from the $\mathcal{V}_k$ its vertical paths). The result will then follow from Lemma 5.7.

Direct W_1 from Y_{pq}^{ℓ} to Y_{qp}^{ℓ}. Let e^1 be the first edge of W_1 such that the initial component W^1 of $W_1 - e^1$ meets $\geq c|\mathcal{P}_{pq}^{\ell}|/r^2$ paths from $\mathcal{V}_k$ for some k, and so that e^1 does not lie on one of these paths. Without loss of generality we may assume that $k = 1$. Note that $e^1 \notin E(F^{\ell})$, since $\mathcal{V}_1 \subseteq \mathcal{Q}_i^{\ell}$ and the paths in $\bigcup_{a=1}^{r} \mathcal{Q}_a^{\ell}$ are disjoint. Let e^2 be the first edge of $W_1 - W^1 - e^1 =: W'$ such that the initial component W^2 of $W' - e^2$ meets $\geq c|\mathcal{P}_{pq}^{\ell}|/r^2$ paths from $\mathcal{V}_k$ for some $k \geq 2$, and so that e^2 does not lie on one of these paths. Continuing in this fashion, define $e^1, \dots, e^{r^2-1}$ and $W^1, \dots, W^{r^2-1}$, and let W^{r^2} be the final component of $W_1 - e^{r^2-1}$. Without loss of generality we may assume that each W^k meets $\geq c|\mathcal{P}_{pq}^{\ell}|/r^2$ paths from $\mathcal{V}_k$.

Let $m := \lfloor \sqrt{c}/r^2 \rfloor$. For every $k = 1, \dots, r^2$ let e_1^k be the first edge of W^k such that the initial component W_1^k of $W^k - e_1^k$ meets $\geq \lfloor \sqrt{c} \rfloor |\mathcal{P}_{pq}^{\ell}|$ paths from

$\mathcal{V}_k$, and so that e_1^k does not lie on one of these paths. Let e_2^k be the first edge of $W^k - W_1^k - e_1^k =: R^k$ such that the initial component W_2^k of $R^k - e_2^k$ meets $\geq \lfloor\sqrt{c}\rfloor|\mathcal{P}_{pq}^\ell|$ paths from $\mathcal{V}_k$, and so that e_2^k does not lie on one of these paths. Continuing in this fashion, define $e_1^k, \dots, e_{m-1}^k$ and $W_1^k, \dots, W_{m-1}^k$, and let W_m^k denote the final component of $W^k - e_m^k$. Thus each W_a^k meets $\geq \lfloor\sqrt{c}\rfloor|\mathcal{P}_{pq}^\ell|$ paths from $\mathcal{V}_k$, and $e_a^k \notin E(F^\ell)$. For $k < r^2$, let $e_m^k := e^k$. Menger's theorem and condition (iv) now imply that for each e_a^k there is a set S_a^k of $< |\mathcal{P}_{pq}^\ell|$ vertices separating Y_{pq}^ℓ from Y_{qp}^ℓ in the graph $(H_{pq}^\ell \cup F^\ell) - e_a^k$. Let S be the union of all these S_a^k. Then

$$|S| \leq (r^2m - 1)(|\mathcal{P}_{pq}^\ell| - 1) \leq \lfloor\sqrt{c}\rfloor(|\mathcal{P}_{pq}^\ell| - 1).$$

Hence each W_a^k meets at least one path $V_a^k \in \mathcal{V}_k$ that avoids S. Clearly, S_a^k must consist of exactly one vertex $v_a^k(P)$ on each path $P \in \mathcal{P}_{pq}^\ell \backslash \{W_1\}$. For all P and $1 \leq k < r^2$ let $e_0^{k+1} := e_m^k$ and $v_0^{k+1}(P) := v_m^k(P)$. Let $v_0^1(P)$ be the endvertex of P in Y_{pq}^ℓ and $v_m^{r^2}(P)$ that in Y_{qp}^ℓ. Note that V_a^k meets P neither in the initial component of $P - v_{a-1}^k(P)$ nor in the final component of $P - v_a^k(P)$ (here both the initial component of $P - v_0^1(P)$ and the final component of $P - v_m^{r^2}(P)$ are defined to be the empty set)—otherwise there would be a Y_{pq}^ℓ-Y_{qp}^ℓ path in the graph $(H_{pq}^\ell \cup F^\ell) - e_{a-1}^k$ or $(H_{pq}^\ell \cup F^\ell) - e_a^k$ avoiding S. This implies that $v_a^k(P)$ precedes $v_{a+1}^k(P)$ when P is directed from Y_{pq}^ℓ to Y_{qp}^ℓ. Thus for all $1 \leq a \leq m$, $1 < b \leq r$ and $1 \leq k \leq r^2$ the path V_a^k meets W_b exactly in the segment of W_b strictly between $v_{a-1}^k(W_b)$ and $v_a^k(W_b)$, and V_a^k meets W^1 exactly in the segment W_a^k. That means that the binary trees in the H_{ik} $(1 \leq k \leq r^2)$ are m-attached to an $r \times mr^2$ pseudogrid whose vertical paths are the V_a^k and whose horizontal paths are those obtained from $W_1, \dots, W_r$ by deleting their endvertices. Thus the horizontal paths are disjoint from all H_{ik}, and the vertical paths meet the H_{ik} only in their first vertices. Since

$$\begin{aligned}
\lceil \log_2 n \rceil^{2r-1} &= \lceil \log_2 (2^{r^9} r^{r^5+4r^4+2}) \rceil^{2r-1} \\
&= \lceil r^9 + (r^5 + 4r^4 + 2)\log_2 r \rceil^{2r-1} \\
&\leq (2r^9)^{2r-1} \\
&\leq 2^{r^5/4} \\
&\leq \sqrt{m},
\end{aligned}$$

we can apply Lemma 5.7 with n replaced by $\lceil \log_2 n \rceil$ to find the desired K_r minor in G.

Case 2. All paths in $\mathcal{P}$ avoid less than $|\mathcal{P}_{ij}^\ell|/c$ paths in $\mathcal{P}_{ij}^\ell$.

As in Case 1, we first apply Lemma 5.8 to find a set $\mathcal{V}$ of $\geq |\mathcal{P}|/2$ paths from $\mathcal{P}$ and a set $\mathcal{W}$ of r paths from $\mathcal{P}_{ij}^\ell$ such that every path in $\mathcal{V}$ meets every path in $\mathcal{W}$. Again, we then find a set $\mathcal{V}'$ of $\geq |\mathcal{V}|/r^r$ paths from $\mathcal{V}$ and a labelling of the paths from $\mathcal{W}$ by $W_1, \dots, W_r$ such that on its way from Y_{pq}^ℓ to Y_{qp}^ℓ every path from $\mathcal{V}'$ meets W_a before it meets W_{a+1} (for all $1 \leq a < r$). Recall that by condition (ii) each H_{pk} $(k = 1, \dots, s)$ contains endvertices of $\leq n/r^2sc^{\ell r^2 + \lfloor r^2/2 \rfloor - \ell}$ paths from $\mathcal{P}^{\ell}_{pq}$, and $\lfloor r^2/2 \rfloor - 2 \geq \binom{r}{2}$ since $r \geq 4$. Together with

$$\begin{aligned} |\mathcal{V}'|/2 &\geq |\mathcal{P}^{\ell}_{pq}|/4r^{r+2} = \frac{n}{4r^{r+4}c^{\ell r^2+\lfloor r^2/2 \rfloor}} \geq \frac{n}{c^{\ell r^2+\lfloor r^2/2 \rfloor+1}} \\ &\geq r^2 \frac{n}{r^2 s c^{\ell r^2+\lfloor r^2/2 \rfloor-\ell}} \end{aligned}$$

and

$$|\mathcal{V}'|/2 \geq (s-r^2)c\frac{n}{r^2 c^{(\ell+1)r^2}} = (s-r^2)c|\mathcal{P}^{\ell}_{ij}|,$$

this implies that there are r^2 of the H_{pk}, without loss of generality $H_{p1}, \ldots, H_{pr^2}$, such that each of them contains endvertices of $\geq c|\mathcal{P}^{\ell}_{ij}|$ paths from $\mathcal{V}'$. For all $k = 1, \ldots, r^2$, let $\mathcal{V}_k$ be the set of all paths begining at H_{pk}. Similarly as in Case 1, the paths in $\mathcal{W}$ together with $\lfloor \sqrt{c}/r^2 \rfloor$ paths from each $\mathcal{V}_k$ form a pseudogrid, and we can apply Lemma 5.7 to find a K_r minor in G. □

5.4 A tree attached to many non-outerplanar graphs

Let $\mathcal{H}_n$ be the class of all graphs H which can be obtained from the binary tree T of height n by identifying each leaf v of T with a vertex of a connected non-outerplanar graph $K(v)$ (where the $K(v)$ are disjoint from each other and from the rest of T). T is called the binary tree in H, the leaves of T are called *leaves of* H, and $K(v)$ is said to be the non-outerplanar graph *glued to* v.

Theorem 5.10 *Given an integer r, there exist integers d and n with the following property. Whenever a graph G contains an externally d-connected set X such that $G[X]$ has some graph in $\mathcal{H}_n$ as a minor, there is a K_r minor in G.*

Lemma 5.4 together with the fact that every non-outerplanar graph contains a subdivision of K_4 or $K_{2,3}$ implies that for $n \geq 2$ every graph in $\mathcal{H}_n$ contains H^4_k, $H^{2,3}_k$ or $H^{3,2}_k$ as a minor where $k := \lfloor (n-2)/\log_2 n \rfloor$. Thus in the statement of Theorem 5.10 one could have alternatively required that $G[X]$ contains either H^4_n, $H^{2,3}_n$ or $H^{3,2}_n$ as a minor.

Actually, the property that will be used in the proof of Theorem 5.10 is not that every graph $K(v)$ is non-outerplanar, but that each $K(v)$ has three distinct vertices $x, y, z \neq v$ such that any two vertices of x, y, z can be joined by a path P while the third can be joined to v by a path not meeting P. Indeed, since every non-outerplanar graph contains a subdivision of K_4 or $K_{2,3}$, such vertices x, y, z can be found.

Conversely, note that every graph K containing distinct vertices v, x, y, z satisfying the above property cannot be outerplanar. Indeed, suppose that K is outerplanar. Then adding a new vertex a to K and joining it to v, x, y, z yields a planar graph K'. Consider a drawing of K', and let aa_1, aa_2, aa_3, aa_4 be the edges of K' incident to a in clockwise order (thus $\{a_1, a_2, a_3, a_4\} = \{v, x, y, z\}$). Then $K' - a$ contains an a_1-a_3 path Q_1 and an a_2-a_4 path Q_2 such that Q_1 and Q_2 are disjoint, contradicting the planarity of K'.

For every leaf v of $H \in \mathcal{H}_n$ choose a set $B(v)$ consisting of three such vertices x, y, z of $K(v)$. For the proof of Theorem 5.10 we shall need the following lemmas.

Lemma 5.11 *Let n, r, s be positive integers such that $s^{1/4} \geq n^{2r-1+\lceil \log_2 \binom{r}{2} \rceil}$ and $n \geq 2$. Let G be a graph that contains a binary tree T of height n together with a path P such that T and P are disjoint and there is a set A of s leaves of T which are joined (by edges) injectively to the vertices of P. Suppose that T is s-attached in $G - P$ to an $r \times s$ pseudogrid G' so that A is precisely the set of attached leaves of T. Then G contains a K_r minor.*

Proof. First apply Lemma 5.6 to obtain a set A' of $\geq \sqrt{s}$ leaves in A such that their order in T corresponds to the order of their neighbours on P. Apply Lemma 5.6 once more to obtain a set A'' of $\geq s^{1/4}$ leaves in A' such that their order in T corresponds to the order of the vertical paths in G' which begin in A''. Lemma 5.4 now gives us a subdivision S of the binary tree of height $2r-1+\lceil \log_2 \binom{r}{2} \rceil$ in T such that all leaves of S are contained in A''. S contains $\binom{r}{2}$ disjoint subdivisions $S_1, \ldots, S_{\binom{r}{2}}$ of the binary tree of height $2r-1$ such that each of them is a branch of S. Note that $S_1, \ldots, S_{\binom{r}{2}}$ are tidily 2^{2r-1}-attached in $G-P$ to an $r \times 2^{2r-1}\binom{r}{2}$ subpseudogrid of G', and that there are $\binom{r}{2}$ disjoint segments of P, each containing all the neighbours of leaves of some S_i on P. Lemma 5.11 now follows from Lemma 5.5. $\square$

Lemma 5.12 *Let $\mathcal{M}$ and $\mathcal{H} = \{H_1, H_2, H_3\}$ be sets of disjoint directed paths such that $|\mathcal{M}| = 3$ and every path from $\mathcal{M}$ meets every path from $\mathcal{H}$. Then there are vertices x_1, x_2, x_3 on H_1, H_2, H_3 respectively, and a labelling of the paths in $\mathcal{M}$ as M_1, M_2, M_3, such that, for all $i = 1, 2, 3$, the vertex x_i lies on M_i and $M_i\mathring{x}_i$ does not meet any $H_j x_j$ $(j = 1, 2, 3)$.*

Proof. For every $i = 1, 2, 3$, let y_i^1 be the first vertex of H_i that lies on a path from $\mathcal{M}$. Given $k \geq 1$, assume inductively that for every $i = 1, 2, 3$ we have constructed a sequence $y_i^1, \ldots, y_i^k$ of vertices on H_i satisfying the following conditions.

(i) If $k > 1$, then $y_i^k \in y_i^{k-1}H_i$, and $y_i^k \neq y_i^{k-1}$ for at least one $i \in \{1, 2, 3\}$.

(ii) The vertex y_i^k lies on some path $M_i^k \in \mathcal{M}$.

(iii) For every $M \in \mathcal{M}$, the initial component of $M - \{y_1^k, y_2^k, y_3^k\}$ meets none of the paths $H_j y_j^k$ $(j = 1, 2, 3)$.

If the paths M_1^k, M_2^k, M_3^k are all distinct, then $x_i := y_i^k$ and $M_i := M_i^k$ satisfy the assertion of the lemma. We show that if M_1^k, M_2^k, M_3^k are not distinct then there are vertices y_i^{k+1} extending our three sequences $y_i^1, \ldots, y_i^k$; by (i), this can happen only finitely often.

If M_1^k, M_2^k, M_3^k are not distinct, then there exists a path $M \in \mathcal{M}$ containing more than one of the vertices y_1^k, y_2^k, y_3^k (let y_i^k be the last of these on M), as well as a path $M' \in \mathcal{M}$ avoiding $\{y_1^k, y_2^k, y_3^k\}$. By (iii), M' avoids $H_i y_i^k$ and hence meets H_i in $\mathring{y}_i^k H_i$. So $\mathring{y}_i^k H_i$ has a first vertex in $\bigcup \mathcal{M}$; we choose this vertex as y_i^{k+1} and put $y_j^{k+1} := y_j^k$ for $j \neq i$. Then conditions (i) and (ii) hold for $k+1$. Condition (iii) for $k+1$ holds for our M because $M\mathring{y}_i^k$ contains another $y_j^k = y_j^{k+1}$. Condition (iii) for the other two paths in $\mathcal{M}$ is again clear: as they do not contain y_i^k, their initial components in (iii) did not get longer when y_i^k was replaced by y_i^{k+1}, so they satisfy (iii) for $k+1$ because they did for k. $\square$

Lemma 5.13 *Let n, r, s be positive integers such that $\sqrt{s} \geq n^{2r-1+\lceil \log_2 \binom{r}{2} \rceil}$ and $n \geq 2$, and let G be a graph containing a graph $H \in \mathcal{H}_n$. Let $A = \{v_1, \ldots, v_s\}$ be a set of leaves of H, and let $B := \bigcup_{v \in A} B(v)$. Suppose that there is a set $\mathcal{V}$ of $3s$ disjoint directed paths in G starting in B and having no other vertices in H, and a set $\mathcal{W} = \{W_0, \ldots, W_r\}$ of disjoint directed paths in $G - H$ such that every path in $\mathcal{V}$ meets each W_i, and it does so before it meets W_{i+1}. Suppose furthermore that each path $W_i \in \mathcal{W}$ consists of s consecutive (vertex disjoint) segments such that, for all $j = 1, \ldots, s$, every path from $\mathcal{V}$ that starts in $B(v_j)$ meets W_i exactly in its jth segment. Then G contains a K_r minor.*

Proof. First apply Lemma 5.6 to obtain a set A' of $\geq \sqrt{s}$ leaves in A such that their order in the binary tree T of H corresponds to the order of the paths from $\mathcal{V}$ that begin in $\bigcup_{v \in A'} B(v)$, i.e. if $v, w \in A'$ and v is the successor of w in the ordering σ_T restricted to A', then no path from $\mathcal{V}$ that starts in $B(u)$ for some $v, w \neq u \in A'$ lies between paths from $\mathcal{V}$ starting in $B(v)$ and $B(w)$. Moreover, reversing the orientation of the paths from $\mathcal{W}$ if necessary, we may assume that for each W_i the segment of W_i belonging to the paths from $\mathcal{V}$ starting in $B(w)$ precedes that belonging to the paths from $\mathcal{V}$ starting in $B(v)$. By Lemma 5.4, there is a subdivision S of the binary tree of height $2r - 1 + \lceil \log_2 \binom{r}{2} \rceil$ in T such that all leaves of S are contained in A'.

Let $S_1, \ldots, S_{\binom{r}{2}}$ be disjoint subdivisions of the binary tree of height $2r - 1$ in S such that each of them is a branch of S. Each S_ℓ has a set $\mathcal{P}_\ell$ of r nested paths. Let G' be the graph consisting of the paths from $\mathcal{W}$ together with all paths from $\mathcal{V}$ starting in non-outerplanar graphs glued to leaves of the S_ℓ.

We now construct the branch sets for our K_r minor. Each of these branch sets will contain a path Q_i running alternately through an S_ℓ and G' in a similar way as in the proof of Lemma 5.5. In particular, each Q_i will contain exactly one path from each $\mathcal{P}_\ell$. For the edges of the K_r we need disjoint paths, one between any two of the Q_i. Each of these paths we will find in one of the trees $S_1, \ldots, S_{\binom{r}{2}}$. Indeed, in each S_ℓ we can join two neighbouring Q_i (i.e. two Q_i containing paths from $\mathcal{P}_\ell$ lying next to each other), and we will show that we can also use a non-outerplanar graph glued to a leaf of S_ℓ to 'switch' two neighbouring Q_i (Fig. 5.3). Together this will imply that for every edge of the K_r we can find a path connecting the corresponding Q_i.

To make this more precise, suppose that we have partially constructed such paths $Q_1, \ldots, Q_r$, which have their current endvertices on different paths from $\mathcal{W} \backslash \{W_0\}$, and that next we want to let them run through S_ℓ, and that furthermore we want to switch Q_i and Q_j, where Q_j currently runs along W_k and Q_i along W_{k+1} for some $k > 0$. Let A_ℓ be the set of all those leaves of S_ℓ in which there begins a path from $\mathcal{P}_\ell$. Let v be the $(k+1)$th leaf in A_ℓ (in the ordering σ_T restricted to A_ℓ), and let w be the predecessor of v in A_ℓ. We will use the non-outerplanar graph $K(v)$ to switch Q_i and Q_j.

Let V_x, V_y, V_z denote the paths from $\mathcal{V}$ starting in $B(v) =: \{x, y, z\}$. Apply Lemma 5.12 to $\mathcal{H} := \{W_0, W_k, W_{k+1}\}$ and the set $\mathcal{M}$ consisting of the subpaths

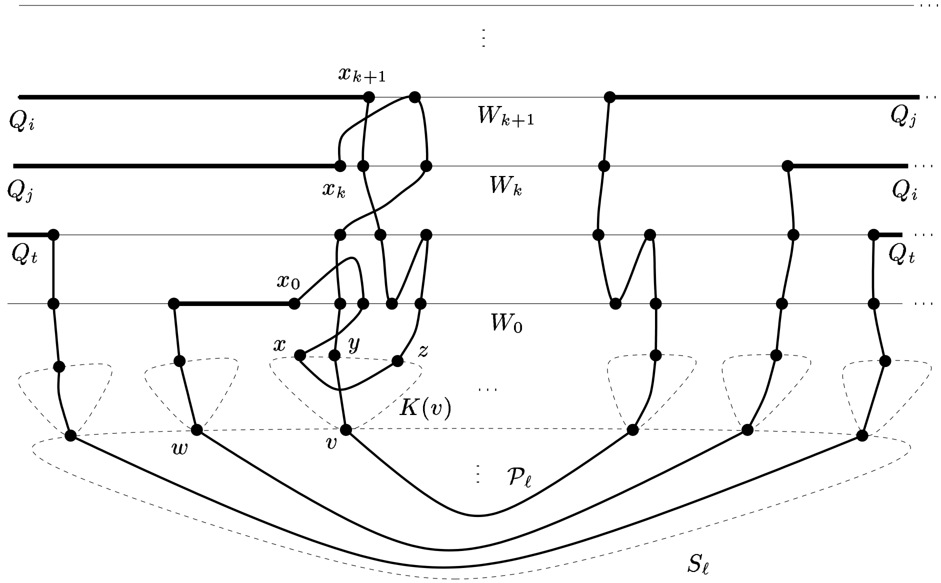

Figure 5.3: Switching Q_i and Q_j

of V_x, V_y, V_z between their endvertices in $B(v)$ and their first vertices on W_{k+1} to obtain vertices x_0, x_k, x_{k+1} and a labelling M_0, M_k, M_{k+1} of the paths from $\mathcal{M}$. Then for all $i = 0, k, k+1$, traversing M_i from $B(v)$ to x_i, and then moving backwards along W_i gives disjoint paths. Thus we may extend Q_j by traversing W_k as far as x_k, and then moving along M_k to $B(v)$ and further through $K(v)$ to v, and then along the path P_v from $\mathcal{P}_\ell$ that begins in v and up to W_{k+1} along a path from $\mathcal{V}$ starting in the non-outerplanar graph glued to the endvertex of P_v. Extend Q_i by traversing W_{k+1} as far as x_{k+1}, then moving along M_{k+1} to $B(v)$ and through $K(v)$ to the endvertex of M_0, then along M_0 to x_0, then backwards along W_0 and down through $K(w)$, and along the path P_w from $\mathcal{P}_\ell$ starting in w, and then up to W_k along a path from $\mathcal{V}$ starting in the non-outerplanar graph glued to the endvertex of P_w. From the definition of $B(v)$, it follows that we may choose the subpaths of Q_i and Q_j running through $K(v)$ so that Q_i and Q_j remain disjoint.

Moreover, note that $\binom{r}{2}$ switchings suffice to ensure that for any two Q_i there is a tree S_ℓ in which they lie next to each other (and thus can be joined by a path). Indeed, first switch the lowest Q_i with all higher ones, then the one which is now the lowest with all but the highest. Continuing in this fashion, we need $(r-1)+(r-2)+\cdots+1 = \binom{r}{2}$ switchings. Thus the Q_i can be constructed to form the branch sets of a (subdivided) K_r minor in G. $\square$

Lemma 5.14 *Let T be a tree with $\Delta(T) \leq 3$. Suppose that $B_1, \ldots, B_k$ are disjoint sets of leaves of T such that for all $1 \leq i < k$ the leaves of T in B_i either all precede or all succeed those in B_{i+1} in the ordering σ_T. Then there are $\lceil k/3 \rceil$ disjoint subtrees of T such that each of them contains all vertices of some B_i.*

Proof. Induction on k. We may assume that $k \geq 4$. Let e be an edge of T such that the branch T' above e contains all vertices of some B_i, and such that T' is as small as possible. As $\Delta(T) \leq 3$, the endvertex x of e in T' has degree at most two in T'. It is easy to see that the minimality of T' now implies that T' meets at most three of the B_i. T' will be one of the desired subtrees of T, and applying the induction hypothesis on $T - T' - e$ and the B_i's not meeting T' gives the remaining $\lceil (k-3)/3 \rceil = \lceil k/3 \rceil - 1$ subtrees. □

Lemma 5.15 *Let n, r, s be positive integers such that $s^{1/4} \geq n^{2r-1+\lceil \log_2 3\binom{r}{2} \rceil}$ and $n \geq 2$. Let G be a graph containing two disjoint copies T_1 and T_2 of the binary tree of height n, and let A be a set of s leaves of T_1. Suppose that there is a set $\mathcal{P}$ of s disjoint paths joining the vertices in A to leaves of T_2 and meeting $T_1 \cup T_2$ only in their endvertices. Let G_0 be the graph obtained from G by deleting T_2 and all inner vertices of paths from $\mathcal{P}$. Furthermore, suppose that T_1 is s-attached in G_0 to an $r \times s$ pseudogrid G' such that A is precisely the set of attached leaves of T_1. Then G contains a K_r minor.*

Proof. By Lemma 5.6, there exists a set A' of $\geq s^{1/4}$ leaves in A such that their order in σ_{T_1} corresponds (or is opposite) to both the order of the leaves of T_2 joined to A' by paths from $\mathcal{P}$ and the order of the vertical paths in G' that begin in A'. Let $k := 3\binom{r}{2}$. Lemma 5.4 implies that T_1 contains a subdivision S of the binary tree of height $2r - 1 + \lceil \log_2 k \rceil$ such that all leaves of S are contained in A'. Let $S_1, \ldots, S_k$ be disjoint subdivisions of the binary tree of height $2r-1$ in S such that each S_i is a branch of S. Then $S_1, \ldots, S_k$ are tidily 2^{2r-1}-attached to an $r \times k2^{2r-1}$ subpseudogrid of G'. For all $i = 1, \ldots k$, let B_i be the set of all leaves of T_2 which are endvertices of those paths from $\mathcal{P}$ that start in a leaf of S_i. From the choice of A', it follows that for all $1 \leq i < k$, the leaves in B_i either all precede or all succeed those in B_{i+1}. Lemma 5.14 now implies that there are $\lceil k/3 \rceil = \binom{r}{2}$ disjoint subtrees of T_2 such that each of them contains all vertices of some B_i. Lemma 5.15 thus follows from Lemma 5.5. □

Lemma 5.16 *Let W be a directed path. Let $\ell_1, \ldots, \ell_s$ and $r_1, \ldots, r_s$ be vertices on W such that $\ell_i \in Wr_i$ for all $i = 1, \ldots, s$. Let $S_1, \ldots, S_s$ be non empty disjoint segments of W such that S_i precedes S_{i+1} for all $1 \leq i < s$. Let $t := \lfloor (s/4)^{1/3} \rfloor$. Then there exists $I \subseteq \{1, \ldots, s\}$ with $|I| \geq t$ such that one of the following conditions holds.*

(a) *Either $S_{\max I} \subseteq W\mathring{r}_i$ for all $i \in I$ or $S_{\min I} \subseteq \mathring{\ell}_i W$ for all $i \in I$.*

(b) *For all $i \in I$, there is a segment A_i of W such that each A_i contains ℓ_i, r_i and S_i, and $A_i \cap A_j = \emptyset$ for all $i, j \in I$ with $i \neq j$.*

Proof. We may assume that $t \geq 2$. Moreover, let us first assume that for a set I_1 of $\geq s/2$ elements $i \in \{1, \ldots, s\}$ either $r_i \in S_i$ or $S_i \subseteq W\mathring{r}_i$. If for $\geq t$ elements $j \in I_1$ the segment S_j precedes the first r_i with $i \in I_1$ on W, then the subset of I_1 consisting of all these j satisfies condition (a). Thus, denoting the first of the r_i with $i \in I_1$ on W by r_{i_1}, we may assume that $S_j \subseteq \mathring{r}_{i_1} W$ for a set I_1' of $\geq |I_1| - t$ elements $j \in I_1$. Note that $i_1 \notin I_1'$. As before, we are done if $\geq t$ segments S_j with $j \in I_1'$ precede the first r_i with $i \in I_1'$ on W, r_{i_2} say. Thus we may assume that $S_j \subseteq \mathring{r}_{i_2} W$ for a set I_1'' of $\geq |I_1'| - t$ elements $j \in I_1'$. Continuing in this fashion, we may assume that there is a set $I_2 \subseteq I_1$ with $|I_2| \geq |I_1|/t$ such that $S_j \subseteq \mathring{r}_i W$ for all $i < j$ in I_2. (Indeed, let I_2 be the set consisting of $i_1, i_2, \ldots$.)

Let S_i' be the smallest segment of W containing S_i and r_i. Then for all $i \in I_2$ the segments S_i' are pairwise disjoint. If for a set I_3 of $\geq |I_2|/2 \geq t$ elements $i \in I_2$ either $\ell_i \in S_i$ or $S_i \subseteq W\mathring{\ell}_i$, then I_3 satisfies condition (b) (since $\ell_i \in W r_i$ we may take A_i to be S_i'). Thus we may assume that there is a set I_3 of $\geq |I_2|/2$ elements of I_2 such that $S_i \subseteq \mathring{\ell}_i W$ for all $i \in I_3$. Considering the vertices ℓ_i and the segments S_i' for all $i \in I_3$ and arguing similarly as before, we may assume that there is a set I of $\geq |I_3|/t \geq t$ elements of I_3 such that $S_i' \subseteq W\mathring{\ell}_j$ for all $i < j$ in I. Thus I satisfies condition (b). The case that $S_i \subseteq \mathring{r}_i W$ for $\geq s/2$ elements $i \in \{1, \ldots, s\}$ is similar. $\square$

Proof of Theorem 5.10. It suffices to show the following assertion.

Let $c := r^{16r^7 3^{2r+2}}$ and $n := c^{2\binom{r}{2}}$. Suppose G contains a graph H consisting of r disjoint graphs $H_1, \ldots, H_r \in \mathcal{H}_{\lceil \log_2 n \rceil}$ as an externally $3n$-connected subgraph. Then G contains a K_r minor. $(**)$

We may assume that $r \geq 4$. The first part of the proof of $(**)$ is similar to (but much easier than) that of $(*)$, and we will only sketch it. For every $i = 1, \ldots, r$, choose n leaves of H_i and let Y_i denote the union of the sets $B(v)$ over all chosen leaves v. Since H is externally $3n$-connected in G, for all pairs $1 \leq i < j \leq r$ there is a set $\mathcal{P}_{ij}$ of $|Y_i| = 3n$ disjoint Y_i-Y_j paths in G which have no inner vertices in H.

As in the proof of $(*)$, we will try to find single paths $P_{ij} \in \mathcal{P}_{ij}$ that are disjoint for different pairs ij, and thus link up the graphs H_i to form a K_r minor in G. If that is not possible, there will be two sets $\mathcal{P}_{pq}$ and $\mathcal{P}_{ij}$, such that many paths of $\mathcal{P}_{pq}$ together with many paths of $\mathcal{P}_{ij}$ form a pseudogrid, which we shall then use to find a K_r minor within the graph consisting of H_p, H_q and the paths in $\mathcal{P}_{pq} \cup \mathcal{P}_{ij}$. Let $\sigma : \{ij \mid 1 \leq i < j \leq r\} \to \{0, 1, \ldots, \binom{r}{2} - 1\}$ be any bijection. Let $\ell^* \leq \binom{r}{2}$ be maximal such that for all $0 \leq \ell < \ell^*$ and $1 \leq i < j \leq r$ there exist sets $\mathcal{P}_{ij}^\ell$ satisfying the following conditions.

(i) $\mathcal{P}_{ij}^\ell$ is a non-empty set of disjoint Y_i-Y_j paths having no inner vertices in H.

As soon as $\mathcal{P}_{ij}^\ell$ is defined, let H_{ij}^ℓ be the graph consisting of all paths in $\mathcal{P}_{ij}^\ell$. Furthermore, let Y_{ij}^ℓ be the set of all endvertices of paths from $\mathcal{P}_{ij}^\ell$ in Y_i and Y_{ji}^ℓ the set of those in Y_j.

(ii) If $\sigma(ij) < \ell$, then $\mathcal{P}_{ij}^\ell$ has exactly one element P_{ij}^ℓ, and P_{ij}^ℓ is disjoint from any path belonging to a set $\mathcal{P}_{ab}^\ell$ with $ab \neq ij$.

(iii) If $\sigma(ij) \geq \ell$ and v is a leaf of H_i, then either each of the three vertices in $B(v)$ belongs to Y^ℓ_{ij} or non of them does. In the first case we will say that the three paths of $\mathcal{P}^\ell_{ij}$ that start from the vertices in $B(v)$ form a *bundle*. Thus (iii) says that $\mathcal{P}^\ell_{ij}$ consists only of bundles.

(iv) If $\sigma(ij) = \ell$, then $|\mathcal{P}^\ell_{ij}| = 3n/c^{2\ell}$.

(v) If $\sigma(ij) > \ell$, then $|\mathcal{P}^\ell_{ij}| = 3n/c^{2\ell+1}$.

(vi) If $\ell = \sigma(pq) < \sigma(ij)$, then for every edge $e \in E(H^\ell_{ij}) \backslash E(H^\ell_{pq})$ there are no $|\mathcal{P}^\ell_{ij}|$ disjoint Y^ℓ_{ij}-Y^ℓ_{ji} paths in the graph $(H^\ell_{ij} \cup H^\ell_{pq}) - e$.

If $\ell^* = \binom{r}{2}$, then by (i) and (ii) we have a (subdivided) K_r minor with branch sets H_i. Thus suppose that $\ell^* < \binom{r}{2}$, and note that as in the previous section $\ell^* > 0$. Let $\ell := \ell^* - 1$ and $pq := \sigma^{-1}(\ell)$. Similarly as in the proof of $(*)$, for every path $P \in \mathcal{P}^\ell_{pq}$ there exists a pair ij with $\sigma(ij) > \ell$ such that P avoids $\leq |\mathcal{P}^\ell_{ij}|/3c$ bundles from $\mathcal{P}^\ell_{ij}$ (where P *avoids a bundle* if it avoids every path in it). Thus there is a set $\mathcal{P}$ of $\geq |\mathcal{P}^\ell_{pq}|/3^2\binom{r}{2}$ paths from $\mathcal{P}^\ell_{pq}$ for which this pair ij can be chosen to be the same, and such that any two paths from $\mathcal{P}$ belong to different bundles of $\mathcal{P}^\ell_{pq}$ and end in different non-outerplanar graphs glued to leaves of H_q. Again, we now use Lemma 5.8 to find a set $\mathcal{V}'$ of $\geq |\mathcal{P}|/2$ paths from $\mathcal{P}$ and a set $\mathcal{W}'$ of $t := (2r)^4 + 2r$ bundles from $\mathcal{P}^\ell_{ij}$ such that no path in $\mathcal{V}'$ avoids a bundle in $\mathcal{W}'$. Then there is a set $\mathcal{V}''$ of $\geq |\mathcal{V}'|/3^t$ paths from $\mathcal{V}'$ and a set $\mathcal{W}''$ of t paths, one from each bundle in $\mathcal{W}'$, such that every path in $\mathcal{V}''$ meets every path in $\mathcal{W}''$. We now find a set $\mathcal{V}'''$ of $\geq |V''|/t^t \geq |\mathcal{P}^\ell_{ij}|c/3^{t+2}t^tr^2$ paths from $\mathcal{V}''$ and a labelling of the paths in $\mathcal{W}''$ by $W''_1, \ldots, W''_t$ such that on its way from H_p to H_q every path in $\mathcal{V}'''$ meets W''_a before it meets W''_{a+1} (for all $1 \leq a < t$). Similarly as in the proof of $(*)$, condition (vi) now yields a set $\mathcal{V}^*$ of $\geq \lfloor \sqrt{c}/3^{(t+2)/2}rt^{t/2} \rfloor$ paths from $\mathcal{V}'''$ which form a pseudogrid together with the paths in $\mathcal{W}''$. To make the horizontal paths of the pseudogrid disjoint from $H_p \cup H_q$, we direct each path in $\mathcal{W}''$ from H_i to H_j, and let $\mathcal{W}$ be the set of all (directed) paths obtained from paths in $\mathcal{W}''$ by deleting their endvertices. Let

$$s := \left\lceil \frac{\sqrt{c}}{t^{2t}} \right\rceil \leq \left\lfloor \frac{\sqrt{c}}{3^{(t+2)/2}rt^{t/2}} \right\rfloor - 2.$$

Discarding also the leftmost and the rightmost path from $\mathcal{V}^*$ in the pseudogrid (as well as any $|\mathcal{V}^*| - s - 2$ other paths), we have thus found sets $\mathcal{V} = \{V_1, \ldots, V_s\} \subseteq \mathcal{V}^* \subseteq \mathcal{P}^\ell_{pq}$ and $\mathcal{W} = \{W_1, \ldots, W_t\} \subseteq \mathcal{P}^\ell_{ij}$ satisfying the following conditions.

- No two paths from $\mathcal{V}$ belong to the same bundle of $\mathcal{P}^\ell_{pq}$. Furthermore, no two paths of $\mathcal{V}$ end in the same non-outerplanar graph glued to a leaf of H_q.
- Every path in $\mathcal{W}$ is disjoint from $H_p \cup H_q$.
- The paths in $\mathcal{V}$ and $\mathcal{W}$ together form an $s \times t$ pseudogrid; where the paths in $\mathcal{V}$ are its vertical and those in $\mathcal{W}$ its horizontal paths, every path

from $\mathcal{V}$ on its way from H_p to H_q meets W_a before it meets W_{a+1} (for all $1 \leq a < t$), and for every $i = 1, \ldots, t$ the segment of W_i belonging to V_a precedes that belonging to V_{a+1} (for all $1 \leq a < s$).

For all $a = 1, \ldots, s$, denote the two paths from $\mathcal{P}^\ell_{pq}$ that are in the same bundle as V_a by V'_a and V''_a, where we may assume that V''_a meets as least as many paths from $\mathcal{W}$ as V'_a.

Case 1. There are at least $s/2$ of the V'_a such that each of them avoids at least r paths in $\mathcal{W}$.

Then there is a set $I \subseteq \{1, \ldots, s\}$ with $|I| \geq s/2t^r$ such that each V'_a with $a \in I$ avoids the same r paths in $\mathcal{W}$, $W_{b_1}, \ldots, W_{b_r}$ say (where $b_1 < \cdots < b_r$). For all $a \in I$, let V^*_a be the subpath of V_a between its endvertex in H_p and its first vertex on W_{b_r}. Note that $\lceil \log_2 3\binom{r}{2} \rceil \leq r + 1$ since $r \geq 4$, and thus

$$\lceil \log_2 n \rceil^{2r-1+\lceil \log_2 3\binom{r}{2} \rceil} \;\leq\; \lceil r^2 \log_2 c \rceil^{3r} \;\leq\; \lceil 16 r^9 3^{2r+2} \log_2 r \rceil^{3r} \;\leq\; r^{r^7}. \quad (5.1)$$

Together with $t + r \leq r^7$ this implies

$$|I|^{1/4} \;\geq\; \left(\frac{\sqrt{c}}{2t^{2t+r}} \right)^{1/4} \;\geq\; \left(\frac{\sqrt{c}}{r^{14r^7}} \right)^{1/4} \;\geq\; r^{r^7} \;\stackrel{(5.1)}{\geq}\; \lceil \log_2 n \rceil^{2r-1+\lceil \log_2 3\binom{r}{2} \rceil}.$$

Thus we may apply Lemma 5.15 (with n replaced by $\lceil \log_2 n \rceil$) to the minor of G obtained by contracting every $K(v)$ to find a K_r minor in G. (The binary tree in H_p plays the role of T_1 in Lemma 5.15, the binary tree in H_q that of T_2, the set $\{V'_a \mid a \in I\}$ that of $\mathcal{P}$, and the pseudogrid formed by the V^*_a (for all $a \in I$) and $W_{b_1}, \ldots, W_{b_r}$ that of G'.)

Case 2. There are at least $s/2$ of the V'_a such that each of them meets at least $t - r$ paths in $\mathcal{W}$.

Note that if V'_a meets $\geq t - r$ paths in $\mathcal{W}$, then so does V''_a. Thus there is a set $I \subseteq \{1, \ldots, s\}$ with $|I| \geq s/2t^{2r}$ such that all V'_a and V''_a with $a \in I$ meet the same $t - 2r = (2r)^4$ paths of $\mathcal{W}$. Noting that there are $\leq t^t$ permutations of a $(t - 2r)$-element set and using Lemma 5.6 twice, we can find a set $I' \subseteq I$ with $|I'| \geq s/2t^{2t+2r}$ and paths $W_{b_1}, \ldots, W_{b_{2r}}$ with $b_1 < \cdots < b_{2r}$ such that

- either, on its way from H_p to H_q, every V'_a with $a \in I'$ meets W_{b_k} before it meets $W_{b_{k+1}}$ or every V'_a with $a \in I'$ meets $W_{b_{k+1}}$ before it meets W_{b_k} (for all $1 \leq k < 2r$), and
- the analogous condition holds for all V''_a with $a \in I'$.

Case 2.1. Every V'_a with $a \in I'$ meets $W_{b_{k+1}}$ before it meets W_{b_k} (for all $1 \leq k < 2r$).

For all $a \in I'$, let V^*_a be the subpath of V_a between its endvertex in H_p and its first vertex on W_{b_r}. Using that $t + r \leq r^7$ since $r \geq 4$, we have

$$|I'|^{1/4} \;\geq\; \left(\frac{\sqrt{c}}{2t^{4t+2r}} \right)^{1/4} \;\geq\; \left(\frac{\sqrt{c}}{r^{28r^7}} \right)^{1/4} \;\geq\; r^{r^7} \;\stackrel{(5.1)}{\geq}\; \lceil \log_2 n \rceil^{2r-1+\lceil \log_2 \binom{r}{2} \rceil}.$$

Thus we may apply Lemma 5.11 to (a minor of) G to find a K_r minor in G. (The binary tree in H_p plays the role of T in Lemma 5.11, $W_{b_{2r}}$ that of P, and the pseudogrid formed by all the V_a^* with $a \in I'$ and $W_{b_1}, \ldots, W_{b_r}$ that of G'.)

Case 2.2. Every V_a'' with $a \in I'$ meets $W_{b_{k+1}}$ before it meets W_{b_k} (for all $1 \leq k < 2r$).

The proof of this case is the same as that of Case 2.1.

Case 2.3. Neither Case 2.1 nor Case 2.2 hold.

For all $a \in I'$, let ℓ_a' (respectively ℓ_a'') be the first vertex of W_{b_1} that lies on the subpath of V_a' (respectively V_a'') between its endvertex in H_p and its first vertex on $W_{b_{r+1}}$. Similarly define r_a' and r_a'' to be the respective last vertices. Let S_a be the segment of W_{b_1} belonging to V_a in the pseudogrid formed by the paths in $\mathcal{V}$ and $\mathcal{W}$ (i.e. S_a is the segment between the first and the last vertex of W_{b_1} on V_a). Let $J \subseteq I'$ be obtained by applying Lemma 5.16 to ℓ_a', r_a' and S_a (for all $a \in I'$). Thus $|J| \geq \lfloor (|I'|/4)^{1/3} \rfloor \geq (|I'|/32)^{1/3}$.

Suppose first that J satisfies condition (a) of Lemma 5.16, say $S_{\max J} \subseteq W \mathring{r}_a'$ for all $a \in J$. Then the binary tree T in H_p is $|J|$-attached to an $r \times |J|$ pseudogrid whose vertical paths are the V_a with $a \in J$ together with (for each of these V_a) a path in the non-outerplanar graph glued to a leaf of H_p which joins this leaf to the endvertex of V_a in H_p, and whose horizontal paths are $W_{b_{r+1}}, \ldots, W_{b_{2r}}$. Note that

$$\begin{aligned}\left(\frac{|I'|^{1/3^{2r+2}}}{32}\right)^{1/4} &\geq \left(\frac{1}{32}\left(\frac{\sqrt{c}}{2t^{4t+2r}}\right)^{1/3^{2r+2}}\right)^{1/4} \\ &\geq \frac{c^{1/(8\cdot 3^{2r+2})}}{32^{1/4}(r^{28r^7})^{1/(4\cdot 3^{2r+2})}} \geq r^{r^7} \overset{(5.1)}{\geq} \lceil \log_2 n \rceil^{2r-1+\lceil \log_2 \binom{r}{2} \rceil}.\end{aligned} \tag{5.2}$$

Hence in particular

$$|J|^{1/4} \geq \left(\frac{|I'|}{32}\right)^{1/3 \cdot 1/4} \overset{(5.2)}{\geq} \lceil \log_2 n \rceil^{2r-1+\lceil \log_2 \binom{r}{2} \rceil},$$

and we may apply Lemma 5.11 to a minor of G to find a K_r minor in G. (The role of P in Lemma 5.11 is played by the subpath of W_{b_1} that starts in the first r_a' on W_{b_1} with $a \in J$.)

Hence we may assume that J satisfies condition (b) of Lemma 5.16. For all $a \in J$, let A_a be as in condition (b). Now let $J' \subseteq J$ be obtained by applying Lemma 5.16 to ℓ_a'', r_a'' and A_a (for all $a \in J$). Thus $|J'| \geq \lfloor (|J|/4)^{1/3} \rfloor \geq |I'|^{1/3^2}/32$. As before, since $|J'|^{1/4} \geq \lceil \log_2 n \rceil^{2r-1+\lceil \log_2 \binom{r}{2} \rceil}$ by (5.2), we may assume that J' satisfies condition (b). Applying the same argument to every W_{b_k} with $k \leq r+1$, we may assume that there exists $J'' \subseteq J'$ such that

$$|J''| \geq \frac{|I'|^{1/3^{2r+2}}}{32^{(3^0+3^1+\cdots+3^{2r+1})/3^{2r+2}}} \geq \frac{|I'|^{1/3^{2r+2}}}{32},$$

and that to every $a \in J''$ there belongs a segment B_{ak} on every W_{b_k} $(k \le r+1)$ that contains not only every vertex of V_a on W_{b_k} but also all vertices of W_{b_k} that lie on the subpaths of V_a' and V_a'' between their endvertices in H_p and their first vertices on $W_{b_{r+1}}$, and such that the B_{ak} are disjoint for different $a \in J''$. (Indeed, as $|J''|^{1/4} \ge \lceil \log_2 n \rceil^{2r-1+\lceil \log_2 \binom{r}{2} \rceil}$ by (5.2), we may assume that each application of Lemma 5.16 yields a subset of J' satisfying condition (b).)

Since $|J''|^{1/2} \ge \lceil \log_2 n \rceil^{2r-1+\lceil \log_2 \binom{r}{2} \rceil}$ by (5.2), we may apply Lemma 5.13 to find the desired K_r minor in G. $\square$

5.5 Proof of Theorem 5.2

In the proof of Theorem 5.2 we use the notion of a k-embedding of a given graph in the plane with one vortex. (It is a special case of the definition of a near-embedding of a graph in a surface introduced by Robertson and Seymour.) Let S be the surface obtained from the sphere by removing the interior of a closed disc. So the boundary of S consists of one component D, which is homeomorphic to the unit circle. We say that a graph G is *k-embeddable in the plane with one vortex* if G has a set X of at most k vertices such that $G - X$ can be written as $G^0 \cup G^1$ and the following conditions hold:

- there is an embedding of G^0 in S with $V(G^0) \cap D = V(G^0) \cap V(G^1) =: \{v_1, \dots, v_\ell\}$,
- G^1 has a path-decomposition $(p_1 \dots p_\ell, (X_{p_i})_{i=1}^{\ell})$ of width at most k and such that $v_i \in X_{p_i}$ for all $i = 1, \dots, \ell$, and the points $v_1, \dots, v_\ell$ occur on D in this order (for one of the two orientations of D).

The idea for the proof of Theorem 5.2 is to show that if $\mathcal{H}$ and $\mathcal{H}^*$ are inequivalent, then there exists $k \in \mathbb{N}$ such that the vertices of each graph $H \in \mathcal{H}$ can be joined bijectively to the bottom row of a grid to obtain a graph G_H in which H is externally $|H|$-connected but which is k-embeddable in the plane with one vortex. Lemma 2.3 from [15] then implies that no graph G_H $(H \in \mathcal{H})$ contains a large complete minor. Hence $\mathcal{H}$ does not force large complete minors, contradicting our assumption.

Let us first prove a simple lemma, which implies that if T, T' are trees and T is a minor of T', then T does not lie 'upside down' in T'.

Lemma 5.17 *Suppose that T, T' are trees and T is a minor of T'. Let $X \subseteq V(T')$ be the branch set corresponding to a leaf x of T. Then T' contains a path which joins X to a leaf of T' but does not meet another branch set.*

Proof. Suppose that $|T| \ge 2$, let x be a leaf of T and let y be the neighbour of x in T. Let e be the unique edge of T' corresponding to xy (i.e. e is the edge joining the branch set X of x to that of y). Write C for the component of $T - e$ which contains X. Then the (connected) subgraph of T' which contracts to $T - x$ is disjoint from C. Thus joining x to any leaf of T contained in C yields a path as desired. $\square$

Proof of Theorem 5.2. We consider only the case where $\mathcal{H}^* = \mathcal{H}'$, the other three cases are similar. Let $H_n^\sharp$ denote the graph obtained from the binary tree T of height n by adding a new vertex x and joining it to all leaves of T. Thus H'_n is the disjoint union of n copies of $H_n^\sharp$. The vertices in H'_n of the form x are called the *new vertices* of H'_n.

Since $\mathcal{H} \preccurlyeq \mathcal{H}'$, we find for every $H \in \mathcal{H}$ an $n \in \mathbb{N}$ such that H is a minor of H'_n. Let X_H be the set of all vertices of H whose branch sets contain a new vertex of H'_n. Thus $H - X_H$ is a forest, every component of which is joined to at most one vertex in X_H by edges of H. For every $x \in X_H$, let F_x be the forest consisting of all paths in $H - X_H$ between neighbours (in H) of x. So we have that

(i) the F_x are disjoint for different $x \in X_H$,

(ii) for each $x \in X_H$ the neighbours of x in H all lie in F_x,

(iii) for each $x \in X_H$ every component C of F_x is a subtree of some component of $H - X_H$, and this component meets F_x precisely in C and does not meet any $F_{x'}$ with $x' \neq x$.

Case 1. For every $k \in \mathbb{N}$ there is an $H \in \mathcal{H}$ such that at least k of the forests F_x $(x \in X_H)$ have path-width at least k.

We will show that $\mathcal{H}' \preccurlyeq \mathcal{H}$. Our assumption implies that for every $n \in \mathbb{N}$ there exists a graph $H \in \mathcal{H}$ and distinct $x_1, \dots, x_n \in X_H$ such that $\mathrm{pw}(F_{x_i}) \geq 2^{n+1}$ for every $i = 1, \dots, n$. Thus each F_{x_i} contains the binary tree of height n as a minor (see e.g. [6, Thm. 12.4.5]). Lemma 5.17 now says that every branch set corresponding to a leaf of this binary tree can be reached from a leaf of F_{x_i} by a path in F_{x_i} not meeting another branch set. Clearly, these paths are disjoint for different leaves of the binary tree. As in H every leaf of F_{x_i} is a neighbour of x_i, it follows that $H_n^\sharp$ is a minor of the graph obtained from F_{x_i} by including x_i and joining it to all its neighbours in H. Thus (i) implies that H'_n is a minor of H, and so $\mathcal{H}' \preccurlyeq \mathcal{H}$, as desired.

Case 2. There exists a $k \in \mathbb{N}$ such that for every $H \in \mathcal{H}$ all but at most k of the forests F_x $(x \in X_H)$ have path-width less than k.

For every $H \in \mathcal{H}$, let X'_H denote the set of all vertices $x \in X_H$ for which $\mathrm{pw}(F_x) < k$, and let $X''_H := X_H \setminus X'_H$. Put

$$G_H^1 := H[X'_H \cup \bigcup_{x \in X'_H} V(F_x)].$$

Let us first show that the vertices of H can be joined bijectively to the bottom row of the $|H| \times |H|$ grid in such a way that the graph G_H thus obtained can be k-embedded in the plane with one vortex, namely as $G_H - X''_H = G_H^0 \cup G_H^1$ where X''_H is the deleted set and G_H^0 is the subgraph of $G_H - X''_H$ obtained by deleting all edges of G_H^1. Indeed, since $H - X_H$ is a forest, the graph H' obtained from $H - X''_H$ by deleting all vertices of G_H^1 is also a forest. Since H' is obtained from $H - X_H$ by deleting $\bigcup_{x \in X'_H} F_x$, it follows from (ii) and (iii) that each component of H' is joined to G^1_H by at most one edge of H. Moreover, (i)–(iii) imply that G^1_H is obtained from the disjoint union of all F_x with $x \in X'_H$ by including every vertex $x \in X'_H$ and joining it to all its neighbours in F_x. Hence $\mathrm{pw}(G^1_H) \leq k$. Put $\ell := |G^1_H|$. Then there exists a path-decomposition $(p_1 \dots p_\ell, (X_{p_i})_{i=1}^{\ell})$ of G^1_H of width at most k and such that to each X_{p_i} one can assign a vertex $v_i \in X_{p_i}$ with $v_i \neq v_j$ for $i \neq j$. So $H - X''_H$ can be thought of as a 'wide path' (consisting of G^1_H) with trees added, each of them joined to the 'wide path' by at most one edge. Now join the vertices of $H - X''_H$ injectively to the bottom row R of the $|H| \times |H|$ grid so that the ordering of the neighbours of $V(G^1_H) = \{v_1, \dots, v_\ell\}$ in R corresponds to $v_1, \dots, v_\ell$, and so that the vertices of every component C of H' are joined without crossings to vertices of R, and if C is joined to v_i by an edge in H then these vertices of R lie between the neighbours of v_{i-1} and v_{i+1} in R (Fig. 5.4). Joining the vertices in X''_H to the

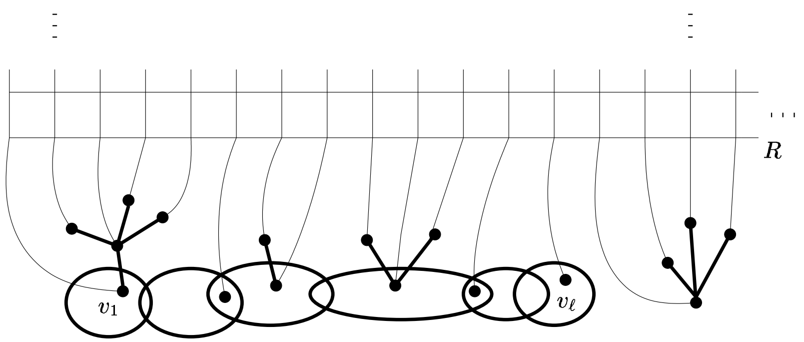

Figure 5.4: Joining the vertices of $H - X''_H$ to the bottom row of a grid.

remaining vertices of R then yields a graph G_H as desired.

The first part of the proof of Lemma 2.3 from [15] now implies that none of the graphs G_H ($H \in \mathcal{H}$) contains a K_{4k+8} minor. But since the bottom row of an $s \times s$ grid G is externally s-connected in G, it follows that every $H \in \mathcal{H}$ is externally $|H|$-connected in G_H, contradicting our assumption that $\mathcal{H}$ forces large complete minors. $\square$

Chapter 6

On well-quasi-ordering infinite trees

6.1 Introduction and terminology

A fundamental result of Nash-Williams [33] states that the infinite trees are well-quasi-ordered by the topological minor relation. To prove this, he introduced the stronger concept of better-quasi-ordered sets, and showed that the infinite trees are even better-quasi-ordered. In this chapter we give an essentially self-contained proof of this theorem. In general, the proof follows the lines of the original one. Nash-Williams's definition of a better-quasi-ordering is purely combinatorial; however, we use an equivalent topological concept, which is due to Simpson [44]. We remark that Laver [30] generalized Nash-Williams's result to a certain class of order theoretic trees. Thomas [47] extended Nash-Williams's result by proving that every class of infinite graphs with linked tree decompositions of bounded width is well-quasi-ordered by the minor relation.

We write $[n]$ for the set $\{1, \dots, n\}$. We denote by $\mathcal{C}$ the class of all cardinals, and by $\mathcal{O}$ that of all ordinals. We denote the domain of a function f by $\mathrm{D}f$.

For an infinite set $X \subseteq \mathbb{N}$ we define $X^{(\omega)}$ to be the set of all infinite subsets of X. We often identify an element $s \in X^{(\omega)}$ with the strictly ascending sequence whose elements are those of s; and conversely. Thus, if we write $s = (s_1, s_2, \dots)$ for an element of $X^{(\omega)}$, we mean that $s_1 < s_2 < \dots$. The *Ellentuck topology* on $X^{(\omega)}$ is defined by taking as basic open neighbourhoods of an element $s \in X^{(\omega)}$ all sets of the form $\{t \in s^{(\omega)} \mid u \subseteq t\}$, where u is a finite initial segment of s. Thus the Ellentuck topology is a refinement of the Tychonov (product) topology. Given a function $f : X^{(\omega)} \to D$, where D is some topological space, we say that f is *Ellentuck-continuous*, if f is continuous when we impose the Ellentuck topology on $X^{(\omega)}$. In particular, if D is discrete, then f is Ellentuck-continuous if and only if for every $s \in X^{(\omega)}$ there exists a finite initial segment u of s such that $f(s) = f(t)$ for all infinite subsequences t of s beginning with u.

We will repeatedly make use of the following theorem of Ellentuck, which says that Ellentuck-open sets are Ramsey (for a proof see e.g. [4, §20]). Apart from this, our presentation is self-contained.

Theorem 6.1 *Let $X \in \mathbb{N}^{(\omega)}$. For every Ellentuck-open set $A \subseteq X^{(\omega)}$ there exists $B \in X^{(\omega)}$ such that either $B^{(\omega)} \subseteq A$ or $B^{(\omega)} \cap A = \emptyset$.*

A reflexive and transitive relation is called a *quasi-ordering.* A quasi-ordered set $Q, \leq$ is *well-quasi-ordered* (*wqo*), if for every infinite sequence $q_1, q_2, \dots$ in Q there are indices $i < j$ such that $q_i \leq q_j$. In what follows Q will always denote a quasi-ordered set, and we also view Q as a discrete topological space. Q is *better-quasi-ordered* (*bqo*) if for every $X \in \mathbb{N}^{(\omega)}$ and for every Ellentuck-continuous function $f : X^{(\omega)} \to Q$ there exists an $s \in X^{(\omega)}$ such that $f(s) \leq f(s \backslash \{\min s\})$. We remark that a result of Mathias [31] implies that one obtains an equivalent definition by replacing Ellentuck-continuity by Tychonov-continuity, or by requiring Borel measurability. A *Q-array* is an Ellentuck-continuous function $f : X^{(\omega)} \to Q$, for some $X \in \mathbb{N}^{(\omega)}$. If there is no $s \in X^{(\omega)}$ such that $f(s) \leq f(s \backslash \{\min s\})$, then f is a *bad* Q-array. Thus Q is bqo if and only if there is no bad Q-array.

All trees considered in this chapter will have a root. For two trees T and U with roots t and u, respectively, we call an injective mapping $\varphi : V(T) \to V(U)$ an *embedding* of T into U, if φ can be extended to an isomorphism between a subdivision of T and the smallest subtree U' of U containing all vertices in $\varphi(V(T))$, and furthermore, the path between $\varphi(t)$ and u in U contains no vertex of U' other than $\varphi(t)$. We say that T is a *rooted topological minor* of U, abbreviated by $T \preccurlyeq U$, if there is an embedding of T into U. This defines a quasi-ordering on the class of all trees.

Given two vertices x and y of a tree T, we say that x is *above* y if y lies on the path from x to the root of T. If x and y are adjacent and x is above y, we call y the *predecessor* of x and x the *successor* of y. The *branch above* x, abbreviated by $\mathrm{br}(x)$, is the subtree of T spanned by all vertices above x (including x itself). For the root of $\mathrm{br}(x)$ we choose x.

6.2 Better-quasi-ordering infinite trees

Lemma 6.2 *Every bqo set Q is wqo.*

Proof. Let $q_1, q_2, \dots$ be any infinite sequence in Q. Define a function $f : \mathbb{N}^{(\omega)} \to Q$ by $f(s) := q_{\min s}$. Then f is Ellentuck-continuous, and thus a Q-array. Hence, since Q is bqo, there exists an $s \in \mathbb{N}^{(\omega)}$ such that $f(s) \leq f(s \backslash \{\min s\})$. But this means that $q_{s_1} \leq q_{s_2}$, where $s = (s_1, s_2, \dots)$. Thus Q is wqo. □

If Q is a quasi-ordered set, then we may quasi-order the elements of the power set of Q by saying that $A \leq B$ if for all $a \in A$ there exists $b \in B$ such that $a \leq b$ in Q. We denote the power set of Q with this quasi-ordering by $\mathcal{S}(Q)$. The following lemma implies that if Q is bqo then so is $\mathcal{S}(Q)$.

Lemma 6.3 *If f is a bad $\mathcal{S}(Q)$-array, then there exists a bad Q-array g such that $Dg = Df$ and $g(s) \in f(s)$ for all $s \in Dg$.*

Proof. Let $s \in \mathrm{D}f$. Since $f(s) \not\leq f(s \backslash \{\min s\})$ there exists an $x_s \in f(s)$ such that $x_s \not\leq y$ for all $y \in f(s \backslash \{\min s\})$. We can choose x_s such that it depends only on the pair $f(s), f(s\backslash\{\min s\})$ and not on s itself, i.e. if $f(s) = f(t)$ and $f(s\backslash\{\min s\}) = f(t\backslash\{\min t\})$, then $x_s = x_t$. We now define a function $g : \mathrm{D}f \to Q$ by setting $g(s) := x_s$. Then the Ellentuck-continuity of f and the fact that x_s depends only on the pair $f(s), f(s\backslash\{\min s\})$ imply that g is Ellentuck-continuous, and thus a Q-array. It is also bad, since $g(s) \leq g(s\backslash\{\min s\})$ would contradict the choice of x_s. $\square$

Given two quasi-ordered sets Q and Q', we define a quasi-ordering on $Q \times Q'$ by saying that $(q_1, q_1') \leq (q_2, q_2')$ if $q_1 \leq q_2$ and $q_1' \leq q_2'$.

Lemma 6.4 *If $f = (f_1, f_2)$ is a bad $\mathcal{C} \times Q$-array, then there exists a bad Q-array g such that $Dg \subseteq Df$ and $g(s) = f_2(s)$ for all $s \in Dg$.*

Proof. Let $A := \{s \in \mathrm{D}f \mid f_1(s) \leq f_1(s\backslash\{\min s\})\}$. Then the Ellentuck-continuity of f implies that A is Ellentuck-open. Hence by Theorem 6.1, there exists a $B \in \mathrm{D}f$ such that either $B^{(\omega)} \subseteq A$ or $B^{(\omega)} \cap A = \emptyset$. But the latter cannot hold, since then for $s = (s_1, s_2, \dots) \in B^{(\omega)}$ we would have $f_1(s_1, s_2, \dots) > f_1(s_2, s_3, \dots) > f_1(s_3, s_4, \dots) > \dots$, contradicting the fact that $\mathcal{C}$ is well-ordered. Thus $B^{(\omega)} \subseteq A$, and so $g : B^{(\omega)} \to Q$ defined by $g(s) := f_2(s)$ must be a bad Q-array, as required. $\square$

Let $\mathrm{Seq}(Q)$ be the set of all transfinite sequences with elements in Q. For a transfinite sequence $F : \alpha \to Q$ we define $\mathrm{length}(F)$ to be α. If $\beta < \alpha$, we write $F|_\beta$ for the restriction of F to β. Given $F, G \in \mathrm{Seq}(Q)$, we call a mapping $\varphi : \mathrm{length}(F) \to \mathrm{length}(G)$ an *embedding* of F into G if φ is strictly increasing and $F(\alpha) \leq G(\varphi(\alpha))$ for all $\alpha < \mathrm{length}(F)$. We impose a quasi-ordering on $\mathrm{Seq}(Q)$ by saying that $F \leq G$ if there exists an embedding from F into G.

The following lemma implies that if a set Q is bqo then so is $\mathrm{Seq}(Q)$, which is also a result due to Nash-Williams [34]. In the proof we present here, we closely follow Prömel and Voigt [35].

Lemma 6.5 *If f is a bad $\mathrm{Seq}(Q)$-array, then there exists a bad Q-array g such that $Dg \subseteq Df$ and $g(s) \in f(s)$ for all $s \in Dg$.*

Proof. For sequences $F, G \in \mathrm{Seq}(Q)$ we write $F \leq^* G$ if F is an initial segment of G, and $F <^* G$ if F is a proper initial segment of G. If h and h' are $\mathrm{Seq}(Q)$-arrays, we write $h \leq^* h'$ if $\mathrm{D}h \subseteq \mathrm{D}h'$ and $h(s) \leq^* h'(s)$ for all $s \in \mathrm{D}h$. Furthermore, we write $h <^* h'$ if $h \leq^* h'$ and there exists an $s \in \mathrm{D}h$ such that $h(s) <^* h'(s)$. We will first prove the following claim.

There exists a minimal bad $\mathrm{Seq}(Q)$-array h such that $h \leq^* f$. $\quad (*)$

We may assume that f itself is not minimal. Put $f_0 := f$ and $X_0^{(\omega)} := \mathrm{D}f_0$. For a $\mathrm{Seq}(Q)$-array g and $s \in \mathrm{D}g$ we define

$$k_{g,s} := \min\{k \mid k \in s \text{ such that } g(s) = g(t) \text{ for all } t \in s^{(\omega)} \text{ with } s \cap [k] = t \cap [k]\}.$$

Thus $k_{g,s}$ is the smallest integer $k \in s$ such that g is constant on the set of all $t \in s^{(\omega)}$ that begin with the initial segment $s \cap [k]$ of s. (Note that $k_{g,s}$ exists, since g is Ellentuck-continuous.) We now choose a bad Seq(Q)-array $f_1' <^* f_0$ such that

$$\min\{k_{f_1',s} \mid s \in \mathrm{D}f_1' \text{ with } f_1'(s) <^* f_0(s)\} \ =: \ k_1$$

is minimal. Choose an element $s_1 \in \mathrm{D}f_1'$ such that $f_1'(s_1) <^* f_0(s_1)$ and $k_{f_1',s_1} = k_1$. Define a function $f_1 : (s_1 \cup (X_0 \cap [k_1]))^{(\omega)} \to \mathrm{Seq}(Q)$ by

$$f_1(s) := \begin{cases} f_1'(s) & \text{if } s \in s_1^{(\omega)}; \\ f_0(s) & \text{otherwise.} \end{cases}$$

It is easily checked that f_1 is a bad Seq(Q)-array and $f_1 <^* f_0$. If f_1 is not minimal, we continue in this fashion to construct f_2', f_2, s_2 and k_2. Thus we may assume that we have constructed infinite sequences $f_1', f_2', \ldots$ and $f_1, f_2, \ldots$ and $s_1, s_2, \ldots$ and $k_1, k_2, \ldots$. Then $k_{i+1} \geq k_i$ for all $i \geq 1$, since f_{i+1}' was a candidate for the choice of f_i'. Moreover, the sequence (k_i) is unbounded. Indeed, suppose that there is an i such that $k_i = k_j$ for all $j \geq i$. Then there exists an infinite sequence $i < j_1 < j_2 < \ldots$ such that $s_{j_1} \cap [k_i] = s_{j_\ell} \cap [k_i]$ for all $\ell \geq 1$. This yields $s_{j_{\ell+1}} \subseteq s_{j_\ell}$ for all $\ell \geq 1$. Hence the definition of $k_{f_{j_\ell}', s_{j_\ell}}$ implies that

$$f_{j_\ell}(s_{j_\ell}) \ = \ f_{j_\ell}'(s_{j_\ell}) \ = \ f_{j_\ell}'(s_{j_{\ell+1}}) \ = \ f_{j_\ell}(s_{j_{\ell+1}}).$$

By the choice of $s_{j_{\ell+1}}$ it follows that

$$f_{j_{\ell+1}}(s_{j_{\ell+1}}) \ = \ f_{j_{\ell+1}}'(s_{j_{\ell+1}}) \ <^* \ f_{j_\ell}(s_{j_{\ell+1}}) \ = \ f_{j_\ell}(s_{j_\ell}).$$

Thus $\mathrm{length}(f_{j_1}(s_{j_1})), \mathrm{length}(f_{j_2}(s_{j_2})), \ldots$ is an infinite strictly descending chain of ordinals, a contradiction.

Let $X := \bigcap_{i \geq 1} X_i$, where $X_i^{(\omega)} := \mathrm{D}f_i$. Since X contains every k_i, the unboundedness of the sequence (k_i) implies that X is infinite. Also, note that for all $s \in X^{(\omega)}$ there exists an integer $i = i(s)$ such that $f_i(s) = f_j(s)$ for all $j \geq i$. (Otherwise there would be an infinite strictly descending chain of ordinals, since $f_{j+1}(s) \leq^* f_j(s)$.) Define a function $h' : X^{(\omega)} \to \mathrm{Seq}(Q)$ by putting $h'(s) := f_{i(s)}(s)$.

We will now find an Ellentuck-continuous restriction of h' that will do for h in $(*)$. Let A be the set of all $s \in X^{(\omega)}$ such that h' is Ellentuck-continuous in s. Thus A is Ellentuck-open. By Theorem 6.1 there exists a $B \in X^{(\omega)}$ such that either $B^{(\omega)} \subseteq A$ or $B^{(\omega)} \cap A = \emptyset$. Suppose first that the latter holds, and let $t_1 \in B^{(\omega)}$. Since $f_{i(t_1)}$ is Ellentuck-continuous, there is a basic Ellentuck-neighbourhood N_1 of t_1 on which $f_{i(t_1)}$ is constant. Since h' is not Ellentuck-continuous in t_1, there exists an $t_2 \in N_1$ such that $h'(t_2) \neq h'(t_1)$, and thus from the definition of h' it follows that $h'(t_2) <^* f_{i(t_1)}(t_2) = f_{i(t_1)}(t_1) = h'(t_1)$. But t_2 is a subsequence of t_1 (since it lies in a basic Ellentuck-neighbourhood of t_1), and so $t_2 \in B^{(\omega)}$. Continuing in this fashion we obtain an infinite sequence $t_1, t_2, \ldots$ such that $h'(t_1) >^* h'(t_2) >^* \ldots$, i.e. $\mathrm{length}(h'(t_1)), \mathrm{length}(h'(t_2)), \ldots$ is an infinite strictly descending chain of ordinals, a contradiction. Thus $B^{(\omega)} \subseteq A$, and hence the restriction h of h' on $B^{(\omega)}$ is Ellentuck-continuous.

The definition of h' implies that h is a bad Seq(Q)-array and $h \leq^* f$ (in fact, $h \leq^* f_i$ for all $i \geq 0$). Suppose that h is not minimal, and let φ be a bad

Seq(Q)-array such that $\varphi <^* h$. Let

$$k := \min\{k_{\varphi,s} \mid s \in \mathrm{D}\varphi \text{ and } \varphi(s) <^* h(s)\}.$$

Since the sequence (k_i) is unbounded, there is an i with $k_i > k$, contradicting the fact that φ was a candidate for the choice of f'_i. This shows that h is also minimal, and thus h is as required in $(*)$.

We now use $(*)$ to complete the proof of the lemma. For all $s \in \mathrm{D}h$ define

$$\psi(s) := \sup\{\alpha \in \mathcal{O} \mid h(s)|_\alpha \leq h(s\backslash\{\min s\})\}.$$

Then $\psi(s) < \mathrm{length}(h(s))$, since h is a bad Seq(Q)-array; and the Ellentuck-continuity of h implies that of ψ. Moreover, it is straightforward to show that

$$h(s)|_{\psi(s)} \leq h(s\backslash\{\min s\}),$$

but

$$h(s)|_{\psi(s)+1} \not\leq h(s\backslash\{\min s\}).$$

Let

$$C := \{s \in \mathrm{D}h \mid h(s)|_{\psi(s)} \leq h(s\backslash\{\min s\})|_{\psi(s\backslash\{\min s\})}\}.$$

Since both h and ψ are Ellentuck-continuous, C is Ellentuck-open. Thus by Theorem 6.1 there exists an $D \in \mathrm{D}h$ such that either $D^{(\omega)} \subseteq C$ or $D^{(\omega)} \cap C = \emptyset$. If the latter holds, then $\chi : D^{(\omega)} \to \mathrm{Seq}(Q)$ defined by $\chi(s) := h(s)|_{\psi(s)}$ would be a bad Seq(Q)-array with $\chi <^* h$, contradicting the choice of h.

Thus $D^{(\omega)} \subseteq C$. We now define $g : D^{(\omega)} \to Q$ by putting $g(s) := h(s)(\psi(s))$, the value of $h(s)$ at $\psi(s)$. Then g is Ellentuck-continuous, since h and ψ are. Moreover, $\mathrm{D}g \subseteq \mathrm{D}f$ and $g(s) \in f(s)$ for all $s \in \mathrm{D}g$. If there were an $s \in \mathrm{D}g$ with $g(s) \leq g(s\backslash\{\min s\})$, then we could define an embedding of $h(s)|_{\psi(s)+1}$ into $h(s\backslash\{\min s\})$ by first embedding $h(s)|_{\psi(s)}$ into $h(s\backslash\{\min s\})|_{\psi(s\backslash\{\min s\})}$ (this is possible since $s \in D^{(\omega)} \subseteq C$), and secondly, by sending $h(s)(\psi(s)) = g(s)$ to $h(s\backslash\{\min s\})(\psi(s\backslash\{\min s\})) = g(s\backslash\{\min s\})$. This contradicts the definition of ψ. Thus g is a bad Seq(Q)-array as required. □

If Q is a quasi-ordered set, we may quasi-order the elements of the power set of Q by saying that $A \leq B$ if there is an injective function $f : A \to B$ such that $a \leq f(a)$ in Q for all $a \in A$. Let $\mathcal{S}^\sharp(Q)$ denote the power set of Q with this quasi-ordering. Lemma 6.5 implies the following assertion.

Corollary 6.6 *If f is a bad $\mathcal{S}^\sharp(Q)$-array, then there exists a bad Q-array g such that $Dg \subseteq Df$ and $g(s) \in f(s)$ for all $s \in Dg$.*

An example of Rado [38] shows that there are wqo sets Q such that $\mathcal{S}(Q)$ (and thus also $\mathcal{S}^\sharp(Q)$ and Seq(Q)) are not wqo. This lack of closure properties under certain infinite operations is the reason why the stronger concept of bqo was introduced.

Denote the class of all trees by $\mathcal{R}$, and recall that the elements of $\mathcal{R}$ are quasi-ordered by the rooted topological minor relation. Let $\mathcal{R}_0$ be the subclass containing all trees T with the property that there is no infinite sequence

$x_1, x_2, \dots$ of vertices in T such that x_{i+1} is above x_i and $\mathrm{br}(x_i) \not\preccurlyeq \mathrm{br}(x_{i+1})$ for all $i \geq 1$. Given a tree T, let $\mathrm{S}(T)$ be the set of all its vertices x for which $T \not\preccurlyeq \mathrm{br}(x)$. If $x \in \mathrm{S}(T)$, we call $\mathrm{br}(x)$ a *strict branch* of T. For a vertex $x \in T$ we denote the set of its successors by $\mathrm{succ}(x)$, and let

$$\Gamma(x) \ := \ (\,|\{\mathrm{succ}(x)\backslash \mathrm{S}(T)\}|\,,\ \{\mathrm{br}(y)\,|\,y \in \mathrm{succ}(x) \cap \mathrm{S}(T)\}\,).$$

We view $\Gamma(x)$ as an element of the quasi-ordered set $\mathcal{C} \times \mathcal{S}^\sharp(\mathcal{R})$.

Lemma 6.7 *Suppose that T and U are trees such that for every vertex $x \in T$ there exists a vertex $y \in U$ with $\Gamma(x) \leq \Gamma(y)$. Then $T \preccurlyeq U$.*

Proof. For $n = 0, 1, \dots$, let W_n denote the set of all vertices of T which have distance at most n from the root of T. We shall inductively define an embedding φ of T into U such that, at stage n, we have defined φ on a set $V_n \subseteq V(T)$ satisfying the following conditions:

(i) $W_n \subseteq V_n$, and if $x \in V_n$ then the predecessor of x in T lies in V_n. If $x \in V_n \backslash W_n$, then $V(\mathrm{br}(x)) \subseteq V_n$.

(ii) Suppose that $x \in W_{n+1} \backslash V_n$, and let z be the predecessor of x. Then $x \notin \mathrm{S}(T)$ and there exists a vertex $v_x^n \in \mathrm{succ}(\varphi(z))\backslash \mathrm{S}(U)$ such that no vertex of $\mathrm{br}(v_x^n)$ lies in $\varphi(V_n)$. Furthermore, the vertices v_x^n are distinct for distinct $x \in W_{n+1} \backslash V_n$.

Let x_0 be the root of T. Then by the assumptions of the lemma, there is a vertex $y_0 \in U$ such that $\Gamma(x_0) \leq \Gamma(y_0)$. Thus for all $x \in \mathrm{succ}(x_0)$, there is a vertex $v_x^0 \in \mathrm{succ}(y_0)$ such that, firstly, the vertices v_x^0 are distinct for distinct x, secondly, if $x \notin \mathrm{S}(T)$, then $v_x^0 \notin \mathrm{S}(U)$, and thirdly, if $x \in \mathrm{S}(T)$, then $\mathrm{br}(x) \preccurlyeq \mathrm{br}(v_x^0)$. Put $\varphi(x_0) := y_0$, and extend φ by embedding $\mathrm{br}(x)$ into $\mathrm{br}(v_x^0)$ for all $x \in \mathrm{succ}(x_0) \cap \mathrm{S}(T)$. Setting

$$V_0 \ := \ \{x_0\} \cup \bigcup \{V(\mathrm{br}(x)) \mid x \in \mathrm{succ}(x_0) \cap \mathrm{S}(T)\}$$

starts the induction. Suppose that $n > 0$ and conditions (i) and (ii) hold for $n-1$. If $W_n \subseteq V_{n-1}$, then $V_{n-1} = V(T)$ by (i), and we are done. Thus let us assume that $W_n \not\subseteq V_{n-1}$, and let x be any vertex in $W_n \backslash V_{n-1}$. By the assumption of the lemma there is a vertex $y \in U$ such that $\Gamma(x) \leq \Gamma(y)$. Let v_x^{n-1} be as in condition (ii). Then $U \preccurlyeq \mathrm{br}(v_x^{n-1})$, since $v_x^{n-1} \notin \mathrm{S}(U)$. Let y' be the image of y in $\mathrm{br}(v_x^{n-1})$ under this embedding. The fact that $\Gamma(x) \leq \Gamma(y)$ now implies that for all $a \in \mathrm{succ}(x)$ there exists a vertex $v_a^n \in \mathrm{succ}(y')$ satisfying the following three conditions. Firstly, the v_a^n are distinct for distinct a. Secondly, if $a \notin \mathrm{S}(T)$ then $v_a^n \notin \mathrm{S}(U)$, and thirdly, if $a \in \mathrm{S}(T)$, then $\mathrm{br}(a) \preccurlyeq \mathrm{br}(v_a^n)$. Put $\varphi(x) := y'$ and extend φ further by embedding $\mathrm{br}(a)$ into $\mathrm{br}(v_a^n)$ for all $a \in \mathrm{succ}(x) \cap \mathrm{S}(T)$. Proceed similarly for every $x \in W_n \backslash V_{n-1}$. Then, setting

$$V_n \ := \ V_{n-1} \cup W_n \cup \bigcup \{V(\mathrm{br}(a))\,|\,a \in \mathrm{succ}(x) \cap \mathrm{S}(T) \text{ for some } x \in W_n \backslash V_{n-1}\}$$

completes the induction step. □

Lemma 6.8 *If f is a bad $\mathcal{R}_0$-array, then there exists a bad $\mathcal{R}_0$-array g such that $Dg \subseteq Df$ and $g(s)$ is a strict branch of $f(s)$ for all $s \in Dg$.*

Proof. For a tree $T \in \mathcal{R}_0$ we define $\Sigma(T) := \{\Gamma(x) \mid x \in T\}$ and think of it as an element of the quasi-ordered set $\mathcal{S}(\mathcal{C} \times \mathcal{S}^\sharp(\mathcal{R}_0))$. Lemma 6.7 implies that for all $T, U \in \mathcal{R}_0$,

$$\Sigma(T) \leq \Sigma(U) \;\Rightarrow\; T \preccurlyeq U.$$

Hence $\Sigma \circ f$ is a bad $\mathcal{S}(\mathcal{C} \times \mathcal{S}^\sharp(\mathcal{R}_0))$-array. By Lemma 6.3, there is a bad $\mathcal{C} \times \mathcal{S}^\sharp(\mathcal{R}_0)$-array φ such that $\mathrm{D}\varphi = \mathrm{D}\Sigma \circ f = \mathrm{D}f$ and $\varphi(s) \in \Sigma \circ f(s)$ for all $s \in \mathrm{D}\varphi$. Now Lemma 6.4 implies that there is a bad $\mathcal{S}^\sharp(\mathcal{R}_0)$-array ψ such that $\mathrm{D}\psi \subseteq \mathrm{D}\varphi$ and $\psi(s) = \varphi_2(s)$ for all $s \in \mathrm{D}\psi$. Finally, by Corollary 6.6, there is a bad $\mathcal{R}_0$-array g such that $\mathrm{D}g \subseteq \mathrm{D}\psi$ and $g(s) \in \psi(s)$ for all $s \in \mathrm{D}g$. Clearly, $\mathrm{D}g \subseteq \mathrm{D}f$. Furthermore, for all $s \in \mathrm{D}g$, $g(s)$ is an element of the second component of an element of $\Sigma \circ f(s)$, and thus a strict branch of $f(s)$, as required. $\square$

If h and h' are $\mathcal{R}_0$-arrays, we write $h \leq' h'$ if $\mathrm{D}h \subseteq \mathrm{D}h'$, and if $h(s)$ is a branch of $h'(s)$ for all $s \in \mathrm{D}h$. Furthermore, we write $h <' h'$ if $h \leq' h'$ and there exists an $s \in \mathrm{D}h$ such that $h(s)$ is a strict branch of $h'(s)$.

Lemma 6.9 *If f is a bad $\mathcal{R}_0$-array, then there exists a minimal bad $\mathcal{R}_0$-array h such that $h \leq' f$.*

We omit the proof, since it is an easy modification of the proof of assertion $(*)$ in the proof of Lemma 6.5. Indeed, the only difference is the following. In Lemma 6.5 we repeatedly made use of the fact that we could not have an infinite sequence $F_1, F_2, \ldots$ in $\mathrm{Seq}(Q)$ such that F_{i+1} is a proper initial segment of F_i for all $i \geq 1$, since $\mathrm{length}(F_1), \mathrm{length}(F_2), \ldots$ would then have been an infinite strictly descending chain of ordinals. In the proof of Lemma 6.9 an infinite sequence $F_1, F_2, \ldots$ in $\mathcal{R}_0$ such that F_{i+1} is a strict branch of F_i for all $i \geq 1$ would contradict the definition of $\mathcal{R}_0$.

Lemmas 6.8 and 6.9 immediately imply the following result.

Corollary 6.10 *$\mathcal{R}_0$ is bqo.*

Given a tree T, let $\mathrm{F}(T) := \{x \in T \mid \mathrm{br}(x) \in \mathcal{R}_0\}$ and $I(T) := V(T)\backslash F(T)$. For a vertex $x \in T$ define

$$\Delta(x) \;:=\; (\,|\mathrm{succ}(x) \cap \mathrm{I}(T)|\,,\, \{\mathrm{br}(z) \,|\, z \in \mathrm{succ}(x) \cap \mathrm{F}(T)\}\,).$$

We view $\Delta(x)$ as an element of $\mathcal{C} \times \mathcal{S}^\sharp(\mathcal{R}_0)$.

Lemma 6.11 *Suppose that T is a tree and $x_0, y_0 \in I(T)$ are such that*

$$\forall x \in br(x_0) \cap I(T) \;\; \forall y \in br(y_0) \cap I(T) \;\; \exists z \in br(y) : \;\; \Delta(x) \leq \Delta(z).$$

Then $br(x_0) \preccurlyeq br(y_0)$.

Proof. The proof is very similar to that of Lemma 6.7. For $n = 0, 1, \ldots,$ let W_n denote the set of all vertices of $\mathrm{br}(x_0)$ which have distance at most n from x_0. We shall inductively define an embedding φ of $\mathrm{br}(x_0)$ into $\mathrm{br}(y_0)$ such that, at stage n, we have defined φ on a set $V_n \subseteq V(\mathrm{br}(x_0))$ satisfying the following conditions:

(i) $W_n \subseteq V_n$, and if $x \in V_n$ then the predecessor of x in $\mathrm{br}(x_0)$ lies in V_n. If $x \in V_n \backslash W_n$, then $V(\mathrm{br}(x)) \subseteq V_n$.

(ii) Suppose that $x \in W_{n+1} \backslash V_n$, and let y be the predecessor of x. Then $x \in \mathrm{I}(T)$, and there exists a vertex $v_x^n \in \mathrm{succ}(\varphi(y)) \cap \mathrm{I}(U)$ such that no vertex of $\mathrm{br}(v_x^n)$ lies in $\varphi(V_n)$. Furthermore, the vertices v_x^n are distinct for distinct $x \in W_{n+1} \backslash V_n$.

By the assumptions of the lemma, there is a vertex $z_0 \in \mathrm{br}(y_0)$ such that $\Delta(x_0) \leq \Delta(z_0)$. Thus for all $x \in \mathrm{succ}(x_0)$ there is a vertex $v_x^0 \in \mathrm{succ}(z_0)$ such that, firstly, the vertices v_x^0 are distinct for distinct x, secondly, if $x \in \mathrm{I}(T)$ then $v_x^0 \in \mathrm{I}(T)$, and thirdly, if $x \in \mathrm{F}(T)$, then $\mathrm{br}(x) \preccurlyeq \mathrm{br}(v_x^0)$. Put $\varphi(x_0) := z_0$, and extend φ by embedding $\mathrm{br}(x)$ into $\mathrm{br}(v_x^0)$ for all $x \in \mathrm{succ}(x_0) \cap \mathrm{F}(T)$. Setting

$$V_0 \ := \ \{x_0\} \cup \bigcup \{V(\mathrm{br}(x)) \mid x \in \mathrm{succ}(x_0) \cap \mathrm{F}(T)\}$$

starts the induction. Suppose that $n > 0$ and conditions (i) and (ii) hold for $n-1$. If $W_n \subseteq V_{n-1}$, then $V_{n-1} = V(\mathrm{br}(x_0))$ by (i), and we are done. Thus we may assume that $W_n \not\subseteq V_{n-1}$. Let x be any vertex in $W_n \backslash V_{n-1}$, and let v_x^{n-1} be as in condition (ii). Then by the assumption of the lemma there is a vertex $z \in \mathrm{br}(v_x^{n-1})$ such that $\Delta(x) \leq \Delta(z)$. Thus for all $a \in \mathrm{succ}(x)$ there exists a vertex $v_a^n \in \mathrm{succ}(z)$ such that, firstly, the v_a^n are distinct for distinct a, secondly, if $a \in \mathrm{I}(T)$ then $v_a^n \in \mathrm{I}(T)$, and thirdly, if $a \in \mathrm{F}(T)$, then $\mathrm{br}(a) \preccurlyeq \mathrm{br}(v_a^n)$. Put $\varphi(x) := z$, and extend φ further by embedding $\mathrm{br}(a)$ into $\mathrm{br}(v_a^n)$ for all $a \in \mathrm{succ}(x) \cap \mathrm{F}(T)$. Proceed similarly for every $x \in W_n \backslash V_{n-1}$. Then, setting

$$V_n \ := \ V_{n-1} \cup W_n \cup \bigcup \{V(\mathrm{br}(a)) \,|\, a \in \mathrm{succ}(x) \cap \mathrm{F}(T) \text{ for some } x \in W_n \backslash V_{n-1}\}$$

completes the induction step. $\square$

Theorem 6.12 *The infinite trees are bqo by the rooted topological minor relation.*

Proof. By Corollary 6.10 it suffices to show that every tree lies in $\mathcal{R}_0$. Suppose not, and let T be a tree that does not lie in $\mathcal{R}_0$. Let x_0 be the root of T. Since $T \notin \mathcal{R}_0$, there is a vertex $y_1 \in \mathrm{br}(x_0) \cap \mathrm{I}(T)$ such that $\mathrm{br}(x_0) \not\preccurlyeq \mathrm{br}(y_1)$. Then Lemma 6.11 implies that there exist vertices $z_1 \in \mathrm{br}(x_0) \cap \mathrm{I}(T)$ and $x_1 \in \mathrm{br}(y_1) \cap \mathrm{I}(T)$ such that $\Delta(z_1) \not\leq \Delta(z)$ for all $z \in \mathrm{br}(x_1)$. Since $x_1 \in \mathrm{I}(T)$, there is a vertex $y_2 \in \mathrm{br}(x_1) \cap \mathrm{I}(T)$ such that $\mathrm{br}(x_1) \not\preccurlyeq \mathrm{br}(y_2)$. Again, Lemma 6.11 implies that there exist vertices $z_2 \in \mathrm{br}(x_1) \cap \mathrm{I}(T)$ and $x_2 \in \mathrm{br}(y_2) \cap \mathrm{I}(T)$ such that $\Delta(z_2) \not\leq \Delta(z)$ for all $z \in \mathrm{br}(x_2)$. Continuing in this fashion, we obtain an infinite sequence $z_1, z_2, \ldots$ such that $\Delta(z_i) \not\leq \Delta(z_j)$ in $\mathcal{C} \times \mathcal{S}^\sharp(\mathcal{R}_0)$ for all $1 \leq i < j$. But since $\mathcal{R}_0$ is bqo by Corollary 6.10, $\mathcal{C} \times \mathcal{S}^\sharp(\mathcal{R}_0)$ is bqo by Lemma 6.4 and Corollary 6.6, and thus it is wqo by Lemma 6.2, a contradiction. $\square$

Bibliography

[1] M. Ajtai, J. Komlós and E. Szemerédi, A dense infinite Sidon sequence, *Europ. J. Comb.* **2** (1981), 1–11.

[2] N. Alon and J. Spencer, *The Probabilistic Method* (2nd edition), Wiley-Interscience 2000.

[3] M.A. Armstrong, *Basic Topology*, Springer-Verlag 1983.

[4] B. Bollobás, *Combinatorics*, Cambridge University Press 1986.

[5] B. Bollobás, *Graph Theory*, Graduate Texts in Mathematics 63, Springer-Verlag 1979.

[6] R. Diestel, *Graph Theory*, Graduate Texts in Mathematics 173, Springer-Verlag 1997.

[7] R. Diestel, The end structure of a graph: recent results and open problems, *Discrete Mathematics* **100** (1992), 313–327.

[8] R. Diestel, Spanning trees and k-connectedness, *J. Combin. Theory B* **56** (1992), 263–277.

[9] R. Diestel, K.Yu. Gorbunov, T.R. Jensen and C. Thomassen, Highly connected sets and the excluded grid theorem, *J. Combin. Theory B* **75** (1999), 61–73.

[10] R. Diestel and D. Kühn, On infinite cycles I, *Combinatorica* (to appear).

[11] R. Diestel and D. Kühn, On infinite cycles II, submitted.

[12] R. Diestel and D. Kühn, Graph minor hierarchies, submitted.

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

[14] R. Diestel and I. Leader, Normal spanning trees, Aronszajn trees and excluded minors, *J. London Math. Soc.* **63** (2001), 16–32.

[15] R. Diestel and R. Thomas, Excluding a countable clique, *J. Combin. Theory B* **76** (1999), 41–67.

[16] P. Erdős, Problems and results in chromatic graph theory, in: F. Harary (ed.), *Proof Techniques in Graph Theory*, Academic Press 1969, 27–35.

[17] F. Galvin, I. Rival and B. Sands, A Ramsey-type theorem for traceable graphs, *J. Combin. Theory B* **33** (1982), 7–16.

[18] M. Grohe, Communication at the 1999 Oberwolfach conference on Graph Theory.

[19] P. Hajnal, Partition of graphs with condition on the connectivity and minimum degree, *Combinatorica* **3** (1983), 95–99.

[20] D.W. Hall and G.L. Spencer, *Elementary Topology*, John Wiley, New York 1955.

[21] H.A. Kierstead and S.G. Penrice, Radius two trees specify χ-bounded classes, *J. Graph Theory* **18** (1994), 119–129.

[22] J.B. Kruskal, Well-quasi-ordering, the tree theorem and Vázsonyi's conjecture, *Trans. Amer. Math. Soc.* **95** (1960), 210-225.

[23] D. Kühn, Forcing a K_r-minor by high external connectivity, submitted.

[24] D. Kühn, On well-quasi-ordering infinite trees – Nash-Williams's theorem revisited, *Math. Proc. Cambridge Philos. Soc.* **130** (2001), 401–408.

[25] D. Kühn and D. Osthus, Every graph of sufficiently large average degree contains a C_4-free subgraph of large average degree, *Combinatorica* (to appear).

[26] D. Kühn and D. Osthus, Induced cycles in $K_{s,s}$-free graphs of large average degree, submitted.

[27] D. Kühn and D. Osthus, Induced subdivisions in $K_{s,s}$-free graphs of large average degree, submitted.

[28] D. Kühn and D. Osthus, Partitions of graphs with high minimum degree or connectivity, submitted.

[29] D.G. Larman and P. Mani, On the existence of certain configurations within graphs and the 1-skeletons of polytopes, *Proc. London Math. Soc.* **20** (1970), 144–160.

[30] R. Laver, Better-quasi-orderings and a class of trees, *Studies in Foundations and Combinatorics*, Adv. in Math. Supplementary Studies **1** (1978), 31–48.

[31] A.R.D. Mathias, Happy families, *Ann. Math. Logic* **12** (1977), 59–111.

[32] C.St.J.A. Nash-Williams, Decompositions of graphs into closed and endless chains, *J. London Math. Soc.* (3) **10** (1960), 221–238.

[33] C.St.J.A. Nash-Williams, On well-quasi-ordering infinite trees, *Proc. Cambridge Philos. Soc.* **61** (1965), 697–720.

[34] C.St.J.A. Nash-Williams, On better-quasi-ordering transfinite sequences, *Proc. Cambridge Philos. Soc.* **64** (1968), 273–290.

[35] H.J. Prömel and B. Voigt, From wqo to bqo via Ellentuck's theorem, *Discrete Mathematics* **108** (1992), 83–106.

[36] L. Pyber, Regular subgraphs of dense graphs, *Combinatorica* **5** (1985), 347–349.

[37] L. Pyber, V. Rödl and E. Szemerédi, Dense graphs without 3-regular subgraphs, *J. Combin. Theory B* **63** (1995), 41–54.

[38] R. Rado, Partial well-ordering of sets of vectors, *Mathematika* **1** (1954), 89–95.

[39] B. Richter, Communication at the 1999 Oberwolfach conference on Graph Theory.

[40] N. Robertson and P.D. Seymour, Graph minors V. Excluding a planar graph, *J. Combin. Theory B* **41** (1986), 92–114.

[41] N. Robertson and P.D. Seymour, Graph minors XVI. Excluding a non-planar graph, *J. Combin. Theory B* (to appear).

[42] V. Rödl, On the chromatic number of subgraphs of a given graph, *Proc. Amer. Math. Soc.* **64** (1977), 370–371.

[43] A. Scott, Induced trees in graphs of large chromatic number, *J. Graph Theory* **24** (1997), 297–311.

[44] S.G. Simpson, Bqo theory and Fraïssé's conjecture, in: R. Mansfield and G. Weitkamp (eds.), *Recursion theoretic Aspects of Descriptive Set Theory*, Oxford University Press 1985, Chapter 9.

[45] J.M. Steele, Variations on the monotone subsequence theme of Erdős and Szekeres, in: D. Aldous et al. (ed.), *Discrete Probability and Algorithms*, *IMA Vol. 72*, Springer Verlag 1995, pp. 111–131.

[46] R. Thomas, A counterexample to "Wagner's conjecture" for infinite graphs, *Math. Proc. Cambridge Philos. Soc.* **103** (1988), 55–57.

[47] R. Thomas, Well-quasi-ordering infinite graphs with forbidden finite planar minor, *Trans. Amer. Math. Soc.* **312** (1989), 279–313.

[48] C. Thomassen, Embeddings and minors, in: R.L. Graham et al. (ed.), *Handbook of Combinatorics*, *Vol. 1*, Elsevier 1995, pp. 301–349.

[49] C. Thomassen, Non-separating cycles in k-connected graphs, *J. Graph Theory* **5** (1981), 351–354.

[50] C. Thomassen, Graph decomposition with constraints on the connectivity and minimum degree, *J. Graph Theory* **7** (1983), 165–167.

[51] C. Thomassen, Graph decomposition with applications to subdivisions and path systems modulo k, *J. Graph Theory* **7** (1983), 261–271.

[52] C. Thomassen, Girth in graphs, *J. Combin. Theory B* **35** (1983), 129–141.

[53] C. Thomassen, Subdivisions of graphs with large minimum degree, *J. Graph Theory* **8** (1984), 23–28.

Zusammenfassung

Kapitel 1 Unendliche Kreise

Das Ziel dieses Kapitels ist die Erweiterung von verschiedenen Sätzen über den Zyklenraum von endlichen Graphen auf unendliche Graphen. Hierzu wird auch der Begriff des Kreises verallgemeinert: als Kreise eines unendlichen Graphen G betrachten wir alle homöomorphen Bilder des Einheitskreises in einem topologischen Raum, der aus G und den Enden von G besteht.

Kapitel 2 Induzierte Unterteilungen in $K_{s,s}$-freien Graphen mit hohem Durchschnittsgrad

Ein klassischer Satz von Mader sagt, daß jeder Graph G eine Unterteilung eines gegebenen Graphen H enthält, wenn G nur genügend hohen Durchschnittsgrad hat. Das Hauptergebnis des zweiten Kapitels ist eine analoge Aussage für induzierte Unterteilungen: *Für jeden Graphen H und jedes $s \in \mathbb{N}$ gibt es ein $d \in \mathbb{N}$, so daß jeder Graph mit Durchschnittsgrad mindestens d entweder einen $K_{s,s}$ als Teilgraphen enthält oder eine induzierte Unterteilung von H.*

Kapitel 3 Teilgraphen mit hohem Durchschnittsgrad ohne Kreise der Länge < 6

In Kapitel 3 wird der Fall $g \leq 6$ der folgenden Vermutung von Thomassen bewiesen: *Für alle $k, g \in \mathbb{N}$ gibt es ein $d \in \mathbb{N}$, so daß jeder Graph mit Durchschnittsgrad mindestens d einen Teilgraphen mit Durchschnittsgrad mindestens k enthält, der keinen Kreis der Länge $< g$ hat.*

Kapitel 4 Partitionen von Graphen mit hohem Minimalgrad oder Zusammenhang

Eines der Ergebnisse aus Kapitel 4 ist die folgende Verallgemeinerung eines Satzes von Hajnal und Thomassen: *Für jedes $\ell \in \mathbb{N}$ gibt es ein $k \in \mathbb{N}$, so daß die Eckenmenge jedes k-zusammenhängenden Graphen G in zwei nicht-leere Mengen S und T zerlegt werden kann, so daß die von S und T induzierten Teilgraphen von G beide ℓ-zusammenhängend sind und jede Ecke aus S mindestens ℓ Nachbarn in T hat.* Zusammen mit dem oben erwähnten Satz von Mader folgt daraus, daß jeder genügend hoch zusammenhängende Graph G eine Unterteilung TH eines gegebenen Graphen H enthält, so daß $G - V(TH)$ noch immer hoch zusammenhängend ist.

Kapitel 5 Erzwingung von vollständigen Minoren durch hohen externen Zusammenhang

Kapitel 5 beschäftigt sich mit der Frage, welche Strukturen ein Graph H enthalten muß, wenn jeder Graph G, der H als extern hoch zusammenhängenden Teilgraphen enthält, einen großen vollständigen Minor hat. Es werden vier minimale Strukturen mit dieser Eigenschaft identifiziert.

Kapitel 6 Wohlquasiordnung unendlicher Bäume

Das letzte Kapitel enthält einen kurzen Beweis des folgenden Satzes von Nash-Williams: *Die unendlichen Bäume sind wohlquasigeordnet, d.h. zu jeder unendlichen Folge* $T_1, T_2, \ldots$ *von unendlichen Bäumen gibt es Indizes* $i < j$, *so daß* T_j *eine Unterteilung von* T_i *enthält.*

Curriculum Vitae

Daniela Kühn

Persönliche Daten

Geboren	am 12.12.1973 in Chemnitz
Staatsangehörigkeit	deutsch
Familienstand	verheiratet mit Deryk Osthus

Ausbildung

1980 – 1990 Besuch der Polytechnischen Oberschule in Chemnitz

1990 – 1992 Besuch der Erweiterten Oberschule (Gymnasium) in Chemnitz

Juni 1992 Abitur (Durchschnittsnote 1,3)

Okt. 1992 – Juli 1998
Mathematikstudium an der TU Chemnitz und an der University of Cambridge
Nebenfach: Betriebswirtschaftslehre

Aug. 1994 Diplomvorprüfung ("sehr gut")

Juli 1995 Aufnahme in die Studienstiftung

Okt. 1995 – Juni 1997
Studium in Cambridge (St. John's College)

Juni 1996 Teilnahme an der Abschlußprüfung des B.A. Mathematikstudiengangs (Note: 1st class)

Juni 1997 "Certificate of Advanced Studies in Mathematics" bestanden mit Auszeichnung ("Distinction")

Juli 1998 Diplom an der TU Chemnitz "mit Auszeichnung"
Diplomarbeit über "Minoren-universelle plättbare Graphen"

Sept. 1998 – Juni 2001
Promotionsstudium an der TU Chemnitz und an der Universität Hamburg als Stipendiat der Studienstiftung
Betreuer: Professor Reinhard Diestel, PhD