

UBIQUITY OF LOCALLY FINITE GRAPHS WITH EXTENSIVE TREE-DECOMPOSITIONS

Nathan Bowler¹, Christian Elbracht², Joshua Erde^{*3}, J. Pascal Gollin^{†4}, Karl Heuer^{‡5},
Max Pitz⁶, and Maximilian Teegen⁷

^{1,2,6,7}Universität Hamburg, Department of Mathematics, Bundesstraße 55 (Geomatikum), 20146 Hamburg, Germany
nathan.bowler@uni-hamburg.de, christian.elbracht@uni-hamburg.de, max.pitz@uni-hamburg.de,
maximilian.teegen@uni-hamburg.de

³Graz University of Technology, Institute of Discrete Mathematics, Steyrergasse 30, 8010 Graz, Austria
erde@math.tugraz.at

⁴FAMNIT, University of Primorska, Glagoljaška 8, 6000 Koper, Slovenia
pascal.gollin@famnit.upr.si

⁵Technical University of Denmark, Department of Applied Mathematics and Computer Science, Richard Petersens Plads,
Building 322, 2800 Kongens Lyngby, Denmark
karheu@dtu.dk

Submitted: Jan 19, 2022; Accepted: Apr 10, 2024; Published: Sep 30, 2024

© The authors. Released under the CC BY license (International 4.0).

Abstract. A graph G is said to be *ubiquitous*, if every graph Γ that contains arbitrarily many disjoint G -minors automatically contains infinitely many disjoint G -minors. The well-known *Ubiquity conjecture* of Andreae says that every locally finite graph is ubiquitous.

In this paper we show that locally finite graphs admitting a certain type of tree-decomposition, which we call an *extensive tree-decomposition*, are ubiquitous. In particular this includes all locally finite graphs of finite tree-width, and also all locally finite graphs with finitely many ends, all of which have finite degree. It remains an open question whether every locally finite graph admits an extensive tree-decomposition.

Keywords. Graph minors, infinite graphs, ubiquity

Mathematics Subject Classifications. 05C83, 05C63

*Supported by the Austrian Science Fund (FWF) 10.55776/P36131. For the purpose of open access, the author has applied a CC BY public copyright licence to any Author Accepted Manuscript version arising from this submission.

†Supported by the Institute for Basic Science (IBS-R029-Y3) and in part by the Slovenian Research and Innovation Agency (research project N1-0370).

‡Supported by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (ERC consolidator grant DISTRUCT, agreement No. 648527).

1. Introduction

Given a graph G and some relation \triangleleft between graphs, we say that G is \triangleleft -ubiquitous if whenever Γ is a graph such that $nG \triangleleft \Gamma$ for all $n \in \mathbb{N}$, then $\aleph_0 G \triangleleft \Gamma$, where αG is the disjoint union of α many copies of G . A classic result of Halin [Hal65, Satz 1] says that the ray, i.e. a one-way infinite path, is \subseteq -ubiquitous, where \subseteq is the subgraph relation. That is, any graph which contains arbitrarily large collections of vertex-disjoint rays must contain an infinite collection of vertex-disjoint rays. Later, Halin showed that the double ray, i.e. a two-way infinite path, is also \subseteq -ubiquitous [Hal70]. However, not all graphs are \subseteq -ubiquitous, and in fact even trees can fail to be \subseteq -ubiquitous (see for example [Woo76]).

The question of ubiquity for classes of graphs has also been considered for other graph relations. In particular, whilst there are still reasonably simple examples of graphs which are not \leq -ubiquitous (see [Lak76, And77]), where \leq is the topological minor relation, it was shown by Andreae that all rayless countable graphs [And80] and all locally finite trees [And79] are \leq -ubiquitous. The latter result was recently extended to the class of all trees by the present authors [BEE⁺22].

In [And02] Andreae initiated the study of ubiquity of graphs with respect to the minor relation \preceq . He constructed a graph which is not \preceq -ubiquitous, however the construction relies on the existence of a counterexample to the well-quasi-ordering of infinite graphs under the minor relation, for which only examples of uncountable size are known [Kom95, Pit23, Tho88]. In particular, the question of whether there exists a countable graph which is not \preceq -ubiquitous remains open.

Andreae conjectured that at least all *locally finite* graphs, those with all degrees finite, should be \preceq -ubiquitous.

The Ubiquity Conjecture. *Every locally finite connected graph is \preceq -ubiquitous.*

In [And13] Andreae established the following pair of results, demonstrating that his conjecture holds for wide classes of locally finite graphs. Recall that a *block* of a graph is a maximal 2-connected subgraph, and that a graph has *finite tree-width* if there is an integer k such that the graph has a tree-decomposition of width k .

Theorem 1.1 (Andreae, [And13, Corollary 1]). *Let G be a locally finite, connected graph with finitely many ends such that every block of G is finite. Then G is \preceq -ubiquitous.*

Theorem 1.2 (Andreae, [And13, Corollary 2]). *Let G be a locally finite, connected graph of finite tree-width such that every block of G is finite. Then G is \preceq -ubiquitous.*

Note, in particular, that if G is such a graph, then the degree of every end in G must be one.¹ The main result of this paper is a far-reaching extension of Andreae's results, removing the assumption of finite blocks.

¹A precise definitions of the ends of a graph and their degree can be found in Section 3.

Theorem 1.3. *Let G be a locally finite, connected graph with finitely many ends such that every end of G has finite degree. Then G is \preceq -ubiquitous.*

Theorem 1.4. *Every locally finite, connected graph of finite tree-width is \preceq -ubiquitous.*

The reader may have noticed that these results are of a similar flavour: they all make an assertion that locally finite graphs which are built by pasting finite graphs in a tree like fashion are ubiquitous – with differing requirements on the size of the finite graphs, how far they are allowed to overlap, and the structure of the underlying decomposition trees. And indeed, behind all the above results there are unifying but more technical theorems, the strongest of which is the true main result of this paper:

Theorem 1.5. *Every locally finite connected graph admitting an extensive tree-decomposition is \preceq -ubiquitous.*

The precise definition of an extensive tree-decomposition is somewhat involved and will be given in detail in Section 4 up to Theorem 4.6. Roughly, however, it implies that we can find many self-minors of the graph at spots whose precise positions are governed by the decomposition tree. We hope that the proof sketch in Section 2 is a good source for additional intuition before the reader delves into the technical details. In fact, we note that we do not know of any construction of a locally finite connected graph which *does not* admit an extensive tree-decomposition.

Question 1.6. Do all locally finite connected graphs admit an extensive tree-decomposition?

To summarise, we are facing two main tasks in this paper. One is to prove our main ubiquity result, Theorem 1.5. This will occupy the second part of this paper, Sections 6 to 8. And as our other task, we also need to prove that the graphs in Theorems 1.3 and 1.4 do indeed possess such extensive tree-decompositions.

This analysis occupies Section 4 and 5. The proof uses in an essential way certain results about the well-quasi-ordering of graphs under the minor relation, including Thomas's result [Tho89] that for all $k \in \mathbb{N}$, the classes of graphs of tree-width at most k are well-quasi-ordered under the minor relation. In fact, the class of locally finite graphs having an extensive tree-decomposition is certainly larger than the results stated in Theorems 1.3 and 1.4; for example, it is easy to see that the infinite grid $\mathbb{N} \times \mathbb{N}$ has such an extensive tree-decomposition. It remains an open question whether *every* locally finite graph has an extensive tree-decomposition. A more precise discussion of how this problem relates to the theory of well-quasi- and better-quasi-orderings of finite graphs will be given in Section 9.

But first, in Section 2 we will give a sketch of the key ideas in the proof, at the end of which we will provide a more detailed overview of the structure and the different sections of this paper.

2. Proof sketch

To give a flavour of the main ideas in this paper, let us begin by considering the case of a locally finite connected graph G with a single end ω , where ω has finite degree $d \in \mathbb{N}$ (this means that there is a family $(A_i : 1 \leq i \leq d)$ of d disjoint rays in ω , but no family of more than d such rays).

Our construction will exploit the fact that graphs of this kind have a very particular structure. More precisely, there is a tree-decomposition $(S, (V_s)_{s \in V(S)})$ of G , where $S = s_0 s_1 s_2 \dots$ is a ray and such that, if we denote V_{s_n} by V_n and $G[\bigcup_{l \geq n} V_l]$ by G_n for each n , the following holds:

1. each V_n is finite;
2. $|V_i \cap V_j| = d$ if $|i - j| = 1$, and $|V_i \cap V_j| = 0$ otherwise;
3. all the A_i begin in V_0 ;
4. for each $m \geq 1$ there are infinitely many $n > m$ such that G_m is a minor of G_n , in such a way that for any edge e of G_m and any $i \leq d$, the edge e is contained in A_i if and only if the edge representing it in this minor is.

Property 4 seems rather strong – it is a first glimpse of the strength of extensive tree-decompositions alluded to in Theorem 1.5. The reason it can always be achieved has to do with the well-quasi-ordering of finite graphs. For details of how this works, see Section 5. The sceptical reader who does not yet see how to achieve this may consider the argument in this section as showing ubiquity simply for graphs G with a decomposition of the above kind.

Now we suppose that we are given some graph Γ such that $nG \preceq \Gamma$ for each n , and we wish to show that $\aleph_0 G \preceq \Gamma$. Consider a G -minor H in Γ . Any ray R of G can be expanded to a ray $H(R)$ in the copy H of G in Γ , and since G only has one end, all rays $H(R)$ go to the same end ϵ_H of Γ ; we shall say that H goes to the end ϵ_H .

Techniques from an earlier paper [BEE⁺22] show that we may assume that there is some end ϵ of Γ such that all G -minors in Γ go to ϵ , otherwise it can be shown that $\aleph_0 G \preceq \Gamma$.

From any G -minor H we obtain rays $H(A_i)$ corresponding to our marked rays A_i in G , which by the above all go to ϵ . We will call this family of rays the *bundle* of rays given by H .

Our aim now is to build up an $\aleph_0 G$ -minor of Γ recursively. At stage n we hope to construct n disjoint $G[\bigcup_{m \leq n} V_m]$ -minors $H_1^n, H_2^n, \dots, H_n^n$, such that for each such H_m^n there is a family $(R_{m,i}^n : i \leq d)$ of disjoint rays in ϵ , where the path in H_m^n corresponding to the initial segment of the ray A_i in $\bigcup_{m \leq n} G_m$ is an initial segment of $R_{m,i}^n$, but these rays are otherwise disjoint from the various H_l^n and from each other, see Figure 2.1. We aim to do this in such a way that each H_m^n extends all previous H_m^l for $l \leq n$, so that at the end of our construction we can obtain infinitely many disjoint G -minors as $(\bigcup_{n \geq m} H_m^n : m \in \mathbb{N})$. The rays chosen at later stages need not bear any relation to those chosen at earlier stages; we just need them to exist so that there is some hope of continuing the construction.

We will again refer to the families $(R_{m,i}^n : i \leq d)$ of rays starting at the various H_m^n as the *bundles* of rays from those H_m^n .

The rough idea for getting from the n^{th} to the $n + 1^{\text{st}}$ stage of this construction is now as follows: we choose a very large family \mathcal{H} of disjoint G -minors in Γ . We discard all those which meet any previous H_m^n and we consider the family of rays corresponding to the A_i in the remaining minors. Then it is possible to find a collection of paths transitioning from the $R_{m,i}^n$ from stage n onto these new rays. Precisely what we need is captured in the following definition, which also introduces some helpful terminology for dealing with such transitions:




Figure 2.1: Stage $n = 3$ of the construction with n disjoint $G[V_1 \cup V_2 \cup V_3]$ -minors H_i^n with their bundles of $d = 3$ disjoint rays.

Definition 2.1 (Linkage of families of rays). Let $\mathcal{R} = (R_i : i \in I)$ and $\mathcal{S} = (S_j : j \in J)$ be families of disjoint rays, where the initial vertex of each R_i is denoted x_i . A family of paths $\mathcal{P} = (P_i : i \in I)$, is a *linkage* from \mathcal{R} to \mathcal{S} if there is an injective function $\sigma : I \rightarrow J$ such that

- each P_i goes from a vertex $x'_i \in R_i$ to a vertex $y_{\sigma(i)} \in S_{\sigma(i)}$;
- the family $\mathcal{T} = (x_i R_i x'_i P_i y_{\sigma(i)} S_{\sigma(i)} : i \in I)$ is a collection of disjoint rays.² We write $\mathcal{R} \circ_{\mathcal{P}} \mathcal{S}$ for the family \mathcal{T} as well $R_i \circ_{\mathcal{P}} \mathcal{S}$ for the ray in \mathcal{T} with initial vertex x_i .

We say that \mathcal{T} is obtained by *transitioning* from \mathcal{R} to \mathcal{S} along the linkage. We say the linkage \mathcal{P} *induces* the mapping σ . We further say that \mathcal{P} *links* \mathcal{R} to \mathcal{S} . Given a set X we say that the linkage is *after* X if $X \cap V(R_i) \subseteq V(x_i R_i x'_i)$ for all $i \in I$ and no other vertex in X is used by the members of \mathcal{T} .

Thus, our aim is to find a linkage from the $R_{m,i}^n$ to the new rays after all the H_m^n . That this is possible is guaranteed by the following lemma from [BEE⁺22]:

Lemma 2.2 (Weak linking lemma [BEE⁺22, Lemma 4.3]). *Let Γ be a graph and $\omega \in \Omega(\Gamma)$. Then, for any families $\mathcal{R} = (R_1, \dots, R_n)$ and $\mathcal{S} = (S_1, \dots, S_n)$ of vertex disjoint rays in ω and any finite set X of vertices, there is a linkage from \mathcal{R} to \mathcal{S} after X .*

The aim is now to use property 4 of our tree-decomposition of G to find minor-copies of $G[V_{n+1}]$ sufficiently far along the new rays that we can stick them onto our H_m^n to obtain suitable H_m^{n+1} . There are two difficulties at this point in this argument. The first is that, as well as extending the existing H_m^n to H_m^{n+1} , we also need to introduce an H_{n+1}^{n+1} . To achieve this, we ensure that one of the G -minors in \mathcal{H} is disjoint from all the paths in the linkage, so that we may take an initial segment of it as H_{n+1}^{n+1} . This is possible because of a slight strengthening of the linking lemma above; see [BEE⁺22, Lemma 4.4] or Lemma 3.16 for a precise statement.

A more serious difficulty is that in order to stick the new copy of V_{n+1} onto H_m^n we need the following property:

²Where we use the notation as in [Die17], see also Definition 3.3.

For each of the bundles corresponding to an H_m^n , the rays in the bundle are linked to the rays in the bundle coming from some $H \in \mathcal{H}$. This happens in such a way that each $R_{m,i}^n$ is linked to $H(A_i)$. (*)

Thus we need a great deal of control over which rays get linked to which. We can keep track of which rays are linked to which as follows:

Definition 2.3 (Transition function). Let $\mathcal{R} = (R_i: i \in I)$ and $\mathcal{S} = (S_j: j \in J)$ be families of disjoint rays. We say that a function $\sigma: I \rightarrow J$ is a *transition function* from \mathcal{R} to \mathcal{S} if for any finite set X of vertices there is a linkage from \mathcal{R} to \mathcal{S} after X that induces σ .

So our aim is to find a transition function assigning new rays to the R_m^n so as to achieve (*). One reason for expecting this to be possible is that the new rays all go to the same end, and so they are joined up by many paths. We might hope to be able to use these paths to move between the rays, allowing us some control over which rays are linked to which. The structure of possible jumps is captured by a graph whose vertex set is the set of rays:

Definition 2.4 (Ray graph). Given a finite family of disjoint rays $\mathcal{R} = (R_i: i \in I)$ in a graph Γ the *ray graph*, $\text{RG}_\Gamma(\mathcal{R}) = \text{RG}_\Gamma(R_i: i \in I)$ is the graph with vertex set I and with an edge between i and j if there is an infinite collection of vertex disjoint paths from R_i to R_j which meet no other R_k . When the host graph Γ is clear from the context we will simply write $\text{RG}(\mathcal{R})$ for $\text{RG}_\Gamma(\mathcal{R})$.

Unfortunately, the collection of possible transition functions can be rather limited. Consider, for example, the case of families of disjoint rays in the grid. Any such family has a natural cyclic order, and any transition function must preserve this cyclic order. This paucity of transition functions is reflected in the sparsity of the ray graphs, which are all just cycles.

However, in a previous paper [BEE⁺23] we analysed the possibilities for how the ray graphs and transition functions associated to a given thick³ end may look. We found that there are just three possibilities.

The easiest case is that in which the rays to the end are very joined up, in the sense that any injective function between two families of rays is a transition function. This case was already dealt with in [BEE⁺23], where it was shown that in any graph with such an end we can find a K_{\aleph_0} minor. The second possibility is that which we saw above for the grid: all ray graphs are cycles, and all transition functions between them preserve the cyclic order. The third possibility is that all ray graphs consist of a path together with a bounded number of further ‘junk’ vertices, where these junk vertices are hanging at the ends of the paths (formally: all interior vertices on this *central path* in the ray graph have degree 2). In this case, the transition functions must preserve the linear order along the paths.

The second and third cases can be dealt with using similar ideas, so we will focus on the third one here.

Since we are assuming that all the G -minors in Γ go to ϵ , given a large enough collection of G -minors \mathcal{H} , almost all of the rays from the bundles of the $H \in \mathcal{H}$ lie on the central path of the

³An end is *thick* if it contains infinitely many disjoint rays.

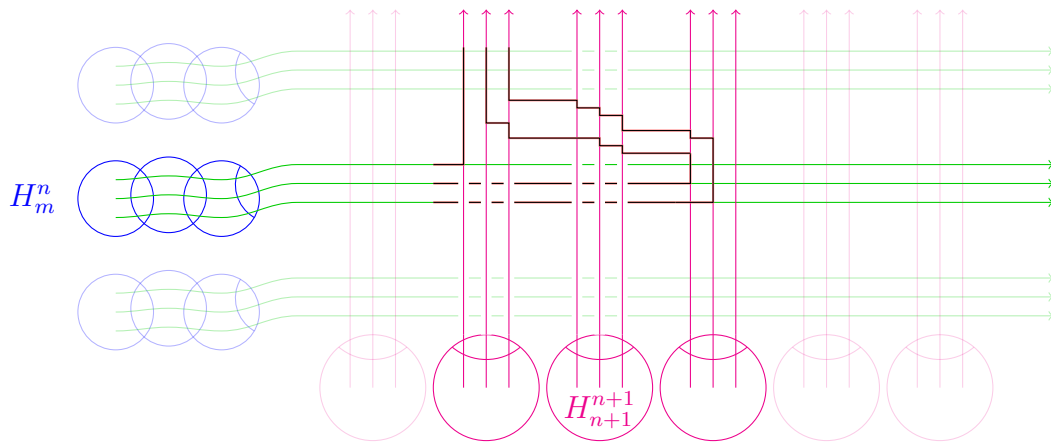


Figure 2.2: Extending the H_m^n by routing onto a set of disjoint G -minors might cause problems with introducing a new H_{n+1}^{n+1} disjoint to the rest.

ray graph of this family of rays, and so in particular by a Ramsey type argument there must be a large collection of $H \in \mathcal{H}$ such that for each H , the rays $H(A_i)$ appear in the same order along the central path.

Since there are only finitely many possible orders, there is some consistent way to order the A_i such that for every n we can find n disjoint G -minors H such that there is some ray graph in which, for each H , the rays $H(A_i)$ appear in this order along the central path, which we can assume, without loss of generality, is from $H(A_1)$ to $H(A_d)$.

This will allow us to recursively maintain a similar property for the rays from the bundles of the H_m^n . More precisely, we can guarantee that there is a slightly larger family \mathcal{R} of disjoint rays, consisting of the $R_{m,i}^n$ and some extra ‘junk’ rays, such that all of the $R_{m,i}^n$ lie on the central path of $\text{RG}(\mathcal{R})$, and for each n and m the $R_{m,i}^n$ appear on this path consecutively in order from $R_{m,1}^n$ to $R_{m,k}^n$.

Then, our extra assumption on the structure of the end ϵ ensures that given a linkage from \mathcal{R} to the bundles from $H \in \mathcal{H}$ which induces a transition function, we can reroute our linkage, using the edges of $\text{RG}(\mathcal{R})$, so that $(*)$ holds.

There is one last subtle difficulty which we have to address, once more relating to the fact that we want to introduce a new H_{n+1}^{n+1} together with its private bundle of rays corresponding to its copies of the A_i , disjoint from all the other H_m^{n+1} and their bundles. Our strengthening of the weak linking lemma allows us to find a linkage which avoids one of the G -minors in \mathcal{H} , but this linkage may not have property $(*)$.

We can, as before, modify it to one satisfying $(*)$ by rerouting the linkage, but this new linkage may then have to intersect some of the rays in the bundle of H_{n+1}^{n+1} , if these rays from H_{n+1}^{n+1} lie between rays linked to a bundle of some H_m^n , see Figure 2.2.

However, we can get around this by instead rerouting the rays in \mathcal{R} before the linkage, so as to rearrange which bundles make use of (the tails of) which rays. Of course, we cannot know before we choose our linkage how we will need to reroute the rays in \mathcal{R} , but we do know that the structure of ϵ restricts the possible reroutings we might need to do.

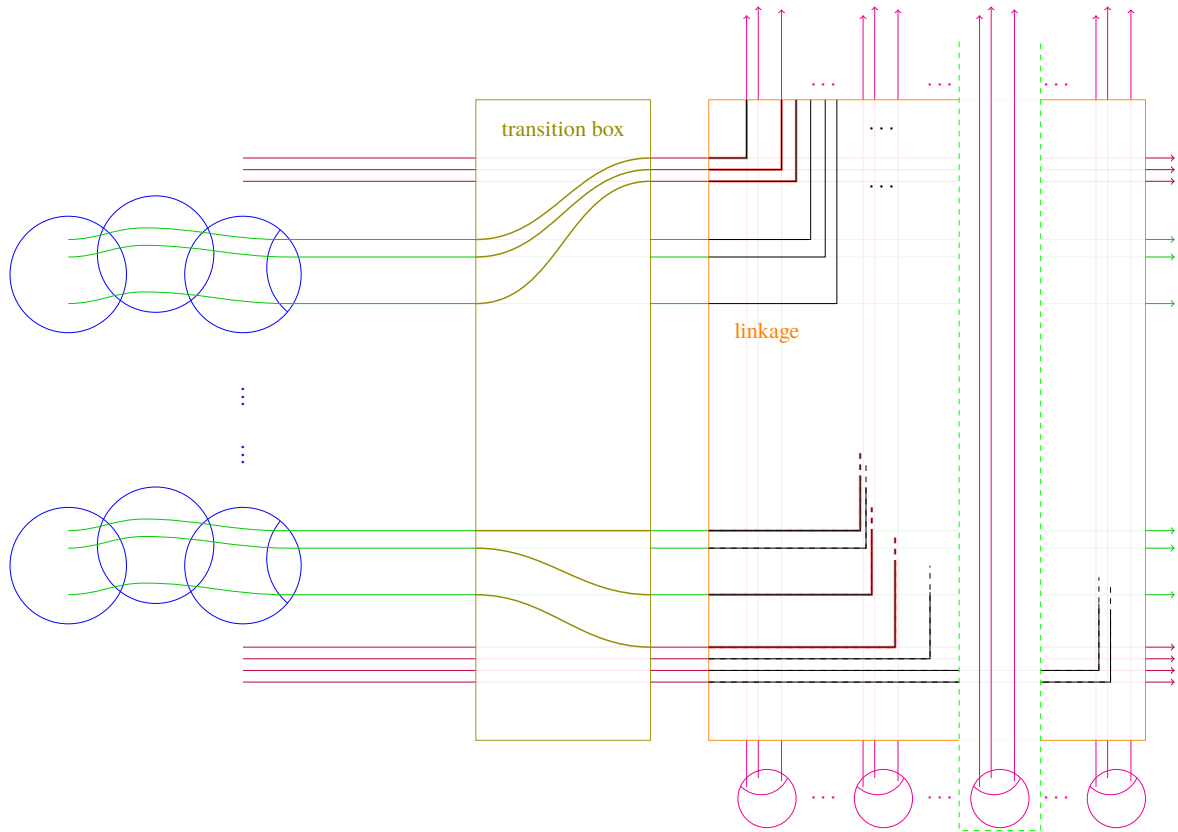


Figure 2.3: The transitioning strategy between the old and new bundles.

Hence, we can avoid this issue by first taking a large, but finite, set of paths between the rays in \mathcal{R} which is rich enough to allow us to reroute the rays in \mathcal{R} in every way which is possible in Γ . Since the rays in \mathcal{R} also go to ϵ , the structure of ϵ will guarantee that this includes all of the possible reroutings we might need to do. We call such a collection a *transition box*.

Only after building our transition box do we choose the linkage from \mathcal{R} to the rays from \mathcal{H} , and we make sure that this linkage is after the transition box. Then, when we later see how the rays in \mathcal{R} should be arranged in order that the rays from the bundle of H_{n+1}^{n+1} do not appear between rays linked to a bundle of some H_m^n , we can go back and perform a suitable rerouting within the transition box, see Figure 2.3.

This completes the sketch of the proof that locally finite graphs with a single end of finite degree are ubiquitous. Our results in this paper are for a more general class of graphs, but one which is chosen to ensure that arguments of the kind outlined above will work for them. Hence we still need a tree-decomposition with properties similar to 1–4 from our ray-decomposition above. Tree-decompositions with these properties are called *extensive*, and the details can be found in Section 4.

However, certain aspects of the sketch above must be modified to allow for the fact that we are now dealing with graphs G with multiple, indeed possibly infinitely many, ends. For any

end δ of G and any G -minor H of Γ , all rays $H(R)$ with R in δ belong to the same end $H(\delta)$ of Γ . If δ and δ' are different ends in G , then $H(\delta)$ and $H(\delta')$ may well be different ends in Γ as well.

So there is no hope of finding a single end ϵ of Γ to which all rays in all G -minors converge. Nevertheless, we can still find an end ϵ of Γ towards which the G -minors are *concentrated*, in the sense that for any finite vertex set X there are arbitrarily large families of G -minors in the same component of $G - X$ as all rays of ϵ have tails in. See Section 7 for details. In that section we introduce the term *tribe* for a collection of arbitrarily large families of disjoint G -minors.

The recursive construction will work pretty much as before, in that at each step n we will again have embedded G^n -minors for some large finite part G^n of G , together with a number of rays in ϵ corresponding to some designated rays going to certain ends δ of G .

In order for this to work, we need some consistency about which ends $H(\delta)$ of Γ are equal to ϵ and which are not. It is clear that for any finite set Δ of ends of G there is some subset Δ' such that there is a tribe of G -minors H converging to ϵ with the property that the set of δ in Δ with $H(\delta) = \epsilon$ is Δ' . This is because there are only finitely many options for this set. But if G has infinitely many ends, there is no reason why we should be able to do this for all ends of G at once.

Our solution is to keep track of only finitely many ends of G at any stage in the construction, and to maintain at each stage a tribe concentrated towards ϵ which is consistent for all these finitely many ends. Thus in our construction consistency of questions such as which ends δ of G converge to ϵ or of the proper linear order in the ray graph of the families of canonical rays to those ends is achieved dynamically during the construction, rather than being fixed in advance. The ideas behind this dynamic process have already been used successfully in our earlier paper [BEE⁺22], where they appear in slightly simpler circumstances.

The paper is then structured as follows. In Section 3 we give precise definitions of some of the basic concepts we will be using, and prove some of their fundamental properties. In Section 4 we introduce extensive tree-decompositions and in Section 5 we illustrate that many locally finite graphs admit such decompositions. In Section 6 we analyse the possible collections of ray graphs and transition functions between them which can occur in a thick end. In Section 7 we introduce the notion of tribes and of their concentration towards an end and begin building some tools for the main recursive construction, which is given in Section 8. We conclude with a discussion of the future outlook in Section 9.

3. Preliminaries

In this paper we will denote by \mathbb{N} the set of positive integers and by \mathbb{N}_0 the set of non-negative integers. In our graph theoretic notation we generally follow the textbook of Diestel [Die17]. For a graph $G = (V, E)$ and $W \subseteq V$ we write $G[W]$ for the induced subgraph of G on W . For two vertices v, w of a connected graph G , we write $\text{dist}(v, w)$ for the edge-length of a shortest v - w path. A path $P = v_0v_1 \dots v_n$ in a graph G is called a *bare path* if $d_G(v_i) = 2$ for all inner vertices v_i for $0 < i < n$.

3.1. Rays and ends

Definition 3.1 (Rays, double rays and initial vertices of rays). A one-way infinite path is called a *ray* and a two-way infinite path is called a *double ray*. For a ray R , let $\text{init}(R)$ denote the *initial vertex* of R , that is the unique vertex of degree 1 in R . For a family \mathcal{R} of rays, let $\text{init}(\mathcal{R})$ denote the set of initial vertices of the rays in \mathcal{R} .

Definition 3.2 (Tail of a ray). Given a ray R in a graph G and a finite set $X \subseteq V(G)$, the *tail of R after X* , written $T(R, X)$, is the unique infinite component of R in $G - X$.

Definition 3.3 (Concatenation of paths and rays). For a path or ray P and vertices $v, w \in V(P)$, let vPw denote the subpath of P with endvertices v and w , and $\overset{\circ}{v}P\overset{\circ}{w}$ the subpath strictly between v and w . If P is a ray, let Pv denote the finite subpath of P between the initial vertex of P and v , and let vP denote the subray (or *tail*) of P with initial vertex v . Similarly, we write $P\overset{\circ}{v}$ and $\overset{\circ}{v}P$ for the corresponding path/ray without the vertex v . For a ray $R = r_0r_1\dots$, let R^- denote the tail r_1R of R starting at r_1 . Given a family \mathcal{R} of rays, let \mathcal{R}^- denote the family $(R^- : R \in \mathcal{R})$.

Given two paths or rays P and Q , which intersect in a single vertex only, which is an endvertex in both P and Q , we write PQ for the *concatenation of P and Q* , that is the path, ray or double ray $P \cup Q$. Moreover, if we concatenate paths of the form vPw and wQx , then we omit writing w twice and denote the concatenation by $vPwQx$.

Definition 3.4 (Ends of a graph, cf. [Die17, Chapter 8]). An *end* of an infinite graph Γ is an equivalence class of rays, where two rays R and S of Γ are *equivalent* if and only if there are infinitely many vertex disjoint paths between R and S in Γ . We denote by $\Omega(\Gamma)$ the set of ends of Γ .

We say that a ray $R \subseteq \Gamma$ *converges* (or *tends*) to an end ϵ of Γ if R is contained in ϵ . In this case, we call R an ϵ -ray. Given an end $\epsilon \in \Omega(\Gamma)$ and a finite set $X \subseteq V(\Gamma)$ there is a unique component of $\Gamma - X$ which contains a tail of every ray in ϵ , which we denote by $C(X, \epsilon)$. Given two ends $\epsilon, \epsilon' \in \Omega(\Gamma)$, we say a finite set $X \subseteq V(\Gamma)$ *separates* ϵ and ϵ' if $C(X, \epsilon) \neq C(X, \epsilon')$.

For an end $\epsilon \in \Omega(\Gamma)$, we define the *degree* of ϵ in Γ , denoted by $\deg(\epsilon)$, as the supremum in $\mathbb{N} \cup \{\infty\}$ of the set $\{|\mathcal{R}| : \mathcal{R} \text{ is a set of disjoint } \epsilon\text{-rays}\}$. Note that this supremum is in fact an attained maximum, i.e. for each end ϵ of Γ there is a set \mathcal{R} of vertex-disjoint ϵ -rays with $|\mathcal{R}| = \deg(\epsilon)$, as proved by Halin [Hal65, Satz 1]. An end with finite degree is called *thin*, otherwise the end is called *thick*.

3.2. Inflated copies of graphs

Definition 3.5 (Inflated graph, branch set). Given a graph G , we say that a pair (H, φ) is an *inflated copy of G* , or an *IG*, if H is a graph and $\varphi: V(H) \rightarrow V(G)$ is a map such that:

- For every $v \in V(G)$ the *branch set* $\varphi^{-1}(v)$ induces a non-empty, connected subgraph of H ;
- There is an edge in H between $\varphi^{-1}(v)$ and $\varphi^{-1}(w)$ if and only if $vw \in E(G)$ and this edge, if it exists, is unique.

When there is no danger of confusion, we will simply say that H is an IG instead of saying that (H, φ) is an IG , and denote by $H(v) = \varphi^{-1}(v)$ the branch set of v .

Definition 3.6 (Minor). A graph G is a *minor* of another graph Γ , written $G \preceq \Gamma$, if there is some subgraph $H \subseteq \Gamma$ such that H is an inflated copy of G . In this case, we also say that H is a G -minor in Γ .

Definition 3.7 (Extension of inflated copies). Suppose $G \subseteq G'$ as subgraphs, and that H is an IG and H' is an IG' . We say that H' *extends* H (or that H' is an *extension* of H) if $H \subseteq H'$ as subgraphs and $H(v) \subseteq H'(v)$ for all $v \in V(G)$. Observe that, since $H \subseteq H'$, for every edge $vw \in E(G)$ the unique edge between the branch sets $H'(v)$ and $H'(w)$ is also the unique edge between $H(v)$ and $H(w)$.

If H' is an extension of H and $X \subseteq V(G)$ is such that $H'(x) = H(x)$ for every $x \in X$, then we say H' is an extension of H *fixing* X .

Definition 3.8 (Tidiness). Let (H, φ) be an IG . We call (H, φ) *tidy* if

- $H[\varphi^{-1}(v)]$ is a tree for all $v \in V(G)$;
- $H[\varphi^{-1}(v)]$ is finite if $d_G(v)$ is finite.

Note that every H which is an IG contains a subgraph H' such that $(H', \varphi \upharpoonright V(H'))$ is a tidy IG , although this choice may not be unique. In this paper we will always assume, without loss of generality, that each IG is tidy.

Definition 3.9 (Restriction). Let G be a graph, $M \subseteq G$ a subgraph of G , and let (H, φ) be an IG . The *restriction of H to M* , denoted by $H(M)$, is the IM given by $(H(M), \varphi')$, where $\varphi'^{-1}(v) = \varphi^{-1}(v)$ for all $v \in V(M)$, and $H(M)$ consists of the union of the subgraphs of H induced on each branch set $\varphi^{-1}(v)$ for each $v \in V(M)$, together with the edge between $\varphi^{-1}(u)$ and $\varphi^{-1}(v)$ in H for each $uv \in E(M)$.

Suppose R is a ray in some graph G . If H is a tidy IG in a graph Γ , then in the restriction $H(R)$ all rays which do not have a tail contained in some branch set will share a tail. Later in the paper, we will want to make this correspondence between rays in G and Γ more explicit, with use of the following definition:

Definition 3.10 (Pullback). Let G be a graph, $R \subseteq G$ a ray, and let (H, φ) be a tidy IG . The *pullback of R to H* is the subgraph $H^\downarrow(R) \subseteq H(R)$, where $H^\downarrow(R)$ is subgraph minimal such that $(H^\downarrow(R), \varphi \upharpoonright V(H^\downarrow(R)))$ is an IR .

Note that, since H is tidy, $H^\downarrow(R)$ is well defined. It can be shown that, in fact, $H^\downarrow(R)$ is also ray.

Lemma 3.11 ([BEE⁺23, Lemma 2.11]). *Let G be a graph and let H be a tidy IG . If $R \subseteq G$ is a ray, then the pullback $H^\downarrow(R)$ is also a ray.*

Definition 3.12. Let G be a graph, \mathcal{R} be a family of disjoint rays in G , and let H be a tidy IG . We will write $H^\downarrow(\mathcal{R})$ for the family $(H^\downarrow(R) : R \in \mathcal{R})$.

It is easy to check that if two rays R and S in G are equivalent, then also $H^\downarrow(R)$ and $H^\downarrow(S)$ are rays (Lemma 3.11) which are equivalent in H , and hence also equivalent in Γ .

Definition 3.13. For an end ω of G and $H \subseteq \Gamma$ a tidy IG , we denote by $H(\omega)$ the unique end of Γ containing all rays $H^\downarrow(R)$ for $R \in \omega$.

3.3. Transitional linkages and the strong linking lemma

The next definition is based on definitions already stated in Section 2 (cf. Definition 2.1, Definition 2.3 and Definition 2.4).

Definition 3.14. We say a linkage between two families of rays is *transitional* if the function which it induces between the corresponding ray graphs is a transition function.

Lemma 3.15. *Let Γ be a graph and let $\epsilon \in \Omega(\Gamma)$. Then, for any finite families $\mathcal{R} = (R_i : i \in I)$ and $\mathcal{S} = (S_j : j \in J)$ of disjoint ϵ -rays in Γ , there is a finite set X such that every linkage from \mathcal{R} to \mathcal{S} after X is transitional.*

Proof. By definition, for every function $\sigma : I \rightarrow J$ which is not a transition function from \mathcal{R} to \mathcal{S} there is a finite set $X_\sigma \subseteq V(\Gamma)$ such that there is no linkage from \mathcal{R} to \mathcal{S} after X_σ which induces σ . If we let Φ be the set of all such σ which are not transition functions, then the set $X := \bigcup_{\sigma \in \Phi} X_\sigma$ satisfies the conclusion of the lemma. \square

In addition to Lemma 2.2, we will also need the following stronger linking lemma, which is a slight modification of [BEE⁺22, Lemma 4.4]:

Lemma 3.16 (Strong linking lemma). *Let Γ be a graph and let $\epsilon \in \Omega(\Gamma)$. Let X be a finite set of vertices and let $\mathcal{R} = (R_i : i \in I)$ a finite family of vertex disjoint ϵ -rays. Let $x_i = \text{init}(R_i)$ and let $x'_i = \text{init}(T(R_i, X))$. Then there is a finite number $N = N(\mathcal{R}, X)$ with the following property: For every collection $(H_j : j \in [N])$ of vertex disjoint connected subgraphs of Γ , all disjoint from X and each including a specified ray S_j in ϵ , there is an $\ell \in [N]$ and a transitional linkage $\mathcal{P} = (P_i : i \in I)$ from \mathcal{R} to $(S_j : j \in [N])$, with transition function σ , which is after X and such that the family*

$$\mathcal{T} = (x_i R_i x'_i P_i y_{\sigma(i)} S_{\sigma(i)} : i \in I)$$

avoids H_ℓ .

Proof. Let $Y \subseteq V(\Gamma)$ be a finite set as in Lemma 3.15. We apply the strong linking lemma established in [BEE⁺22, Lemma 4.4] to the set $X \cup Y$ to obtain this version of the strong linking lemma. \square

Lemma and Definition 3.17. *Let Γ be a graph, $\epsilon \in \Omega(\Gamma)$, $X \subseteq V(\Gamma)$ be finite, and let $\mathcal{R} = (R_i : i \in I)$, $\mathcal{S} = (S_j : j \in J)$ be two finite families of disjoint ϵ -rays with $|I| \leq |J|$. Then there is a finite subgraph Y such that, for any transition function σ from \mathcal{R} to \mathcal{S} , there is a linkage \mathcal{P}_σ from \mathcal{R} to \mathcal{S} inducing σ , with $\bigcup \mathcal{P}_\sigma \subseteq Y$, which is after X .*

We call such a graph Y a transition box between \mathcal{R} and \mathcal{S} (after X).

Proof. Let $\sigma: I \rightarrow J$ be a transition function from \mathcal{R} to \mathcal{S} . By definition, there is a linkage \mathcal{P}_σ from \mathcal{R} to \mathcal{S} after X which induces σ . Let Φ be the set of all transition functions from \mathcal{R} to \mathcal{S} and let $Y = \bigcup_{\sigma \in \Phi} \mathcal{P}_\sigma$. Then Y is a transition box between \mathcal{R} and \mathcal{S} (after X). \square

Remark and Definition 3.18. Let Γ be a graph and $\epsilon \in \Omega(\Gamma)$. Let $\mathcal{R}_1, \mathcal{R}_2, \mathcal{R}_3$ be finite families of disjoint ϵ -rays, \mathcal{P}_1 a transitional linkage from \mathcal{R}_1 to \mathcal{R}_2 , and let \mathcal{P}_2 be a transitional linkage from \mathcal{R}_2 to \mathcal{R}_3 after $V(\bigcup \mathcal{P}_1)$. Then

1. \mathcal{P}_2 is also a transitional linkage from $(\mathcal{R}_1 \circ_{\mathcal{P}_1} \mathcal{R}_2)$ to \mathcal{R}_3 ;⁴
2. The linkage from \mathcal{R}_1 to \mathcal{R}_3 yielding the rays $(\mathcal{R}_1 \circ_{\mathcal{P}_1} \mathcal{R}_2) \circ_{\mathcal{P}_2} \mathcal{R}_3$, which we call the *concatenation* $\mathcal{P}_1 + \mathcal{P}_2$ of \mathcal{P}_1 and \mathcal{P}_2 , is transitional.

The following lemmas are simple exercises.

Lemma 3.19. *Let Γ be a graph and $(R_i: i \in I)$ be a finite family of equivalent disjoint rays. Then the ray graph $\text{RG}(R_i: i \in I)$ is connected. Also, if R'_i is a tail of R_i for each $i \in I$, then we have that $\text{RG}(R_i: i \in I) = \text{RG}(R'_i: i \in I)$. \square*

Lemma 3.20 ([BEE⁺23, Lemma 3.4]). *Let Γ be a graph, $\Gamma' \subseteq \Gamma$, $\mathcal{R} = (R_i: i \in I)$ be a finite family of disjoint rays in Γ' , and let $\mathcal{S} = (S_j: j \in J)$ be a finite family of disjoint rays in $\Gamma - V(\Gamma')$, where I and J are disjoint. Then $\text{RG}_{\Gamma'}(\mathcal{R})$ is a subgraph of $\text{RG}_\Gamma(\mathcal{R} \cup \mathcal{S})[I]$. \square*

3.4. Separations and tree-decompositions of graphs

Definition 3.21. Let $G = (V, E)$ be a graph. A *separation* of G is a pair (A, B) of subsets of vertices such that $A \cup B = V$ and such that there is no edge between $B \setminus A$ and $A \setminus B$. Given a separation (A, B) , we write $\overline{G[B]}$ for the graph obtained by deleting all edges in the *separator* $A \cap B$ from $G[B]$. Two separations (A, B) and (C, D) are *nested* if one of the following conditions hold:

$$\begin{array}{llll} A \subseteq C \text{ and } D \subseteq B, & \text{or} & B \subseteq C \text{ and } D \subseteq A, & \text{or} \\ A \subseteq D \text{ and } C \subseteq B, & \text{or} & B \subseteq D \text{ and } C \subseteq A. & \end{array}$$

Definition 3.22. Let T be a tree with a root $v \in V(T)$. Given nodes $x, y \in V(T)$, let us denote by xTy the unique path in T between x and y , by T_x denote the component of $T - E(vTx)$ containing x , and by $\overline{T_x}$ the tree $T - T_x$.

Given an edge $e = tt' \in E(T)$, we say that t is the *lower vertex* of e if $t \in vTt'$. In this case, we denote t by e^- , and call t' the *higher vertex* of e , denoted by e^+ .

If S is a subtree of a tree T , let us write $\partial(S) = E(S, T \setminus S)$ for the edge cut between S and its complement in T .

We say that S is a *initial subtree* of T if S contains v . In this case, we consider S to be rooted in v as well.

⁴Formally, it is only the subset of \mathcal{P}_2 starting at the rays in $(\mathcal{R}_1 \circ_{\mathcal{P}_1} \mathcal{R}_2)$ which is a linkage from $(\mathcal{R}_1 \circ_{\mathcal{P}_1} \mathcal{R}_2)$ to \mathcal{R}_3 . Here and later in the paper, we will use such abuses of notation, when the appropriate subset of the path family is clear from context.

A reader unfamiliar with tree-decompositions may also consult [Die17, Chapter 12.3].

Definition 3.23 (Tree-decomposition). Given a graph $G = (V, E)$, a *tree-decomposition* of G is a pair (T, \mathcal{V}) consisting of a rooted tree T , together with a family of subsets of vertices $\mathcal{V} = (V_t : t \in V(T))$, such that:

- $V(G) = \bigcup \mathcal{V}$;
- For every edge $e \in E(G)$ there is a $t \in V(T)$ such that e lies in $G[V_t]$;
- $V_{t_1} \cap V_{t_3} \subseteq V_{t_2}$ whenever $t_2 \in V(t_1 T t_3)$.

The vertex sets V_t for $t \in V(T)$ are called the *parts* of the tree-decomposition (T, \mathcal{V}) .

Definition 3.24 (Tree-width). Suppose (T, \mathcal{V}) is a tree-decomposition of a graph G . The *width* of (T, \mathcal{V}) is the number $\sup \{|V_t| - 1 : t \in V(T)\} \in \mathbb{N}_0 \cup \{\infty\}$. The *tree-width* of a graph G is the least width of any tree-decomposition of G .

4. Extensive tree-decompositions and self minors

The purpose of this section is to explain the extensive tree-decompositions mentioned in the proof sketch. Some ideas motivating this definition are already present in Andreae's proof that locally finite trees are ubiquitous under the topological minor relation [And79, Lemma 2].

4.1. Extensive tree-decompositions

Definition 4.1 (Separations induced by tree-decompositions). For a tree-decomposition (T, \mathcal{V}) of a graph G , and an edge $e \in E(T)$, let

- $A(e) := \bigcup \{V_{t'} : t' \notin V(T_{e+})\}$;
- $B(e) := \bigcup \{V_{t'} : t' \in V(T_{e+})\}$;
- $S(e) := A(e) \cap B(e) = V_{e-} \cap V_{e+}$.

Then $(A(e), B(e))$ is a separation of G (cf. [Die17, Chapter 12.3.1]). We call $B(e)$ the *bough* of (T, \mathcal{V}) rooted in e and $S(e)$ the *separator* of $B(e)$. When writing $\overline{G[B(e)]}$ it is implicitly understood that this refers to the separation $(A(e), B(e))$ (cf. Definition 3.21.)

Definition 4.2. Let (T, \mathcal{V}) be a tree-decomposition of a graph G . For a subtree $S \subseteq T$, let us write

$$G(S) = G \left[\bigcup_{t \in V(S)} V_t \right]$$

and, if H is an IG , we write $H(S) = H(G(S))$ for the restriction of H to $G(S)$.

Definition 4.3 (Self-similar bough). Let (T, \mathcal{V}) be a tree-decomposition of a graph G . Given an edge $e \in E(T)$, the bough $B(e)$ is called *self-similar* (towards an end ω of G), if there is a family $\mathcal{R}_e = (R_{e,s} : s \in S(e))$ of disjoint ω -rays in G such that for all $n \in \mathbb{N}$ there is an edge $e' \in E(T_{e^+})$ with $\text{dist}(e^-, e'^-) \geq n$ such that

- for each $s \in S(e)$, the ray $R_{e,s}$ starts in s and meets $S(e')$;
- there is a subgraph $W \subseteq G[B(e')]$ which is an inflated copy of $\overline{G[B(e)]}$;
- for each $s \in S(e)$, we have $V(R_{e,s}) \cap S(e') \subseteq W(s)$.

Such a W is called a *witness for the self-similarity of $B(e)$ (towards an end ω of G) of distance at least n* .

Definition 4.4 (Extensive tree-decomposition). A tree-decomposition (T, \mathcal{V}) of G is *extensive* if

- T is a locally finite, rooted tree;
- each part of (T, \mathcal{V}) is finite;
- every vertex of G appears in only finitely many parts of \mathcal{V} ;
- for each $e \in E(T)$, the bough $B(e)$ is self-similar towards some end ω_e of G .

Remark 4.5. If (T, \mathcal{V}) is extensive then, for each edge $e \in E(T)$ and every $n \in \mathbb{N}$, there is an edge $e' \in E(T_{e^+})$ with $\text{dist}(e^-, e'^-) \geq n$, such that $G[B(e')]$ contains a witness for the self-similarity of $B(e)$. Since T is locally finite, there is some ray R_e in T such that there are infinitely many such e' on R_e .

The following is the main result of this paper.

Theorem 4.6. *Every locally finite connected graph admitting an extensive tree-decomposition is \preceq -ubiquitous.*

4.2. Self minors in extensive tree-decompositions

The existence of an extensive tree-decomposition of a graph G will imply the existence of many self-minors of G , which will be essential to our proof.

Throughout this subsection, let G denote a locally finite, connected graph with an extensive tree-decomposition (T, \mathcal{V}) .

Definition 4.7. Let (A, B) be a separation of G with $A \cap B = \{v_1, v_2, \dots, v_n\}$. Suppose H_1, H_2 are subgraphs of a graph Γ , where H_1 is an inflated copy of $G[A]$, H_2 is an inflated copy of $G[B]$, and for all vertices $x \in A$ and $y \in B$, we have $H_1(x) \cap H_2(y) \neq \emptyset$ only if $x = y = v_i$ for some i . Suppose further that \mathcal{P} is a family of disjoint paths $(P_i : i \in [n])$ in Γ such that each P_i is a path from $H_1(v_i)$ to $H_2(v_i)$, which is otherwise disjoint from $H_1 \cup H_2$. Note that P_i may be a single vertex if $H_1(v_i) \cap H_2(v_i) \neq \emptyset$.

We write $H_1 \oplus_{\mathcal{P}} H_2$ for the IG given by (H, ϕ) , where $H = H_1 \cup H_2 \cup \bigcup_{i \in [n]} P_i$ and

$$H(v) := \begin{cases} H_1(v_i) \cup V(P_i) \cup H_2(v_i) & \text{if } v = v_i \in A \cap B, \\ H_1(v) & \text{if } v \in A \setminus B, \\ H_2(v) & \text{if } v \in B \setminus A. \end{cases}$$

We note that this may produce a non-tidy IG , in which case in practice (in order to maintain our assumption that each IG we consider is tidy) we will always delete some edges inside the branch sets to make it tidy.

We will often use this construction when the family \mathcal{P} consists of certain segments of a family of disjoint rays \mathcal{R} . If \mathcal{R} is such that each R_i has its first vertex in $H_1(v_i)$ and is otherwise disjoint from H_1 , and such that every R_i meets H_2 , and does so first in some vertex $x_i \in H(v_i)$, then we write

$$H_1 \oplus_{\mathcal{R}} H_2 = H_1 \oplus_{(R_i x_i : i \in [n])} H_2.$$

Definition 4.8 (Push-out). A self minor $G' \subseteq G$ (meaning G' is an IG) is called a *push-out of G along e to depth n* for some $e \in E(T)$ if there is an edge $e' \in T_{e^+}$ such that $\text{dist}(e^-, e'^-) \geq n$ and an inflated copy $W \subseteq G[B(e')]$ of $\overline{G[B(e)]}$, such that $G' = G[A(e)] \oplus_{\mathcal{R}_e} W$.

Similarly, if H is an IG , then a subgraph H' of H is a *push-out of H along e to depth n* for some $e \in E(T)$ if there is an edge $e' \in T_{e^+}$ such that $\text{dist}(e^-, e'^-) \geq n$ and a subgraph $W \subseteq H(G[B(e')])$, which is an inflated copy of $\overline{G[B(e)]}$, such that

$$H' = H(G[A(e)]) \oplus_{H^\downarrow(\mathcal{R}_e)} W.$$

Note that if G' is a push-out of G along e to depth n , then $H(G')$ has a subgraph which is a push-out of H along e to depth n .

Lemma 4.9. *For each $e \in E(T)$, each $n \in \mathbb{N}$, and each witness W of the self-similarity of $B(e)$ of distance at least n there is a corresponding push-out $G_W := G[A(e)] \oplus_{\mathcal{R}_e} W$ of G along e to depth n .*

Proof. Let $e' \in E(T_{e^+})$ be the edge in Definition 4.3 such that $W \subseteq G[B(e')]$. By Definition 4.3, each ray $R_{e,s}$ meets $S(e')$ and $R_{e,s} \cap S(e') \subseteq W(s)$. Hence, the initial segment of $R_{e,s}$ up to the first point in W only meets $G[A(e)] \cup W$ in $\{s\} \cup W(s)$. Now, if $s' \in S(e) \cap W(s)$ for some s' , then $s' \in S(e')$, and so $R_{e,s'} \cap S(e') \not\subseteq W(s')$, contradicting Definition 4.3.

Since $G[A(e)]$ is an $IG[A(e)]$ and W is an inflated copy of $\overline{G[B(e)]}$, by Definitions 4.7 and 4.8 $G[A(e)] \oplus_{\mathcal{R}_e} W$ is well-defined and is indeed a push-out of G along e to depth n . \square

The existence of push-outs of G along e to arbitrary depths is in some sense the essence of extensive tree-decompositions, and lies at the heart of our inductive construction in Section 8.

5. Existence of extensive tree-decompositions

The purpose of this section is to examine two classes of locally finite connected graphs that have extensive tree-decompositions: Firstly, the class of graphs with finitely many ends, all of which are thin, and secondly the class of graphs of finite tree-width. In both cases we will show the existence of extensive tree-decompositions using some results about the *well-quasi-ordering* of certain classes of graphs.

A *quasi-order* is a reflexive and transitive binary relation, such as the minor relation between graphs. A quasi-order \preceq on a set X is a *well-quasi-order* if for all infinite sequences $(x_i)_{i \in \mathbb{N}}$ with $x_i \in X$ for every $i \in \mathbb{N}$ there exist $i, j \in \mathbb{N}$ with $i < j$ such that $x_i \preceq x_j$. The following two consequences will be useful.

Remark 5.1. A simple Ramsey type argument shows that if \preceq is a well-quasi-order on X , then every infinite sequence $(x_i)_{i \in \mathbb{N}}$ with $x_i \in X$ for every $i \in \mathbb{N}$ contains an increasing infinite subsequence $x_{i_1}, x_{i_2}, \dots \in X$. That is, an increasing infinite sequence $i_1 < i_2 < \dots$ such that $x_{i_j} \preceq x_{i_k}$ for all $j < k$.

Also, it is simple to show that if \preceq is a well-quasi-order on X , then for every infinite sequence $(x_i)_{i \in \mathbb{N}}$ with $x_i \in X$ for every $i \in \mathbb{N}$ there is an $i_0 \in \mathbb{N}$ such that for every $i \geq i_0$ there are infinitely many $j \in \mathbb{N}$ with $x_i \preceq x_j$.

A famous result of Robertson and Seymour [RS04], proved over a series of 20 papers, shows that finite graphs are well-quasi-ordered under the minor relation. Thomas [Tho89] showed that for any $k \in \mathbb{N}$ the class of graphs with tree-width at most k and arbitrary cardinality is well-quasi-ordered by the minor relation.

We will use slight strengthenings of both of these results, Lemma 5.3 and Lemma 5.11, to show that our two classes of graphs admit extensive tree-decompositions.

In Section 9 we will discuss in more detail the connection between our proof and well-quasi-orderings, and indicate how stronger well-quasi-ordering results could be used to prove the ubiquity of larger classes of graphs.

5.1. Finitely many thin ends

We will consider the following strengthening of the minor relation.

Definition 5.2. Given $\ell \in \mathbb{N}$, an ℓ -pointed graph is a graph G together with a function $\pi: [\ell] \rightarrow V(G)$, called a *point function*. For ℓ -pointed graphs (G_1, π_1) and (G_2, π_2) , we say $(G_1, \pi_1) \preceq_p (G_2, \pi_2)$ if $G_1 \preceq G_2$ and this can be arranged in such a way that $\pi_2(i)$ is contained in the branch set of $\pi_1(i)$ for every $i \in [\ell]$.

Lemma 5.3. For $\ell \in \mathbb{N}$ the set of ℓ -pointed finite graphs is well-quasi-ordered under the relation \preceq_p .

Proof. This follows from a stronger statement of Robertson and Seymour in [RS10, 1.7]. □

We will also need the following structural characterisation of locally finite one-ended graphs with a thin end due to Halin.

Lemma 5.4 ([Hal65, Satz 3']). *Every one-ended, locally finite connected graph G with a thin end of degree $k \in \mathbb{N}$ has a tree-decomposition (R, \mathcal{V}) of G such that $R = t_0 t_1 t_2 \dots$ is a ray, and for every $i \in \mathbb{N}_0$:*

- $|V_{t_i}|$ is finite;
- $|S(t_i t_{i+1})| = k$;
- $S(t_i t_{i+1}) \cap S(t_{i+1} t_{i+2}) = \emptyset$.

Remark 5.5. Note that in the above lemma, for a given finite set $X \subseteq V(G)$, by taking the union over parts corresponding to an initial segment of the ray of the decomposition, one may always assume that $X \subseteq V_{t_0}$. Moreover, note that since $S(t_i t_{i+1}) \cap S(t_{i+1} t_{i+2}) = \emptyset$, it follows that every vertex of G is contained in at most two parts of the tree-decomposition.

Lemma 5.6. *Every one-ended, locally finite connected graph G with a thin end has an extensive tree-decomposition (R, \mathcal{V}) where $R = t_0 t_1 t_2 \dots$ is a ray rooted in its initial vertex.*

Proof. Let $k \in \mathbb{N}$ be the degree of the thin end of G and let $\mathcal{R} = (R_j : j \in [k])$ be a maximal family of disjoint rays in G . Let (R', \mathcal{W}) be the tree-decomposition of G given by Lemma 5.4 where $R' = t'_0 t'_1 \dots$.

By Remark 5.5 (and considering tails of rays if necessary), we may assume that each ray in \mathcal{R} starts in $S(t'_0 t'_1)$. Note that each ray in \mathcal{R} meets the separator $S(t'_{i-1} t'_i)$ for each $i \in \mathbb{N}$. Since \mathcal{R} is a family of k disjoint rays and $|S(t'_{i-1} t'_i)| = k$ for each $i \in \mathbb{N}$, each vertex in $S(t'_{i-1} t'_i)$ is contained in a unique ray in \mathcal{R} .

Let $\ell = 2k$ and consider a sequence $(G_i, \pi_i)_{i \in \mathbb{N}}$ of ℓ -pointed finite graphs defined by $G_i := G[W_{t'_i}]$ and

$$\pi_i : [\ell] \rightarrow V(G_i), j \mapsto \begin{cases} \text{the unique vertex in } S(t'_{i-1} t'_i) \cap V(R_j) & \text{for } 1 \leq j \leq k, \\ \text{the unique vertex in } S(t'_i t'_{i+1}) \cap V(R_{j-k}) & \text{for } k < j \leq 2k = \ell. \end{cases}$$

By Lemma 5.3 and Remark 5.1 there is an $n_0 \in \mathbb{N}$ such that for every $n \geq n_0$ there are infinitely many $m > n$ with $(G_n, \pi_n) \preceq_p (G_m, \pi_m)$.

Let $V_{t_0} := \bigcup_{i=0}^{n_0} W_{t'_i}$ and $V_{t_i} := W_{t'_{n_0+i}}$ for all $i \in \mathbb{N}$. We claim that $(R, (V_{t_i} : i \in \mathbb{N}_0))$ is the desired extensive tree-decomposition of G where $R = t_0 t_1 t_2 \dots$ is a ray with root t_0 . The ray R is a locally finite tree and all the parts are finite. Moreover, every vertex of G is contained in at most two parts by Remark 5.5. It remains to show that for every $i \in \mathbb{N}$, the bough $B(t_{i-1} t_i)$ is self-similar towards the end of G .

Let $e = t_{i-1} t_i$ for some $i \in \mathbb{N}$. For each $s \in S(e)$, we let $p(s) \in [k]$ be such that $s \in R_{p(s)}$ and set $R_{e,s} = s R_{p(s)}$. We wish to show there is a witness W for the self-similarity of $B(e)$ of distance at least n for each $n \in \mathbb{N}$. Note that $B(e) = \bigcup_{j \geq 0} V(G_{n_0+i+j})$. By the choice of n_0 in Remark 5.1, there exists an $m > i + n$ such that $(G_{n_0+i}, \pi_{n_0+i}) \preceq_p (G_{n_0+m}, \pi_{n_0+m})$. Let $e' = t_{m-1} t_m$. We will show that there exists a $W \subseteq G[B(e')]$ witnessing the self-similarity of $B(e)$ towards the end of G .

Recursively, for each $j \geq 0$ we can find $m = m_0 < m_1 < m_2 < \dots$ with

$$(G_{n_0+i+j}, \pi_{n_0+i+j}) \preceq_p (G_{n_0+m_j}, \pi_{n_0+m_j}).$$

In particular, there are subgraphs $H_{m_j} \subseteq G_{n_0+m_j}$ which are inflated copies of G_{n_0+i+j} , all compatible with the point functions, and so

$$S(t'_{n_0+m_{j-1}} t'_{n_0+m_j}) \cup S(t'_{n_0+m_j} t'_{n_0+m_{j+1}}) \subseteq H_{m_j}$$

for each $j \geq 0$.

Hence, for every $j \in \mathbb{N}$ and $p \in [k]$ there is a unique $H_{m_{j-1}}-H_{m_j}$ subpath $P_{p,j}$ of R_p . We claim that

$$W' := \bigcup_{j \geq 0} H_{m_j} \cup \bigcup_{j \in \mathbb{N}} \bigcup_{p \in [k]} P_{p,j}$$

is a subgraph of $G[B(e')]$ that is an $IG[B(e)]$. Hence, the desired W can be obtained as a subgraph of W' .

To prove this claim it is sufficient to check that for each $j \in \mathbb{N}$ and each $s \in S(t_{j-1}t_j)$, the branch sets of s in H_{j-1} and in H_j are connected by $P_{p(s),j}$. Indeed, by construction, every $P_{p,j}$ is a path from $\pi_{n_0+m_{j-1}}(k+p)$ to $\pi_{n_0+m_j}(p)$. And, since the H_{m_j} are pointed minors of $G_{n_0+m_j}$, it follows that $\pi_{n_0+m_{j-1}}(k+p(s)) \in H_{m_{j-1}}(s)$ and $\pi_{n_0+m_j}(p(s)) \in H_{m_j}(s)$ are as desired.

Finally, since $(G_{n_0+i}, \pi_{n_0+i}) \preceq_p (G_{n_0+m}, \pi_{n_0+m})$ as witnessed by H_{m_0} , the branch set of each $s \in S(t_{i-1}t_i)$ must indeed include $V(R_{e,s}) \cap S(e')$. \square

Lemma 5.7. *If G is a locally finite connected graph with finitely many ends, each of which is thin, then G has an extensive tree-decomposition.*

Proof. Let $\Omega(G) = \{\omega_1, \dots, \omega_n\}$ be the set of the ends of G . Let $X \subseteq V(G)$ be a finite set of vertices which separates the ends of G , i.e. so that all $C_i = C(X, \omega_i)$ are pairwise disjoint. Without loss of generality, we may assume that $V(G) = X \cup \bigcup_{i \in [n]} C_i$.

Let $G_i := G[C_i \cup X]$. Then each G_i is a locally finite connected one-ended graph, with a thin end ω_i , and hence by Lemma 5.6 each of the G_i admits an extensive tree-decomposition (R^i, \mathcal{V}^i) , where R^i is rooted in its initial vertex r^i . Without loss of generality, $X \subseteq V_{r^i}^i$ for each $i \in [n]$.

Let T be the tree formed by identifying the family of rays $(R^i : i \in [n])$ at their roots, let r be this identified vertex which we consider to be the root of T , and let (T, \mathcal{V}) be the tree-decompositions whose root part is $\bigcup_{i \in [n]} V_{r^i}^i$, and which otherwise agrees with the (R^i, \mathcal{V}^i) . It is a simple check that (T, \mathcal{V}) is an extensive tree-decomposition of G . \square

5.2. Finite tree-width

Definition 5.8. A rooted tree-decomposition (T, \mathcal{V}) of G is *lean* if for any $k \in \mathbb{N}$, any nodes $t_1, t_2 \in V(T)$, and any $X_1 \subseteq V_{t_1}, X_2 \subseteq V_{t_2}$ such that $|X_1|, |X_2| \geq k$ there are either k disjoint paths in G between X_1 and X_2 , or there is a vertex t on the path in T between t_1 and t_2 such that $|V_t| < k$.

Remark 5.9. Kříž and Thomas [KT91] showed that if G has tree-width at most m for some $m \in \mathbb{N}$, then G has a lean tree-decomposition of width at most m .

Lemma 5.10. *Let G be a locally finite connected graph and let (T, \mathcal{V}) be a lean tree-decomposition of G of width at most m . Then there exists a lean tree-decomposition of G of width at most m such that every bough is connected and the decomposition tree is locally finite. Moreover, we may assume that every vertex appears in only finitely many parts.*

Proof. We begin by defining the underlying tree T' of this decomposition. The root of T' will be the root r of T , and the other vertices will be pairs (e, C) where e is an edge of T and C is a component of $G - S(e)$ meeting (or equivalently, included in) $B(e)$. There is an edge from r to (e, C) whenever $e^- = r$, and from (e, C) to (f, D) whenever $f^- = e^+$ and $D \subseteq C$. For future reference, we define a graph homomorphism π from T' to T by setting $\pi(r) = r$ and $\pi(e, C) = e^+$. Next, we set $V'_r := V_r$ and

$$V'_{(e,C)} := V_{e^+} \cap (V(C) \cup N(V(C))),$$

where $N(V(C))$ is the neighbourhood of $V(C)$. Moreover, we let \mathcal{V}' denote the family of all V'_p for all nodes p of T' .

To see that T' is locally finite, note that for any child (e, C) of p the set C is also a component of $G \setminus V_{\pi(p)}$ and that no two distinct children yield the same component; if (e, C) and (f, C) were distinct children of p , then we would have $V(C) \subseteq B(f) \subseteq A(e)$ and so

$$V(C) \subseteq A(e) \cap B(e) = S(e),$$

which is impossible.

We now analyse, for a given vertex v of G , which of the sets V'_p contain v . Since (T, \mathcal{V}) is a tree-decomposition, T induces a subtree on the set of nodes t of T with $v \in V_t$, and so this set has a minimal element t_v in the tree order. We set $p_v := r$ if $t_v = r$ and otherwise set $p_v := (e, C)$, where e is the unique edge of T with $e^+ = t_v$ and C is the unique component of $G - S(e)$ containing v . This guarantees that $v \in V'_{p_v}$. For any other node p of T' with $v \in V'_p$, we have $p \neq r$ and so p has the form (e, C) . Since $v \in V_{e^+}$ and $p \neq p_v$, it follows that e^- lies on the path from t_v to e^+ and so $v \in V_{e^-}$, from which $v \in N(V(C))$ follows. Thus, some neighbour w of v lies in C . Then $w \in B(e) \setminus S(e) = B(e) \setminus A(e)$ and so t_w lies in T_{e^+} . That is, p lies on the path from p_v to p_w . Conversely, for any $p = (e, C)$ on this path we have $w \in V(C)$ and so $v \in N(V(C)) \subseteq S(e) \subseteq V_{e^+}$, so that $v \in V'_p$.

What we have shown is that v is in V'_p precisely when $p = p_v$ or there is some neighbour w of v in G such that p lies on the path in T' from p_v to $p_w \in V(T'_{p_v})$. Using this information, it is easy to deduce that (T', \mathcal{V}') is a tree-decomposition: A vertex v is in V'_{p_v} and an edge vw with p_v no higher (in the tree order) than p_w in T is also in V'_{p_v} . The third condition in the definition of tree-decompositions follows from the fact that the T' induces a subtree on the set of all nodes p with $v \in V'_p$. These sets are also all finite, since G is locally finite.

Next we examine the boughs of this decomposition. Let $f \in E(T')$ with $f^+ = (e, C)$. Our aim is to show that $B(f) = V(C) \cup N(V(C))$. For any $(e', C') \in V(T'_{f^+})$, we have $V'_{(e',C')} \subseteq V(C') \cup N(V(C')) \subseteq V(C) \cup N(V(C))$, so that $B(f) \subseteq V(C) \cup N(V(C))$. For $v \in V(C)$, we have $p_v \in V(T_{f^+})$ and so $v \in B(f)$ and for $v \in N(V(C))$, there is a neighbour w of v such that f^+ lies on the path from p_v to p_w , yielding once more that $v \in B(f)$. This completes the proof that $B(f) = V(C) \cup N(V(C))$, and in particular $B(f)$ is connected.

Since G is locally finite, for each e , there are only finitely many components of $G - V_{e^-}$, so that T' is also locally finite. The final thing to show is that this decomposition is lean. So, suppose we have $X_1 \subseteq V'_{p_1}$ and $X_2 \subseteq V'_{p_2}$ with $|X_1|, |X_2| \geq k$. Then also $X_1 \subseteq V_{\pi(p_1)}$ and $X_2 \subseteq V_{\pi(p_2)}$, so that if there are no k disjoint paths from X_1 to X_2 in G , then there is some t on the path from $\pi(p_1)$ to $\pi(p_2)$ in T with $|V_t| \leq k$. But then there is some p on the path from p_1 to p_2 in T' with $\pi(p) = t$ and, since $V'_p \subseteq V_t$, we have $|V'_p| \leq k$. \square

Lemma 5.11. *For all $k, \ell \in \mathbb{N}$, the class of ℓ -pointed graphs with tree-width at most k is well-quasi-ordered under the relation \preceq_p .*

Proof. This is a consequence of a result of Thomas [Tho89]. \square

Lemma 5.12. *Every locally finite connected graph of finite tree-width has an extensive tree-decomposition.*

Proof. Let G be a locally finite connected graph of tree-width $m \in \mathbb{N}$. By Remark 5.9, G has a lean tree-decomposition of width at most m and so, by Lemma 5.10, there is a lean tree-decomposition (T, \mathcal{V}) of G with width m in which every bough is connected, every vertex is contained in only finitely many parts, and such that T is a locally finite tree with root r .

Let ϵ be an end of T and let R be the unique ϵ -ray starting at the root of T . Let $d_\epsilon = \liminf_{e \in R} |S(e)|$ and fix a tail $t_0^\epsilon t_1^\epsilon \dots$ of R such that $|S(t_{i-1}^\epsilon t_i^\epsilon)| \geq d_\epsilon$ for all $i \in \mathbb{N}$. Note that, $|S(t_{i_k-1}^\epsilon t_{i_k}^\epsilon)| = d_\epsilon$ for an infinite sequence $i_1 < i_2 < \dots$ of indices.

Since (T, \mathcal{V}) is lean, there are d_ϵ disjoint paths between $S(t_{i_k-1}^\epsilon t_{i_k}^\epsilon)$ and $S(t_{i_{k+1}-1}^\epsilon t_{i_{k+1}}^\epsilon)$ for every $k \in \mathbb{N}$. Moreover, since each $S(t_{i_k-1}^\epsilon t_{i_k}^\epsilon)$ is a separator of size d_ϵ , these paths are all internally disjoint. Hence, since every vertex appears in only finitely many parts, by concatenating these paths we get a family of d_ϵ many disjoint rays in G .

Fix one such family of rays $(R_j^\epsilon : j \in [d_\epsilon])$. We claim that there is an end ω of G such that $R_j^\epsilon \in \omega$ for all $j \in [d_\epsilon]$. Indeed, if not, then there is a finite vertex set X separating some pair of rays R and R' from the family. However, since each vertex appears in only finitely many parts, there is some $k \in \mathbb{N}$ such that $X \cap V_t = \emptyset$ for all $t \in V(T_{t_{i_k-1}^\epsilon}^\epsilon)$. By construction, R and R' have tails in $B(t_{i_k-1}^\epsilon t_{i_k}^\epsilon)$, which is connected and disjoint from X , contradicting the fact that X separates R and R' .

For every $k \in \mathbb{N}$, we define a point function $\pi_{i_k}^\epsilon : [d_\epsilon] \rightarrow S(t_{i_k-1}^\epsilon t_{i_k}^\epsilon)$ by letting $\pi_{i_k}^\epsilon(j)$ be the unique vertex in $V(R_j^\epsilon) \cap S(t_{i_k-1}^\epsilon t_{i_k}^\epsilon)$.

By Lemma 5.11 and Remark 5.1, the sequence $(G[B(t_{i_k-1}^\epsilon t_{i_k}^\epsilon)], \pi_{i_k}^\epsilon)_{k \in \mathbb{N}}$ has an increasing subsequence $(G[B(t_{i_j-1}^\epsilon t_{i_j}^\epsilon)], \pi_{i_j}^\epsilon)_{j \in I_\epsilon}$, i.e. there exists an $I_\epsilon \subseteq \{i_k : k \in \mathbb{N}\}$ such that for any $k, j \in I_\epsilon$ with $k < j$, we have

$$(G[B(t_{k-1}^\epsilon t_k^\epsilon)], \pi_k^\epsilon) \preceq_p (G[B(t_{j-1}^\epsilon t_j^\epsilon)], \pi_j^\epsilon).$$

Let us define $F_\epsilon = \{t_{k-1}^\epsilon t_k^\epsilon : k \in I_\epsilon\} \subseteq E(T)$.

Consider $T^- = T - \bigcup_{\epsilon \in \Omega(T)} F_\epsilon$, and let us write $\mathcal{C}(T^-)$ for the components of T^- . We claim that every component $C \in \mathcal{C}(T^-)$ is a locally finite rayless tree, and hence finite. Indeed, if C contains a ray $R \subseteq T$, then R is in an end ϵ of T and hence $F_\epsilon \cap R \neq \emptyset$, a contradiction. Consequently, each set $\bigcup_{t \in C} V_t$ is finite.

Let us define a tree-decomposition (T', \mathcal{V}') of G with $T' = T/\mathcal{C}(T^-)$, that is where we contract each component $C \in \mathcal{C}(T^-)$ to a single vertex and where $V_{t'}' = \bigcup_{t \in t'} V_t$. We claim this is an extensive tree-decomposition.

Clearly T' is a locally finite tree, each part of (T', \mathcal{V}') is finite, and every vertex of G is contained in only finitely many parts of the tree-decomposition. Given $e \in E(T')$, there is some $\epsilon \in \Omega(T)$ such that $e \in F_\epsilon$. Consider the family of rays $(R_{e,j}: j \in [d_\epsilon])$ given by $R_{e,j} = R_j^\epsilon \cap B(e)$. Let ω_e be the end of G in which the rays $R_{e,j}$ lie.

There is some $k \in \mathbb{N}$ such that $e = t_{k-1}^\epsilon t_k^\epsilon$. Given $n \in \mathbb{N}$, let $k' \in I_\epsilon$ be such that there are at least n indices $\ell \in I_\epsilon$ with $k < \ell < k'$, and let $e' = t_{k'-1}^\epsilon t_{k'}^\epsilon$. Note that, $e' \in F_\epsilon$ and hence $e' \in E(T')$. Furthermore, by construction e'^- has distance at least n from e^- in T' . Then, since $G[B(e)] = G[B(t_{k-1}^\epsilon t_k^\epsilon)]$ and $G[B(e')] = G[B(t_{k'-1}^\epsilon t_{k'}^\epsilon)]$, it follows that $(G[B(e)], \pi_k^\epsilon) \preceq_p (G[B(e')], \pi_{k'}^\epsilon)$, and so suitable subgraphs witness the self-similarity of $B(e)$ towards ω_e with the rays $(R_{e,j}: j \in [d_\epsilon])$, as in Lemma 5.6. \square

Remark 5.13. If for every $\ell \in \mathbb{N}$ the class of ℓ -pointed locally finite graphs without thick ends is well-quasi-ordered under \preceq_p , then every locally finite graph without thick ends has an extensive tree-decomposition. This follows by a simple adaptation of the proof above.

5.3. Sporadic examples

We note that, whilst Lemmas 5.7 and 5.12 show that a large class of locally finite graphs have extensive tree-decompositions, for many other graphs it is possible to construct an extensive tree-decomposition ‘by hand’. In particular, the fact that no graph in these classes has a thick end is an artefact of the method of proof, rather than a necessary condition for the existence of such a tree-decomposition, as is demonstrated by the following examples:

Remark 5.14. The grid $\mathbb{Z} \times \mathbb{Z}$ has an extensive tree-decomposition, which can be seen in Figure 5.1. More explicitly, we can take a ray decomposition of the grid given by a sequence of increasing diamond shaped regions around the origin. It is easy to check that every bough is self-similar towards the end of the grid.

A similar argument shows that the half-grid has an extensive tree-decomposition. However, we note that both of these graphs were already shown to be ubiquitous in [BEE⁺23].

6. The structure of non-pebbly ends

We will need a structural understanding of how the arbitrarily large families of IGs (for some fixed graph G) can be arranged inside some host graph Γ . In particular, we are interested in how the rays of these minors occupy a given end ϵ of Γ . In [BEE⁺23], by considering a pebble pushing game played on ray graphs, we established a distinction between *pebbly* and *non-pebbly* ends. Furthermore, we showed that each non-pebbly end is either *grid-like* or *half-grid-like*.

Theorem 6.1 ([BEE⁺23, Theorem 1.2]). *Let Γ be a graph and let ϵ be a thick end of Γ . Then ϵ is either pebbly, half-grid-like or grid-like.*

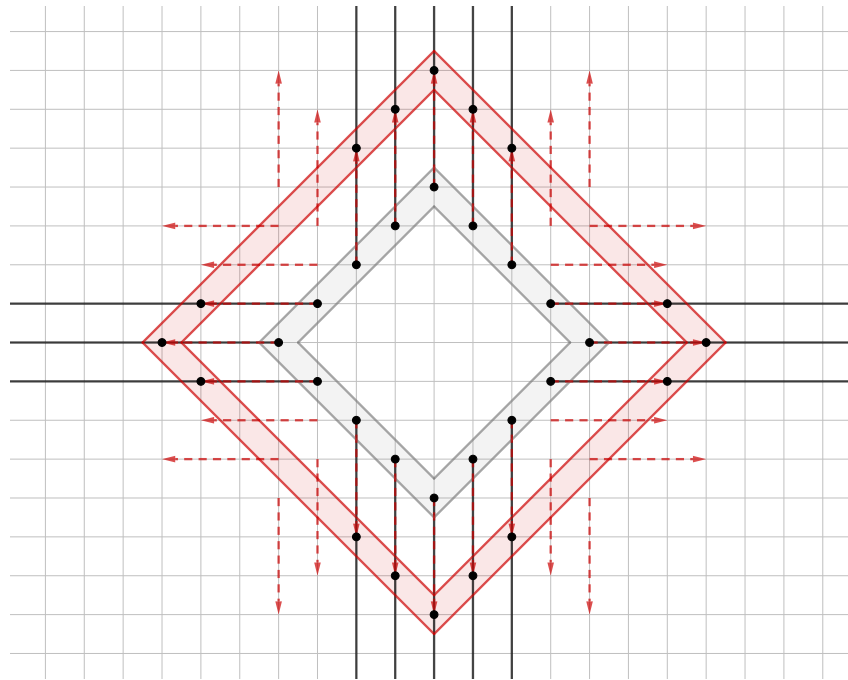


Figure 5.1: In the grid the boughs are self-similar.

Whilst the precise technical definition of such ends will not be important, one can think of this trichotomy roughly in the following way: Pebbly ends are ends which contain an infinite clique minor, and if we exclude such a minor, then it turns out the end has in some sense a ‘planar’ structure and must in some sense resemble a grid or half-grid in terms of how the rays in this end are arranged in the graph. In particular, finite families of rays in a grid-like end have some natural cyclic ordering, as they do in the grid, whereas finite families of rays in a half-grid-like end have some natural linear ordering, as they do in the half-grid, except perhaps for a small ‘exceptional’ set of rays on either end of this ordering. One can think of the structure of the graph towards a grid-like end as resembling a planar graph which doesn’t embed in the half-grid, whereas the structure towards a half-grid-like end resembles a planar graph which does embed in the half-grid, with perhaps some ‘local’ non-planar part which is attached to the x -axis, akin to an infinite vortex of finite depth, where in both cases there might then be ‘local’ non-planar parts attached to 3-separators. In what follows we will simply need to use the following results from [BEE+23].

Corollary 6.2 ([BEE+23, Corollary 5.3]). *Let Γ be a graph with a pebbly end ϵ and let G be a countable graph. Then $\aleph_0 G \preceq \Gamma$.*

Lemma 6.3 ([BEE+23, Lemma 7.1 and Corollary 7.3]). *Let Γ be a graph with a grid-like end ϵ . Then there exists an $N \in \mathbb{N}$ such that the ray graph for any family $(R_i : i \in I)$ of disjoint ϵ -rays in Γ with $|I| \geq N + 2$ is a cycle.*

Furthermore, there is a choice of a cyclic orientation, which we call the correct orientation, of each such ray graph such that any transition function between two families of at least $N + 3$ disjoint ϵ -rays preserves the correct orientation.

Lemma 6.4 ([BEE⁺23, Lemma 7.6, Corollary 7.7 and Corollary 7.9]). *Let Γ be a graph with a half-grid-like end ϵ . Then there exists an $N \in \mathbb{N}$ such that the ray graph K for any family $(R_i: i \in I)$ of disjoint ϵ -rays in Γ with $|I| \geq N + 2$ contains a bare path with at least $|I| - N$ vertices, which we call the central path of K , such that the following statements are true:*

1. *For any $i \in I$, if $K - i$ has precisely two components, each of size at least $N + 1$, then i is an inner vertex of the central path of K .*
2. *There is a choice of an orientation, which we call the correct orientation, of the central path of each such ray graph such that any transition function between two families of at least $N + 3$ disjoint ϵ -rays sends vertices of the central path to vertices of the central path and preserves the correct orientation.*

By Corollary 6.2, if we wish to show that a countable graph G is \preceq -ubiquitous we can restrict our attention to host graphs Γ where each end is non-pebbly. In which case, by Lemmas 6.3 and 6.4 for any end ϵ of Γ , the possible ray graphs, and the possible transition functions between two families of rays, are severely restricted.

Later on in our proof we will be able to restrict our attention to a single end ϵ of Γ and the proof will split into two cases according to whether ϵ is half grid-like or grid-like. However, the two cases are very similar, with the grid-like case being significantly simpler. Therefore, in what follows we will prove only the results necessary for the case where ϵ is half-grid-like, and then later, in Section 8.2, we will briefly sketch the differences for the grid-like case.

6.1. Core rays in the half-grid-like case

By Lemma 6.4, in a half-grid-like end ϵ every ray graph consists, apart for possibly some bounded number of rays on either end, of a bare-path, each of which comes with a correct orientation, which must be preserved by transition functions.

However, in the half-grid itself even more can be seen to be true. There is a natural partial order defined on the set of all rays in the half-grid, where two rays are comparable if they have disjoint tails, and a ray R is less than a ray S if the tail of R lies ‘to the left’ of the tail of S in the half-grid. Then it can be seen that the correct orientations of the central path of any disjoint family of rays can be chosen to agree with this global partial order.

In a general half-grid-like end ϵ a similar thing will be true, but only for a subset of the rays in the end which we call the core rays.

Let us fix for the rest of this section a graph Γ and a half-grid-like end ϵ . By Lemma 6.4, there is some $N \in \mathbb{N}$ such that all but at most N vertices of the ray graph of any large enough family of disjoint ϵ -rays lie on the *central path*.

Definition 6.5 (Core rays). Let R be an ϵ -ray. We say R is a *core ray* (of ϵ) if there is a finite family $\mathcal{R} = (R_i: i \in I)$ of disjoint ϵ -rays with $R = R_c$ for some $c \in I$ such that $\text{RG}(\mathcal{R}) - c$ has precisely two components, each of size at least $N + 1$.

Note that, by Lemma 6.4, such a ray R_c is an inner vertex of the central path of $\text{RG}(\mathcal{R})$. In order to define our partial order on the core rays, we will need to consider what it means for a ray to lie ‘between’ two other rays.

Definition 6.6. Given three ϵ -rays R, S, T such that R, S, T have disjoint tails, we say that S separates R from T if the tails of R and T disjoint from S belong to different ends of $\Gamma - S$.

Lemma 6.7. Let $\mathcal{R} = (R_i: i \in I)$ be a finite family of disjoint ϵ -rays and let $i_1, i_2, j \in I$. Then i_1 and i_2 belong to different components of $\text{RG}(\mathcal{R}) - j$ if and only if R_j separates R_{i_1} from R_{i_2} .

Proof. Suppose that R_{i_1} and R_{i_2} belong to the same end of $\Gamma - V(R_j)$, and let \mathcal{R}' be the subset of $\mathcal{R} \setminus \{R_j\}$ which belong to this end.

Then, \mathcal{R}' is a disjoint family of rays in the same end of $\Gamma - V(R_j)$ and so by Lemma 3.19 the ray graph $\text{RG}_{\Gamma - V(R_j)}(\mathcal{R}')$ is connected. However, it is apparent that $\text{RG}_{\Gamma - V(R_j)}(\mathcal{R}')$ is a subgraph of $\text{RG}_{\Gamma}(\mathcal{R})$, and so i_1 and i_2 belong to the same component of $\text{RG}(\mathcal{R}) - j$.

Conversely, suppose i_1 and i_2 belong to the same component of $\text{RG}(\mathcal{R}) - j$. Then, it is clear that for any two adjacent vertices k and ℓ in $\text{RG}(\mathcal{R}) - j$ the rays R_k and R_ℓ are equivalent in $\Gamma - R_j$, and hence R_{i_1} and R_{i_2} belong to a common end of $\Gamma - R_j$. It follows that R_j does not separate R_{i_1} from R_{i_2} . \square

Lemma 6.8. If R, S, T are ϵ -rays and S separates R from T , then T does not separate R from S and R does not separate S from T .

Proof. As R and T both belong to ϵ , there are infinitely many disjoint paths between them. As S separates R from T , we know that S must meet infinitely many of these paths. Hence, there are infinitely many disjoint paths from S to R , all disjoint from T . Similarly, there are infinitely many disjoint paths from S to T , all disjoint from R . Hence T does not separate R from S and R does not separate S from T . \square

Lemma 6.9. Let R be a core ray of ϵ . Then in $\Gamma - V(R)$ the end ϵ splits into precisely two different ends. (That is, there are two ends ϵ' and ϵ'' of $\Gamma - V(R)$ such that every ϵ -ray in Γ which is disjoint from R is in ϵ' or ϵ'' in $\Gamma - V(R)$.)

Proof. Let $\mathcal{R} = (R_i: i \in I)$ be a finite family of disjoint ϵ -rays witnessing that $R = R_c$ for some $c \in I$ is a core ray. Then there are precisely two ends ϵ' and ϵ'' in $\Gamma - V(R)$ that contain rays in \mathcal{R} , since connected components of $\text{RG}(\mathcal{R}) - c$ are equivalent sets of rays in $\Gamma - V(R)$ and moreover, the two connected components do not contain rays belonging to the same end of $\Gamma - V(R)$ by Lemma 6.7.

Suppose there is a third end in $\Gamma - V(R)$ that contains an ϵ -ray S . We first claim that there is a tail of S which is disjoint from $\bigcup \mathcal{R}$. Indeed, clearly S is disjoint from R , and if S meets $\bigcup \mathcal{R}$ infinitely often then it would meet some $R_i \in \mathcal{R}$ infinitely often, and hence lie in the same end of $\Gamma - V(R)$ as R_i . So let S' be a tail of S which is disjoint from $\bigcup \mathcal{R}$.

Let us consider the family $\mathcal{R}' := \mathcal{R} \cup \{S'\}$, where the ray S' is indexed by some additional index s . Since S' is an ϵ -ray, the ray graph $\text{RG}(\mathcal{R}')$ is connected. Furthermore, since the identity on I is clearly a transition function from \mathcal{R} to \mathcal{R}' , by Lemma 6.4, c is an inner vertex of the central path of $\text{RG}(\mathcal{R}')$, and hence has degree two.

We claim that s is adjacent to some $i \neq c$ in $\text{RG}(\mathcal{R}')$. Indeed, if not, then s must be a leaf of $\text{RG}(\mathcal{R}')$ adjacent to c . In which case, there must be some neighbour i of c in $\text{RG}(\mathcal{R})$ which is not adjacent to c in $\text{RG}(\mathcal{R}')$. However, then s must be adjacent to i in $\text{RG}(\mathcal{R}')$.

However, then clearly s lies in the same end of $\Gamma - V(R)$ as R_i , and hence in either ϵ' or ϵ'' . \square

Hence, every core ray R splits ϵ into two ends. We would like to use this partition to define our partial order on core rays; the core rays in one end will be less than R and the core rays in the other end will be greater than R . However, if we want this partial order to agree with the correct orientation of the central path for any disjoint family of rays in ϵ , then every family of rays $(R_i: i \in I)$ in whose ray graph $R = R_c$ is a vertex of the central path will choose which end of $\Gamma - V(R)$ is less than R and which is greater than R , and we must make sure that this choice is consistent.

So, given a finite family of disjoint ϵ -rays $\mathcal{R} = (R_i: i \in I)$ in whose ray graph $R = R_c$ is a vertex of the central path, we denote by $\top(R, \mathcal{R})$ the end of $\Gamma - V(R)$ containing rays R_i satisfying $i < c$, where $<$ refers to the correct orientation of the vertices of the central path, and with $\perp(R, \mathcal{R})$ the end containing rays R_i satisfying $i > c$. We will show that the labelling \top and \perp is in fact independent of the choice of family \mathcal{R} .

Definition 6.10. Given two (possibly infinite) vertex sets X and Y in Γ , we say that an end ϵ of $\Gamma - X$ is a *sub-end* of an end ϵ' of $\Gamma - Y$ if every ray in ϵ has a tail in ϵ' .

Lemma 6.11. *Let R and S be disjoint core rays of ϵ . Let us suppose that ϵ splits in $\Gamma - V(S)$ into ϵ'_S and ϵ''_S and in $\Gamma - V(R)$ into ϵ'_R and ϵ''_R . If R belongs to ϵ'_S and S belongs to ϵ'_R , then ϵ''_S is a sub-end of ϵ'_R and ϵ''_R is a sub-end of ϵ'_S .*

Proof. Let T be a ray in ϵ''_S . As R belongs to a different end of $\Gamma - V(S)$ than T , there is a tail T' of T which is disjoint from R . As S separates R from T' , we know, by Lemma 6.8, that R does not separate S from T' , hence T' belongs to ϵ'_R . Hence, ϵ''_S is a sub-end of ϵ'_R . Proving that ϵ''_R is a sub-end of ϵ'_S works analogously. \square

Lemma and Definition 6.12. *Let $\mathcal{R}_1 = (R_i: i \in I_1)$ and $\mathcal{R}_2 = (R_i: i \in I_2)$ be two finite families of disjoint ϵ -rays, such for some $c \in I_1 \cap I_2$ the ray R_c lies on the central path of both $\text{RG}(\mathcal{R}_1)$ and $\text{RG}(\mathcal{R}_2)$. Then $\top(R_c, \mathcal{R}_1) = \top(R_c, \mathcal{R}_2)$ and $\perp(R_c, \mathcal{R}_1) = \perp(R_c, \mathcal{R}_2)$.*

We therefore write $\top(\epsilon, R_c)$ for the end $\top(R_c, \mathcal{R}_1)$ and $\perp(\epsilon, R_c)$ accordingly, i.e. $\top(\epsilon, R_c)$ is the end of $\Gamma - V(R_c)$ containing rays that appear on the central path of some ray graph before R_c according to the correct orientation and $\perp(\epsilon, R_c)$ is the end of $\Gamma - V(R_c)$ containing rays that appear on the central path of some ray graph after R_c according to the correct orientation. Note that $\top(\epsilon, R_c) \cap \perp(\epsilon, R_c) = \emptyset$.

Proof. Let ϵ' and ϵ'' be the two ends of $\Gamma - V(R_c)$ and let \mathcal{R}'_1 and \mathcal{R}'_2 be the set of rays in \mathcal{R}_1 and \mathcal{R}_2 respectively that belong to ϵ' , and similarly \mathcal{R}''_1 and \mathcal{R}''_2 be the set of rays in \mathcal{R}_1 and \mathcal{R}_2 respectively that belong to ϵ'' . Let \mathcal{S}' be the larger of \mathcal{R}'_1 and \mathcal{R}'_2 , and similarly \mathcal{S}'' the larger of \mathcal{R}''_1 and \mathcal{R}''_2 .

Let us consider the family of rays $\mathcal{S} := \mathcal{S}' \cup \{R_c\} \cup \mathcal{S}''$. Since the rays in \mathcal{S}' and \mathcal{S}'' belong to different ends of $\Gamma - V(R_c)$, we may, after replacing some of the rays with tails, assume that \mathcal{S} is a family of disjoint rays. We claim that there is a transition function σ_1 from \mathcal{R}_1 to \mathcal{S} which maps R_c to itself, \mathcal{R}'_1 to \mathcal{S}' , and \mathcal{R}''_1 to \mathcal{S}'' .

Let us take a finite separator X which separates ϵ' and ϵ'' in $\Gamma - V(R_c)$. By Lemma 3.15, there is a finite set Y such that any linkage after Y from \mathcal{R}_1 to \mathcal{S} is transitional. Then, since the rays in \mathcal{R}'_1 and \mathcal{S}' belong to the same end of $\Gamma - V(R_c)$ and $|\mathcal{R}'_1| \leq |\mathcal{S}'|$, there is a linkage

after $X \cup Y$ from \mathcal{R}'_1 to S' in $\Gamma - V(R_c)$, and similarly there is a linkage after $X \cup Y$ from \mathcal{R}''_1 to S'' in $\Gamma - V(R_c)$. If we combine these two linkages with a trivial linkage from R_c to itself after $X \cup Y$, we obtain a transitional linkage which induces an appropriate transition function.

The same argument shows that there is a transition function σ_2 from \mathcal{R}_2 to \mathcal{S} which maps R_c to itself, \mathcal{R}'_2 to S' , and \mathcal{R}''_2 to S'' . By Lemma 6.4, both transition functions map vertices of the central path to vertices of the central path and preserve the correct orientation. In particular, R_c lies on the central path of $\text{RG}(\mathcal{S})$.

Moreover, both σ_1 and σ_2 map ϵ' -rays to ϵ' -rays and ϵ'' -rays to ϵ'' -rays. Thus, if $\epsilon' = \top(R_c, \mathcal{S})$, then σ_1 shows that $\epsilon' = \top(R_c, \mathcal{R}_1)$ and σ_2 shows that $\epsilon' = \top(R_c, \mathcal{R}_2)$, and similarly if $\epsilon' = \perp(R_c, \mathcal{S})$. \square

Lemma and Definition 6.13. *Let $\text{core}(\epsilon)$ denote the set of core rays in ϵ . We define a partial order \leq_ϵ on $\text{core}(\epsilon)$ by*

$$R \leq_\epsilon S \text{ if and only if either } R = S, \\ \text{or } R \text{ and } S \text{ have disjoint tails } xR \text{ and } yS \text{ and } xR \in \top(\epsilon, yS)$$

for $R, S \in \text{core}(\epsilon)$.

Proof. We must show that \leq_ϵ is indeed a partial order. For the anti-symmetry, let us suppose that R and S are disjoint rays in $\text{core}(\epsilon)$ such that $R \leq_\epsilon S$ and $S \leq_\epsilon R$, so that $R \in \top(\epsilon, S)$ and $S \in \top(\epsilon, R)$. Let \mathcal{R}_S be a family of rays witnessing that S is a core ray and \mathcal{R}_R a family witnessing that R is a core ray. By Lemma 6.11, $\perp(\epsilon, S)$ is a sub-end of $\top(\epsilon, R)$ and $\perp(\epsilon, R)$ is a sub-end of $\top(\epsilon, S)$. Let $\mathcal{R}_{\perp(S)}$ be the subset of \mathcal{R}_S of rays which belong to $\perp(\epsilon, S)$. Let $\mathcal{R}_{\perp(R)}$ be defined accordingly. By replacing rays with tails, we may assume that all rays in $\mathcal{R} := \mathcal{R}_{\perp(S)} \cup \mathcal{R}_{\perp(R)} \cup \{R\} \cup \{S\}$ are pairwise disjoint. Note that, by the comment after Definition 6.5, both R and S are inner vertices of the central path of $\text{RG}(\mathcal{R})$. Thus, either $S \in \perp(\epsilon, R)$ or $R \in \perp(\epsilon, S)$, contradicting Lemma 6.12.

For the transitivity, let us suppose that R, S, T are rays in $\text{core}(\epsilon)$, such that $R \leq_\epsilon S$ and $S \leq_\epsilon T$. We may assume that R and S , and S and T are disjoint. As \leq_ϵ is anti-symmetric, we have $T \not\leq_\epsilon S$, hence $T \in \perp(\epsilon, S)$. Thus, R and T belong to different ends of $\Gamma - V(S)$, and we may assume that they are also disjoint. As S therefore separates R from T , by Lemma 6.8, we know that T does not separate S from R . Thus, R and S belong to the same end of $\Gamma - V(T)$. Hence $R \in \top(\epsilon, T)$. \square

Remark 6.14. Let $R, S \in \text{core}(\epsilon)$ and let \mathcal{R} be a finite family of disjoint ϵ -rays.

1. Any ray which shares a tail with R is also a core ray of ϵ .
2. If R and S are disjoint, then R and S are comparable under \leq_ϵ .
3. If R and S are on the central path of $\text{RG}(\mathcal{R})$, then $R \leq_\epsilon S$ if and only if R appears before S in the correct orientation of the central path of $\text{RG}(\mathcal{R})$.
4. The maximum number of disjoint rays in $\epsilon \setminus \text{core}(\epsilon)$ is bounded by $2N + 2$.

Lemma 6.15. *Let $R, S \in \text{core}(\epsilon)$ and let $Z \subseteq V(\Gamma)$ be a finite set such that $\top(\epsilon, S)$ and $\perp(\epsilon, S)$ are separated by Z in $\Gamma - V(S)$. Let $H \subseteq \Gamma - Z$ be a connected subgraph which is disjoint to S and contains R and let $T \subseteq H$ be some core ϵ -ray. Then S is in the same relative \leq_ϵ -order to T as to R .*

Proof. Assume $S \leq_\epsilon R$ and hence $R \in \top(\epsilon, S)$. Since H is connected, $T \in \top(\epsilon, S)$ as well and hence $S \leq_\epsilon T$. The other case is analogous. \square

Since, by 6.14 3, the order \leq_ϵ will agree with correct order on the central path, which is preserved by transition functions by Lemma 6.4, the order \leq_ϵ will also be preserved by transition functions, as long as they map core rays to core rays. In order to guarantee that this holds, before linking a family of core rays \mathcal{R} we will first enlarge it slightly by adding some ‘buffer’ rays.

Lemma and Definition 6.16. *Let $\mathcal{R} = (R_i: i \in I)$ be a finite family of disjoint core ϵ -rays. Then there exists a finite family $\overline{\mathcal{R}} \supset \mathcal{R}$ of disjoint ϵ -rays such that*

- *For each $i \in I$, the graph $\text{RG}(\overline{\mathcal{R}}) - i$ has precisely two components, each of size at least $N + 1$;*
- *Each $i \in I$ is an inner vertex of the central path of $\text{RG}(\overline{\mathcal{R}})$;*
- $|\overline{\mathcal{R}}| = |\mathcal{R}| + 2N + 2$.

Even though such a family is not unique, we denote by $\overline{\mathcal{R}}$ an arbitrary such family.

Proof. By Remark 6.142, the rays in \mathcal{R} are linearly ordered by \leq_ϵ . Let R denote the \leq_ϵ -greatest and S denote the \leq_ϵ -smallest element of \mathcal{R} .

As in the proof of Lemma 6.13, let \mathcal{S}_R and \mathcal{S}_S be families of disjoint rays witnessing that R and S are core rays, and let $\mathcal{S}_{\perp(R)}$ be the subset of rays of \mathcal{S}_R belonging to $\perp(\epsilon, R)$ and $\mathcal{S}_{\top(S)}$ be the subset of rays of \mathcal{S}_S belonging to $\top(\epsilon, S)$. Note that, by definition both $\mathcal{S}_{\perp(R)}$ and $\mathcal{S}_{\top(S)}$ contain at least $N + 1$ rays, and we may in fact assume without loss of generality that they both contain exactly $N + 1$ rays.

Now $\mathcal{S}_{\perp(R)} \subseteq \perp(\epsilon, R)$ and $R' \in \top(\epsilon, R)$ for every $R' \in \mathcal{R} \setminus \{R\}$, and each ray in $\mathcal{S}_{\perp(R)}$ has a tail disjoint to $\bigcup \mathcal{R}$. Analogously, $\mathcal{S}_{\top(S)} \subseteq \top(\epsilon, S)$ and $R' \in \perp(\epsilon, S)$ for every $R' \in \mathcal{R} \setminus \{S\}$ and each ray in $\mathcal{S}_{\top(S)}$ has a tail disjoint to $\bigcup \mathcal{R}$. Now, $\mathcal{S}_{\top(S)} \subseteq \top(\epsilon, R)$ and $\mathcal{S}_{\perp(R)} \subseteq \perp(\epsilon, S)$ by Lemma 6.11, yielding that tails of rays in $\mathcal{S}_{\top(S)}$ are necessarily disjoint from tails in $\mathcal{S}_{\perp(R)}$.

Let $\overline{\mathcal{R}}$ be the union of \mathcal{R} with appropriate tails of each ray in $\mathcal{S}_{\perp(R)} \cup \mathcal{S}_{\top(S)}$. Note that $|\overline{\mathcal{R}}| = |\mathcal{R}| + 2N + 2$. For any ray $R_i \in \mathcal{R}$, we first note that that $S \leq_\epsilon R_i$ and so $S \in \top(\epsilon, R_i)$ and $R_i \in \perp(\epsilon, S)$. Then, since $\mathcal{S}_{\top(S)} \subseteq \top(\epsilon, S)$ it follows from Lemma 6.11 that $\mathcal{S}_{\top(S)} \subseteq \top(\epsilon, R_i)$, and hence one of the components of $\text{RG}(\overline{\mathcal{R}}) - i$ has size at least $N + 1$. A similar argument shows that a second component has size at least $N + 1$, and finally, since R_i is a core ray, by Lemma 6.9, there are no other components of $\text{RG}(\overline{\mathcal{R}}) - i$. Finally, by the comment after Definition 6.5, it follows that R_i is an inner vertex of the central path of this ray graph. \square

Lemma 6.17 ([BEE⁺23, Lemma 7.10]). *Let \mathcal{R} and \mathcal{T} be families of disjoint rays, each of size at least $N + 3$, and let σ be a transition function from \mathcal{R} to \mathcal{T} . Let $x \in \text{RG}(\mathcal{R})$ be an inner vertex of the central path. If $v_1, v_2 \in \text{RG}(\mathcal{R})$ lie in different components of $\text{RG}(\mathcal{R}) - x$, then $\sigma(v_1)$ and $\sigma(v_2)$ lie in different components of $\text{RG}(\mathcal{T}) - \sigma(x)$. Moreover, $\sigma(x)$ is an inner vertex of the central path of $\text{RG}(\mathcal{T})$.*

Definition 6.18. Let \mathcal{R}, \mathcal{S} be finite families of disjoint ϵ -rays and let \mathcal{R}' be a subfamily of \mathcal{R} consisting of core rays. A linkage \mathcal{P} between \mathcal{R} and \mathcal{S} is *preserving on \mathcal{R}'* if \mathcal{P} links \mathcal{R}' to core rays and preserves the order \leq_ϵ .

Lemma 6.19. *Let $\mathcal{R} = (R_i : i \in I)$ be a finite family of disjoint core ϵ -rays and $\mathcal{S} = (S_j : j \in J)$ be a finite family of disjoint ϵ -rays. Let $\overline{\mathcal{R}} = (R_i : i \in \overline{I})$ be as in Lemma 6.16 and let \mathcal{P} be a linkage from $\overline{\mathcal{R}}$ to \mathcal{S} . If \mathcal{P} is transitional, then it is preserving on \mathcal{R} .*

Proof. We first note that, by Lemma 6.4, if \mathcal{P} links the rays in \mathcal{R} to core rays, then it will be preserving.

So, let $\sigma : \overline{I} \rightarrow J$ be the transition function induced by \mathcal{P} . For each $i \in I$, since i is an inner vertex of the central path of $\text{RG}(\overline{\mathcal{R}})$, by Lemma 6.17, $\sigma(i)$ is an inner vertex of the central path of $\text{RG}(\mathcal{S})$. Since the central path is a bare path, it follows that $\text{RG}(\mathcal{S}) - \sigma(i)$ has precisely two components.

Furthermore, by Lemma 6.16, the graph $\text{RG}(\overline{\mathcal{R}}) - i$ has precisely two components, each of size at least $N + 1$, and so by Lemma 6.17 the two components of $\text{RG}(\mathcal{S}) - \sigma(i)$ each have size at least $N + 1$. Hence, the family \mathcal{S} witnesses that $S_{\sigma(i)}$ is a core ray. \square

Definition 6.20. If \mathcal{P} is a linkage from \mathcal{R} to \mathcal{S} , then a *sub-linkage* of \mathcal{P} is just a subset of \mathcal{P} , considered as a linkage from the corresponding subset of \mathcal{R} to \mathcal{S} .

Remark 6.21. A sub-linkage of a transitional linkage is transitional.

The following remarks are a direct consequence of the definitions and Lemma 6.4.

Remark 6.22. Let \mathcal{R} be a finite family of disjoint core ϵ -rays and let \mathcal{S} and \mathcal{T} be finite families of disjoint ϵ -rays. Let $\overline{\mathcal{R}}$ be as in Lemma 6.16 and let \mathcal{P}_1 and \mathcal{P}_2 be linkages from $\overline{\mathcal{R}}$ to \mathcal{S} and from $(\overline{\mathcal{R}} \circ_{\mathcal{P}_1} \mathcal{S})$ to \mathcal{T} respectively.

1. If \mathcal{P}_1 is preserving on \mathcal{R} , then any $\mathcal{P}'_1 \subseteq \mathcal{P}_1$ as a linkage between the respective subfamilies is preserving on the respective subfamily of \mathcal{R} .
2. If \mathcal{P}_1 is preserving on \mathcal{R} and \mathcal{P}_2 is preserving on $\overline{\mathcal{R}} \circ_{\mathcal{P}_1} \mathcal{S}$, then the concatenation $\mathcal{P}_1 + \mathcal{P}_2$ is preserving on \mathcal{R} .

Lemma 6.23. *Let \mathcal{R} and \mathcal{S} be finite families of disjoint core rays of ϵ and let $\mathcal{S}' \subseteq \mathcal{S}$ be a subfamily of \mathcal{S} with $|\mathcal{R}| = |\mathcal{S}'|$. Then there is a transitional linkage from $\overline{\mathcal{R}}$ to $\overline{\mathcal{S}}$ which is preserving on \mathcal{R} and links the rays in \mathcal{R} to rays in \mathcal{S}' .*

Proof. Consider $\mathcal{T} := (\overline{\mathcal{S}} \setminus \mathcal{S}) \cup \mathcal{S}' \subseteq \overline{\mathcal{S}}$. It is apparent that the family \mathcal{T} satisfies the conclusions of Lemma 6.16 for \mathcal{S}' .

Let σ be some transition function between $\overline{\mathcal{R}}$ and \mathcal{T} and let \mathcal{P} be a linkage inducing this transition function. By Lemma 6.19 this linkage is preserving on \mathcal{R} . Note that, since σ is a transition function from $\overline{\mathcal{R}}$ to \mathcal{T} , it is also a transition function from $\overline{\mathcal{R}}$ to $\overline{\mathcal{S}}$, and so \mathcal{P} is also a preserving, transitional linkage from $\overline{\mathcal{R}}$ to $\overline{\mathcal{S}}$. We claim further that \mathcal{P} links the rays in \mathcal{R} to the rays in \mathcal{S}' .

Indeed, since $|\overline{\mathcal{R}}| = |\mathcal{T}| = |\mathcal{R}| + 2N + 2$, we may assume for a contradiction that there is some $R_i \in \mathcal{R}$ such that $S_{\sigma(i)} \notin \mathcal{S}'$. Note that, since i is an inner vertex of the central path of $\text{RG}(\overline{\mathcal{R}})$, by Lemma 6.17 $\sigma(i)$ is an inner vertex of the central path of $\text{RG}(\mathcal{T})$, and so in particular $\text{RG}(\mathcal{T}) - \sigma(i)$ has precisely two components.

Since for each $S_j \in \mathcal{S}'$, j lies on the central path of $\text{RG}(\mathcal{T})$, if $S_{\sigma(i)} \notin \mathcal{S}'$ then it is clear that $\text{RG}(\mathcal{T}) \setminus \sigma(i)$ contains one component of size at least $|\mathcal{S}'| + N + 1 = |\mathcal{R}| + N + 1$. However, since i is an inner vertex of the central path of $\text{RG}(\overline{\mathcal{R}})$, by Lemma 6.17 and Lemma 6.19 there must be two components of $\text{RG}(\mathcal{T}) \setminus \sigma(i)$ of size at least $N + 1$, a contradiction. \square

7. G -tribes and concentration of G -tribes towards an end

To show that a given graph G is \preceq -ubiquitous, we shall assume that $nG \preceq \Gamma$ for every $n \in \mathbb{N}$ and need to show that this implies $\aleph_0 G \preceq \Gamma$. To this end we use the following notation for such collections of nG in Γ , which we introduced in [BEE⁺22] and [BEE⁺23].

Definition 7.1 (G -tribes). Let G and Γ be graphs.

- A G -tribe in Γ (with respect to the minor relation) is a family \mathcal{F} of finite collections F of disjoint subgraphs H of Γ , such that each member H of \mathcal{F} is an IG .
- A G -tribe \mathcal{F} in Γ is called *thick* if for each $n \in \mathbb{N}$, there is a layer $F \in \mathcal{F}$ with $|F| \geq n$; otherwise, it is called *thin*.
- A G -tribe \mathcal{F} is *connected* if every member H of \mathcal{F} is connected. Note that, this is the case precisely if G is connected.
- A G -tribe \mathcal{F}' in Γ is a G -subtribe⁵ of a G -tribe \mathcal{F} in Γ , denoted by $\mathcal{F}' \preceq \mathcal{F}$, if there is an injection $\Psi: \mathcal{F}' \rightarrow \mathcal{F}$ such that for each $F' \in \mathcal{F}'$, there is an injection $\varphi_{F'}: F' \rightarrow \Psi(F')$ with $V(H') \subseteq V(\varphi_{F'}(H'))$ for every $H' \in F'$. The G -subtribe \mathcal{F}' is called *flat*, denoted by $\mathcal{F}' \subseteq \mathcal{F}$, if there is such an injection Ψ satisfying $F' \subseteq \Psi(F')$.
- A thick G -tribe \mathcal{F} in Γ is *concentrated at an end* ϵ of Γ if for every finite vertex set X of Γ , the G -tribe $\mathcal{F}_X = \{F_X: F \in \mathcal{F}\}$ consisting of the layers

$$F_X = \{H \in F: H \not\subseteq C(X, \epsilon)\} \subseteq F$$

is a thin subtribe of \mathcal{F} .

We note that, if G is connected, every thick G -tribe \mathcal{F} contains a thick subtribe \mathcal{F}' such that every $H \in \bigcup \mathcal{F}'$ is a tidy IG . We will use the following lemmas from [BEE⁺22].

⁵When G is clear from the context we will often refer to a G -subtribe as simply a subtribe.

Lemma 7.2 (Removing a thin subtribe, [BEE⁺22, Lemma 5.2]). *Let \mathcal{F} be a thick G -tribe in Γ and let \mathcal{F}' be a thin subtribe of \mathcal{F} , witnessed by $\Psi: \mathcal{F}' \rightarrow \mathcal{F}$ and $(\varphi_{F'}: F' \in \mathcal{F}')$. For $F \in \mathcal{F}$, if $F \in \Psi(\mathcal{F}')$, let $\Psi^{-1}(F) = \{F'_F\}$ and set $\hat{F} = \varphi_{F'_F}(F'_F)$. If $F \notin \Psi(\mathcal{F}')$, set $\hat{F} = \emptyset$. Then*

$$\mathcal{F}'' := \{F \setminus \hat{F} : F \in \mathcal{F}\}$$

is a thick flat G -subtribe of \mathcal{F} .

Lemma 7.3 (Pigeon hole principle for thick G -tribes, [BEE⁺22, Lemma 5.3]). *Let $k \in \mathbb{N}$ and let $c: \bigcup \mathcal{F} \rightarrow [k]$ be a k -colouring of the members of some thick G -tribe \mathcal{F} in Γ . Then there is a monochromatic, thick, flat G -subtribe \mathcal{F}' of \mathcal{F} .*

Lemma 7.4 ([BEE⁺22, Lemma 5.4]). *Let G be a connected graph and Γ a graph containing a thick connected G -tribe \mathcal{F} . Then either $\aleph_0 G \preceq \Gamma$, or there is a thick flat subtribe \mathcal{F}' of \mathcal{F} and an end ϵ of Γ such that \mathcal{F}' is concentrated at ϵ .*

Lemma 7.5 ([BEE⁺22, Lemma 5.5]). *Let G be a connected graph and Γ a graph containing a thick connected G -tribe \mathcal{F} concentrated at an end ϵ of Γ . Then the following assertions hold:*

1. *For every finite set X , the component $C(X, \epsilon)$ contains a thick flat G -subtribe of \mathcal{F} .*
2. *Every thick subtribe \mathcal{F}' of \mathcal{F} is concentrated at ϵ .*

The following lemma from [BEE⁺23] shows that we can restrict ourself to thick G -tribes which are concentrated at thick ends.

Lemma 7.6 ([BEE⁺23, Lemma 8.7]). *Let G be a connected graph and Γ a graph containing a thick G -tribe \mathcal{F} concentrated at an end $\epsilon \in \Omega(\Gamma)$ which is thin. Then $\aleph_0 G \preceq \Gamma$.*

Given an extensive tree-decomposition (T, \mathcal{V}) of G , broadly our strategy will be to obtain a family of disjoint IG 's by choosing a sequence of initial subtrees $T_0 \subseteq T_1 \subseteq \dots$ such that $\bigcup T_i = T$ and constructing inductively a family of finitely many $IG(T_{k+1})$'s which extend the $IG(T_k)$'s built previously (cf. Definition 4.2). The extensiveness of the tree-decomposition will ensure that, at each stage, there will be some edges in $\partial(T_i) = E(T_i, T \setminus T_i)$, each of which has in G a family of rays \mathcal{R}_e along which the graph G displays self-similarity.

In order to extend our $IG(T_k)$ at each step, we will want to assume that the IG 's in our thick G -tribe \mathcal{F} lie in a 'uniform' manner in the graph Γ in terms of these rays \mathcal{R}_e .

More specifically, for each edge $e \in \partial(T_i)$, the rays \mathcal{R}_e provided by the extensive tree-decomposition in Definition 4.4 tend to a common end ω_e in G , and for each $H \in \bigcup \mathcal{F}$, the corresponding rays in H converge to an end $H(\omega_e) \in \Omega(\Gamma)$ (cf. Definition 3.13), which might either be the end ϵ of Γ at which \mathcal{F} is concentrated, or another end of Γ . We would like that our G -tribe \mathcal{F} makes a consistent choice across all members H of \mathcal{F} of whether $H(\omega_e)$ is ϵ , for each $e \in \partial(T_i)$.

Furthermore, if $H(\omega_e) = \epsilon$ for every $H \in \bigcup \mathcal{F}$, then this imposes some structure on the end ω_e of G . More precisely, by [BEE⁺23, Lemma 10.1] (see Lemma 7.10), we may assume that $\text{RG}_H(H^\downarrow(\mathcal{R}_e))$ is a path for each member H of the G -tribe \mathcal{F} , or else we immediately find that $\aleph_0 G \preceq \Gamma$ and are done.

By moving to a thick subtribe, we may assume that every ϵ -ray in H is a core ray for every $H \in \bigcup \mathcal{F}$, in which case \leq_ϵ imposes a linear order on every family of rays $H^\downarrow(\mathcal{R}_e)$, which induces one of the two distinct orientations of the path $\text{RG}_H(H^\downarrow(\mathcal{R}_e))$. We will also want that our tribe \mathcal{F} induces this orientation in a consistent manner.

Let us make the preceding discussion precise with the following definitions:

Definition 7.7. Let G be a connected locally finite graph with an extensive tree-decomposition (T, \mathcal{V}) and S be an initial subtree of T . Let $H \subseteq \Gamma$ be a tidy IG , \mathcal{H} be a set of tidy IG s in Γ , and ϵ a prescribed end of Γ .

- Given an end ω of G , we say that ω *converges to ϵ according to H* if $H(\omega) = \epsilon$ (cf. Definition 3.13). The end ω *converges to ϵ according to \mathcal{H}* if it converges to ϵ according to every element of \mathcal{H} .

We say that ω is *cut from ϵ according to H* if $H(\omega) \neq \epsilon$. The end ω is *cut from ϵ according to \mathcal{H}* if it is cut from ϵ according to every element of \mathcal{H} .

Finally, we say that \mathcal{H} *determines whether ω converges to ϵ* if either ω converges to ϵ according to \mathcal{H} or ω is cut from ϵ according to \mathcal{H} .

- Given $E \subseteq E(T)$, we say \mathcal{H} *weakly agrees about E (with respect to ϵ)* if for each $e \in E$, the set \mathcal{H} determines whether ω_e (cf. Definition 4.4) converges to ϵ . When the end ϵ is clear from context, we will simply say that \mathcal{H} weakly agrees about E . If \mathcal{H} weakly agrees about $\partial(S)$ we let

$$\begin{aligned} \partial_\epsilon(S) &:= \{e \in \partial(S) : \omega_e \text{ converges to } \epsilon \text{ according to } \mathcal{H}\}, \\ \partial_{-\epsilon}(S) &:= \{e \in \partial(S) : \omega_e \text{ is cut from } \epsilon \text{ according to } \mathcal{H}\}, \end{aligned}$$

and write

$$\begin{aligned} S^{-\epsilon} &\text{ for the component of the forest } T - \partial_\epsilon(S) \text{ containing the root of } T, \\ S^\epsilon &\text{ for the component of the forest } T - \partial_{-\epsilon}(S) \text{ containing the root of } T. \end{aligned}$$

Note that $S = S^{-\epsilon} \cap S^\epsilon$.

- We say that \mathcal{H} is *well-separated from ϵ at S* if \mathcal{H} weakly agrees about $\partial(S)$ and $H(S^{-\epsilon})$ can be separated from ϵ in Γ for all elements $H \in \mathcal{H}$, i.e. for every H there is a finite $X \subseteq V(\Gamma)$ such that $H(S^{-\epsilon}) \cap C_\Gamma(X, \epsilon) = \emptyset$.

In the case that ϵ is half-grid-like, we say that \mathcal{H} *strongly agrees about $\partial(S)$* (with respect to ϵ) if

- it weakly agrees about $\partial(S)$;
- for each $H \in \mathcal{H}$, every ϵ -ray $R \subseteq H$ is in $\text{core}(\epsilon)$;
- for every $e \in \partial_\epsilon(S)$, there is a linear order $\leq_{\mathcal{H},e}$ on $S(e)$ (cf. Definition 4.4), such that the order induced on $H^\downarrow(\mathcal{R}_e)$ by $\leq_{\mathcal{H},e}$ agrees with \leq_ϵ on $H^\downarrow(\mathcal{R}_e)$ for all $H \in \mathcal{H}$.

Again, when the end ϵ is clear from context, we will simply say that \mathcal{H} strongly agrees about $\partial(S)$.
 If \mathcal{F} is a thick G -tribe concentrated at an end ϵ , we use these terms in the following way:

- Given $E \subseteq E(T)$, we say that \mathcal{F} *weakly agrees about* E if $\bigcup \mathcal{F}$ weakly agrees about E .
- We say that \mathcal{F} is *well-separated from* ϵ at S if $\bigcup \mathcal{F}$ is.
- We say that \mathcal{F} *strongly agrees about* $\partial(S)$ if $\bigcup \mathcal{F}$ does.

For ease of presentation, when a G -tribe \mathcal{F} strongly agrees about $\partial(S)$ we will write $\leq_{\mathcal{F},\epsilon}$ for $\leq_{\bigcup \mathcal{F},\epsilon}$.

Remark 7.8. The properties of weakly agreeing about E , being well-separated from ϵ , and strongly agreeing about $\partial(S)$ are all preserved under taking subsets, and hence under taking flat subtribes.

Note that by the pigeon hole principle for thick G -tribes, given a finite edge set $E \subseteq E(T)$, any thick G -tribe \mathcal{F} concentrated at ϵ has a thick (flat) subtribe which weakly agrees about E .

The next few lemmas show that, with some slight modification, we may restrict to a further subtribe which strongly agrees about E and is also well-separated from ϵ .

Definition 7.9 ([BEE+23, Lemma 3.5]). Let ω be an end of a graph G . We say ω is *linear* if $\text{RG}(\mathcal{R})$ is a path for every finite family \mathcal{R} of disjoint ω -rays.

Lemma 7.10 ([BEE+23, Lemma 10.1]). *Let ϵ be a non-pebbly end of Γ and let \mathcal{F} be a thick G -tribe, such that for every $H \in \bigcup \mathcal{F}$, there is an end $\omega_H \in \Omega(G)$ such that $H(\omega_H) = \epsilon$. Then there is a thick flat subtribe \mathcal{F}' of \mathcal{F} such that ω_H is linear for every $H \in \bigcup \mathcal{F}'$.*

Corollary 7.11. *Let G be a connected locally finite graph with an extensive tree-decomposition (T, \mathcal{V}) , S an initial subtree of T , and let \mathcal{F} be a thick G -tribe which is concentrated at a non-pebbly end ϵ of a graph Γ and weakly agrees about $\partial(S)$. Then ω_e is linear for every $e \in \partial_\epsilon(S)$.*

Proof. Since \mathcal{F} weakly agrees about $\partial(S)$, $H(\omega_e) = \epsilon$ for each $e \in \partial_\epsilon(S)$. Hence, we may apply Lemma 7.10 to \mathcal{F} with $\omega_H = \omega_e$ for each $H \in \bigcup \mathcal{F}$ and conclude that there is a thick flat subtribe \mathcal{F}' of \mathcal{F} for which ω_H is linear for every $H \in \bigcup \mathcal{F}'$. In particular, since $\omega_H = \omega_e$ for each $H \in \bigcup \mathcal{F}'$, it follows that ω_e is linear. \square

Lemma 7.12. *Let G be a connected locally finite graph with an extensive tree-decomposition (T, \mathcal{V}) and let S be an initial subtree of T with $\partial(S)$ finite. Let \mathcal{F} be a thick G -tribe in a graph Γ , which weakly agrees about $\partial(S) \subseteq E(T)$, concentrated at a half-grid-like end ϵ of Γ . Then \mathcal{F} has a thick flat subtribe \mathcal{F}' so that \mathcal{F}' strongly agrees about $\partial(S)$.*

Proof. Since ϵ is half-grid-like, there is some $N \in \mathbb{N}$ as in Lemma 6.4. Then, by Remark 6.144, given any family of disjoint ϵ -rays, at least $m - 2N - 2$ of them are core rays. Thus, since all members of a layer F of \mathcal{F} are disjoint, at least $|F| - 2N - 2$ members of F do not contain any ϵ -ray which is not core. Thus, there is a thick flat subtribe \mathcal{F}^* of \mathcal{F} such that all ϵ -rays in members of \mathcal{F}^* are core.

Given a member H of \mathcal{F}^* and $e \in \partial_\epsilon(S)$, we consider the order $\leq_{H,e}$ induced on $S(e)$ by the order \leq_ϵ on $H^\downarrow(\mathcal{R}_e)$. Let O_e be the set of potential orders on $S(e)$ which is finite since $S(e)$ is finite⁶. Consider the colouring $c: \bigcup \mathcal{F}^* \rightarrow \prod_{e \in \partial_\epsilon(S)} O_e$ where we map every H to the product of the orders $\leq_{H,e}$ it induces. By the pigeon hole principle for thick G -tribes, Lemma 7.3, there is a monochromatic, thick, flat G -subtribe \mathcal{F}' of \mathcal{F}^* . We can now set $\leq_{\mathcal{F}',e} := \leq_{H,e}$ for some $H \in \mathcal{F}'$. Then, by Remark 7.8, this order $\leq_{\mathcal{F}',e}$ witnesses that \mathcal{F}' is a thick flat subtribe of \mathcal{F} which strongly agrees about $\partial(S)$. \square

Lemma 7.13. *Let G be a connected locally finite graph with an extensive tree-decomposition (T, \mathcal{V}) . Let $H \subseteq \Gamma$ be a tidy IG and ϵ an end of Γ . Let e be an edge of T such that $H(\omega_e) \neq \epsilon$. Then there is a finite set $X \subseteq V(\Gamma)$ such that for every finite $X' \supseteq X$, there exists a push-out $H_e = H(G[A(e)]) \oplus_{H^\downarrow(\mathcal{R}_e)} W_e$ of H along e to some depth $n \in \mathbb{N}$ so that*

$$C_\Gamma(X', H(\omega_e)) \neq C_\Gamma(X', \epsilon) \text{ and } W_e \subseteq C_\Gamma(X', H(\omega_e)).$$

Proof. Let $X \subseteq V(\Gamma)$ be a finite vertex set such that $C_\Gamma(X, H(\omega_e)) \neq C_\Gamma(X, \epsilon)$. Then $C_\Gamma(X', H(\omega_e)) \neq C_\Gamma(X', \epsilon)$ holds for any finite vertex set $X' \supseteq X$. Furthermore, since X' is finite, there are only finitely many $v \in V(G)$ whose branch sets $H(v)$ meet X' . By extensiveness, every vertex of G is contained in only finitely many parts of the tree-decomposition, and so there exists an $n \in \mathbb{N}$ such that whenever $e' \in E(T_{e^+})$ is such that $\text{dist}(e^-, e'^-) \geq n$, then

$$H(G[B(e')]) \cap X' = \emptyset, \text{ and so } H(G[B(e')]) \subseteq C_\Gamma(X', H(\omega_e)).$$

Since (T, \mathcal{V}) is an extensive tree-decomposition, there is a witness W of the self-similarity of $B(e)$ at distance at least n . Then by Definition 4.8 and Lemma 4.9, there is a push-out $H_e = H(G[A(e)]) \oplus_{H^\downarrow(\mathcal{R}_e)} H(W)$ of H along e to depth n . Let $W_e = H(W)$, then by Definition 4.8, $V(W_e) \subseteq V(H(G[B(e')])) \subseteq C_\Gamma(X', H(\omega_e))$. \square

Lemma 7.14. *Let G be a connected locally finite graph with an extensive tree-decomposition (T, \mathcal{V}) with root $r \in T$. Let Γ be a graph and \mathcal{F} a thick G -tribe concentrated at a half-grid-like end ϵ of Γ . Then there is a thick subtribe \mathcal{F}' of \mathcal{F} such that*

- (1) \mathcal{F}' is concentrated at ϵ .
- (2) \mathcal{F}' strongly agrees about $\partial(\{r\})$.
- (3) \mathcal{F}' is well-separated from ϵ at $\{r\}$.

Proof. Since T is locally finite, also $d(r)$ is finite, and, by choosing a thick flat subtribe of \mathcal{F} , we may assume that \mathcal{F} weakly agrees about $\partial(\{r\})$. Moreover, by Lemma 7.12, we may even assume that \mathcal{F} strongly agrees about $\partial(\{r\})$. Using Lemma 7.5(2), this \mathcal{F} would then satisfy (1) and (2). So, it remains to arrange for (3):

Given a member H of \mathcal{F} , for every $e \in \partial_{-\epsilon}(\{r\})$, there exists, by Lemma 7.13, a finite set $X_e \subseteq V(\Gamma)$, such that for every finite vertex set $X' \supseteq X_e$ there is a push-out

⁶Note that there are in fact at most two orders of $S(e)$ induced by one of the members of \mathcal{F}^* since ω_e is linear by Corollary 7.11.

$H_e = H(G[A(e)]) \oplus_{H^\downarrow(\mathcal{R}_e)} W_e$ of H along e , so that $C_\Gamma(X', H(\omega_e)) \neq C_\Gamma(X', \epsilon)$ and $W_e \subseteq C_\Gamma(X', H(\omega_e))$. Let X be the union of all these X_e together with $H(\{r\})$. For each $e \in \partial_{-\epsilon}(\{r\})$, let H_e be the push-out whose existence is guaranteed by the above with respect to this set X .

Let us define an IG

$$H' := \bigcup_{e \in \partial_{-\epsilon}(\{r\})} H_e(\{r\}^\epsilon \cup T_{e^+}).$$

It is straightforward, although not quick, to check that this is indeed an IG and so we will not do this in detail. Briefly, this can be deduced from multiple applications of Definition 4.7, and, since each $H_e(G[A(e)])$ extends $H(G[A(e)])$ fixing $A(e) \setminus S(e)$, all that we need to check is that the extra vertices added to the branch sets of vertices in $S(e)$ are distinct for each edge e . However, this follows from Definition 4.8, since these vertices come from $\bigcup H^\downarrow(\mathcal{R}_e)$ and the rays $R_{e,s}$ and $R_{e',s'}$ are disjoint except in their initial vertex when $s = s'$. Let \mathcal{F}' be the tribe given by $\{F' : F \in \mathcal{F}\}$, where $F' = \{H' : H \in F\}$ for each $F \in \mathcal{F}$. We claim that \mathcal{F}' satisfies the conclusion of the lemma.

Firstly, by Lemma 7.5(2), \mathcal{F}' is concentrated at ϵ , i.e. (1) holds. Next, we claim that \mathcal{F}' strongly agrees about $\partial(\{r\})$. Indeed, by construction for each $e \in \partial_{-\epsilon}(\{r\})$ we have $W_e \subseteq C_\Gamma(X, H(\omega_e))$, and hence ω_e is cut from ϵ according to H' . Furthermore, by construction $H(\{r\}^\epsilon) \setminus X = H'(\{r\}^\epsilon) \setminus X$ and so ω_e converges to ϵ according to H' for every $e \in \partial_{-\epsilon}(\{r\})$. In fact, $H^\downarrow(\mathcal{R}_e) = H'^\downarrow(\mathcal{R}_e)$ for every $e \in \partial_{-\epsilon}(\{r\})$. Finally, since $H' \subseteq H$, and \mathcal{F} strongly agrees about $\partial(\{r\})$, it follows that every ϵ -ray in H' is in $\text{core}(\epsilon)$, and so (2) holds.

It remains to show that \mathcal{F}' is well-separated from ϵ at $\{r\}$. Now, $H'(\{r\}^{-\epsilon}) \setminus \bigcup_{e \in \partial_{-\epsilon}(\{r\})} W_e$ is finite, and each W_e is separated from ϵ by X . Hence, there is some finite set Y separating $H'(\{r\}^{-\epsilon})$ from ϵ , and so (3) holds. \square

Lemma 7.15 (Well-separated push-out). *Let G be a connected locally-finite graph with an extensive tree-decomposition (T, \mathcal{V}) . Let $H \subseteq \Gamma$ be a tidy IG and ϵ an end of Γ . Let S be a finite initial subtree of T , such that $\{H\}$ is well-separated from ϵ at S , and let $f \in \partial_\epsilon(S)$. Then there exists a push-out H' of H along f to depth 0 (see Definition 4.8) such that $\{H'\}$ is well-separated from ϵ at $\tilde{S} := S + f \subseteq T$.*

Proof. Let $X' \subseteq V(\Gamma)$ be a finite set with $H(S^{-\epsilon}) \cap C_\Gamma(X', \epsilon) = \emptyset$. If $\partial_{-\epsilon}(\tilde{S}) \setminus \partial(S) = \emptyset$, then $H' = H$ satisfies the conclusion of the lemma, hence we may assume that $\partial_{-\epsilon}(\tilde{S}) \setminus \partial(S)$ is non-empty.

By applying Lemma 7.13 to every $e \in \partial_{-\epsilon}(\tilde{S}) \setminus \partial(S)$, we obtain a finite set $X \supseteq X'$ and a family $(H_e : e \in \partial_{-\epsilon}(\tilde{S}) \setminus \partial(S))$ where each $H_e = H(G[A(e)]) \oplus_{H^\downarrow(\mathcal{R}_e)} W_e$ is a push-out of H along e such that $W_e \subseteq C_\Gamma(X, H(\omega_e)) \neq C_\Gamma(X, \epsilon)$.

Let

$$H' := \bigcup_{e \in \partial_{-\epsilon}(\tilde{S}) \setminus \partial(S)} H_e(S^\epsilon \cup T_{e^+}).$$

As before, it is straightforward to check that H' is an IG , and that H' is a push-out of H along f to depth 0. We claim that H' is well-separated from ϵ at \tilde{S} .

Since X' separates $H(S^{-\epsilon})$ from ϵ , and $\partial_{-\epsilon}(\tilde{S}) \setminus \partial(S)$ is finite, it will be sufficient to show that for each $e \in \partial_{-\epsilon}(\tilde{S}) \setminus \partial(S)$, there is a finite set X_e which separates $H'(G[B(e)])$ from ϵ

in Γ . However, by construction X separates W_e from ϵ , and $H'(G[B(e)]) \setminus W_e$ is finite, and so the claim follows. \square

The following lemma contains a large part of the work needed for our inductive construction. The idea behind the statement is the following: At step n in our construction, we will have a thick G -tribe \mathcal{F}_n which agrees about $\partial(T_n)$, where T_n is an initial subtree of the decomposition tree T with finite $\partial(T_n)$, which will allow us to extend our $IG(T_n)$'s to $IG(T_{n+1})$'s, where T_{n+1} is a larger initial subtree of T , again with finite $\partial(T_{n+1})$. In order to perform the next stage of our construction, we will need to 'refine' \mathcal{F}_n to a thick G -tribe \mathcal{F}_{n+1} which agrees about $\partial(T_{n+1})$.

This would be a relatively simple application of the pigeon hole principle for G -tribes, Lemma 7.3, except that, in our construction, we cannot extend by a member of \mathcal{F}_{n+1} naively. Indeed, suppose we wish to use an IG , say H , to extend an $IG(T_n)$ to an $IG(T_{n+1})$. There is some subgraph, $H(T_{n+1} \setminus T_n)$, of H which is an $IG(T_{n+1} \setminus T_n)$, however in order to use this to extend the $IG(T_n)$ we first have to link the branch sets of the boundary vertices to this subgraph, and there may be no way to do so without using other vertices of $H(T_{n+1} \setminus T_n)$.

For this reason, we will ensure the existence of an 'intermediate G -tribe' \mathcal{F}^* , which has the property that for each member H of \mathcal{F}^* , there are push-outs at arbitrary depth of H which are members of \mathcal{F}_{n+1} . This allows us to first link our $IG(T_n)$ to some $H \in \mathcal{F}^*$ and then choose a push-out $H' \in \mathcal{F}_{n+1}$ of H , such that $H'(T_{n+1} \setminus T_n)$ avoids the vertices we used in our linkage.

Lemma 7.16 (*G -tribe refinement lemma*). *Let G be a connected locally finite graph with an extensive tree-decomposition (T, \mathcal{V}) , let S be an initial subtree of T with $\partial(S)$ finite, and let \mathcal{F} be a thick G -tribe of a graph Γ such that*

- (1) \mathcal{F} is concentrated at a half-grid-like end ϵ ;
- (2) \mathcal{F} strongly agrees about $\partial(S)$;
- (3) \mathcal{F} is well-separated from ϵ at S .

Suppose $f \in \partial_\epsilon(S)$ and let $\tilde{S} := S + f \subseteq T$. Then there is a thick flat subtribe \mathcal{F}^ of \mathcal{F} and a thick G -tribe \mathcal{F}' in Γ with the following properties:*

- (i) \mathcal{F}' is concentrated at ϵ .
- (ii) \mathcal{F}' strongly agrees about $\partial(\tilde{S})$.
- (iii) \mathcal{F}' is well-separated from ϵ at \tilde{S} .
- (iv) $\mathcal{F}' \cup \mathcal{F}$ strongly agrees about $\partial(S) \setminus \{f\}$.
- (v) $S^{-\epsilon}$ w.r.t. \mathcal{F} is a subtree of $\tilde{S}^{-\epsilon}$ w.r.t. \mathcal{F}' .
- (vi) For every $F \in \mathcal{F}^*$ and every $m \in \mathbb{N}$, there is an $F' \in \mathcal{F}'$ such that for all $H \in F$, there is an $H' \in F'$ which is a push-out of H to depth m along f .

Proof. For every member H of \mathcal{F} , consider a sequence $(H^{(i)} : i \in \mathbb{N})$, where $H^{(i)}$ is a push-out of H along f to depth at least i . After choosing a subsequence of $(H^{(i)} : i \in \mathbb{N})$ and relabelling (monotonically), we may assume that for each H , the set $\{H^{(i)} : i \in \mathbb{N}\}$ weakly agrees on $\partial(\tilde{S})$, i.e. for every $e \in \partial(\tilde{S})$ either $H^{(i)\downarrow}(R) \in \epsilon$ for every $R \in \omega_e$ and all i or $H^{(i)\downarrow}(R) \notin \epsilon$ for every $R \in \omega_e$ and all i . Note that a monotone relabelling preserves the property of $H^{(i)}$ being a push-out of H along f to depth at least i .

This uniform behaviour of $(H^{(i)} : i \in \mathbb{N})$ on $\partial(\tilde{S})$ for each member H of \mathcal{F} gives rise to a finite colouring $c : \bigcup \mathcal{F} \rightarrow 2^{\partial(\tilde{S})}$. By Lemma 7.3, we may choose a thick flat subtribe $\mathcal{F}_1 \subseteq \mathcal{F}$ such that c is constant on $\bigcup \mathcal{F}_1$.

Recall that, by Corollary 7.11, for every $e \in \partial_\epsilon(\tilde{S})$ (w.r.t. \mathcal{F}_1), the ray graph $\text{RG}_G(\mathcal{R}_e)$ is a path. We pick an arbitrary orientation of this path and denote by \leq_e the corresponding linear order on \mathcal{R}_e .

Note that, since \mathcal{F}_1 is a flat subtribe of \mathcal{F} which strongly agrees about $\partial(S)$, every ϵ -ray in every member $H \in \bigcup \mathcal{F}_1$ is core. Let us define, for each member $H \in \bigcup \mathcal{F}_1$,

$$d_H : \{H^{(i)} : i \in \mathbb{N}\} \rightarrow \{-1, 1\}^{\partial_\epsilon(\tilde{S})},$$

where

$$d_H(H^{(i)})_e = \begin{cases} 1 & \text{if } \leq_e \text{ agrees with the } \leq_e, \\ -1 & \text{if } \leq_e \text{ agrees with the reverse order } \geq_e \text{ of } \leq_e. \end{cases}$$

Since d_H has finite range, we may assume by Lemma 7.3, after choosing a subsequence and relabelling, that d_H is constant on $\{H^{(i)} : i \in \mathbb{N}\}$ and that $H^{(i)}$ is still a push-out of H along f to depth at least i .

Now, consider $d : \bigcup \mathcal{F}_1 \rightarrow \{-1, 1\}^{\partial_\epsilon(\tilde{S})}$, with $d(H) = d_H(H^{(1)}) (= d_H(H^{(i)}) \text{ for all } i \in \mathbb{N})$. Again, we may choose a thick flat subtribe $\mathcal{F}_2 \subseteq \mathcal{F}_1$ such that d is constant on \mathcal{F}_2 .

Since \mathcal{F} is well-separated from ϵ at S , we get that $\{H^{(i)} : H \in \mathcal{F}\}$ is well-separated from ϵ at S . So, we can now apply Lemma 7.15 to each $H^{(i)}$ to obtain $H'^{(i)}$, yielding a collection which is well-separated from ϵ at \tilde{S} . Note that $H'^{(i)}$ is still a push-out of H along f to depth at least i .

Now, let $\mathcal{F}^* = \mathcal{F}_2$ and $\mathcal{F}' = \{\{H'^{(i)} : H \in F\} : i \in \mathbb{N}, F \in \mathcal{F}^*\}$. Let us verify that these satisfy (i)–(vi). \mathcal{F}^* is concentrated at ϵ because it is a thick flat subtribe of \mathcal{F} by Lemma 7.5. By a comparison, layer by layer, since all members of \mathcal{F}' are push-outs of members of \mathcal{F}^* along f , the tribe \mathcal{F}' is also concentrated at ϵ , satisfying (i).

Property (ii) is satisfied: Since c and d are constant on $\bigcup \mathcal{F}_2$ the collection of the $H^{(i)}$ (for $H \in \bigcup \mathcal{F}_2$) strongly agrees on $\partial(\tilde{S})$, since we have chosen an appropriate subsequence in which $d_H(H^{(i)})$ is constant. The $H'^{(i)}$ are constructed such that this property is preserved. Property (iii) is immediate from the choice of $H'^{(i)}$. Properties (iv) and (v) follow from (2) and the fact that every member of \mathcal{F}' is a push-out of a member of \mathcal{F} along f . Property (vi) is immediate from the construction of \mathcal{F}' . \square

8. The inductive argument

In this section we prove Theorem 4.6, our main result. Given a locally finite connected graph G which admits an extensive tree-decomposition (T, \mathcal{V}) and a graph Γ which contains a thick G -tribe \mathcal{F} , our aim is to construct an infinite family $(Q_i : i \in \mathbb{N})$ of disjoint G -minors in Γ inductively.

Our work so far will allow us to make certain assumptions about \mathcal{F} . For example, by Lemma 7.4, we may assume that \mathcal{F} is concentrated at some end ϵ of Γ , which, by Lemma 7.6, we may assume is a thick end, and, by Lemma 6.2, we may assume is not pebbly. Hence, by Theorem 6.1, we may assume that ϵ is either half-grid-like or grid-like.

At this point our proof will split into two different cases, depending on the nature of ϵ . As we mentioned before, the two cases are very similar, with the grid-like case being significantly simpler. Therefore, we will first prove Theorem 4.6 in the case where ϵ is half-grid-like, and then in Section 8.2 we will briefly sketch the differences for the grid-like case.

So, to briefly recap, in the following section we will be working under the standing assumptions that there is a thick G -tribe \mathcal{F} in Γ , and an end ϵ of Γ such that

- \mathcal{F} is concentrated at ϵ ;
- ϵ is thick;
- ϵ is half-grid-like.

8.1. The half-grid-like case

As explained in Section 2, our strategy will be to take some sequence of initial subtrees $S_1 \subseteq S_2 \subseteq S_3 \dots$ of T such that $\bigcup_{i \in \mathbb{N}} S_i = T$, and to inductively build a collection of n inflated copies of $G(S_n)$, at each stage extending the previous copies. However, in order to ensure that we can continue the construction at each stage, we will require the existence of additional structure.

Let us pick an enumeration $\{t_i : i \geq 0\}$ of $V(T)$ such that t_0 is the root of T and $T_n := T[\{t_i : 0 \leq i \leq n\}]$ is connected for every $n \in \mathbb{N}$. We will not take the S_n above to be the subtrees T_n , but instead the subtrees $T_n^{\neg\epsilon}$ with respect to some tribe \mathcal{F}_n that weakly agrees about $\partial(T_n)$. This will ensure that every edge in the boundary $\partial(S_n)$ will be in $\partial_\epsilon(T_n)$. For every edge $e \in E(T)$, let us fix a family $\mathcal{R}_e = (R_{e,s} : s \in S(e))$ of disjoint rays witnessing the self-similarity of the bough $B(e)$ towards an end ω_e of G , where $\text{init}(R_{e,s}) = s$. By taking $S_n = T_n^{\neg\epsilon}$, we guarantee that for each edge in $e \in \partial(S_n)$, $s \in S(e)$, and every $H \in \bigcup \mathcal{F}_n$, the ray $H^\perp(R_{e,s})$ is an ϵ -ray.

Furthermore, since $\partial(T_n)$ is finite, we may assume by Lemma 7.12, that \mathcal{F}_n strongly agrees about $\partial(T_n)$. We can now describe the additional structure that we require for the induction hypothesis.

At each stage of our construction we will have built some inflated copies of $G(S_n)$, which we wish to extend in the next stage. However, S_n will not in general be a finite subtree, and so we will need some control over where these copies lie in Γ to ensure we have not ‘used up’ all of Γ . The control we will want is that there is a finite set of vertices X , which we call a *bounder*, that

separates all that we have built so far from the end ϵ . This will guarantee, since \mathcal{F} is concentrated at ϵ , that we can find arbitrarily large layers of \mathcal{F} which are disjoint from what we have built so far.

Furthermore, in order to extend these copies in the next step, we will need to be able to link the boundary of our inflated copies of $G(S_n)$ to this large layer of \mathcal{F} . To this end we, will also want to keep track of some structure which allows us to do this, which we call an *extender*. Let us make the preceding discussion precise.

Definition 8.1 (Bounder, extender). Let \mathcal{F} be a thick G -tribe, which is concentrated at ϵ and strongly agrees about $\partial(S)$ for some initial subtree S of T , and let $k \in \mathbb{N}$. Let $\mathcal{Q} = (Q_i : i \in [k])$ be a family of disjoint inflated copies of $G(S^{-\epsilon})$ in Γ (note, $S^{-\epsilon}$ depends on \mathcal{F}).

- A *bounder* for \mathcal{Q} is a finite set X of vertices in Γ separating each Q_i in \mathcal{Q} from ϵ , i.e. such that

$$C(X, \epsilon) \cap \bigcup_{i=1}^k Q_i = \emptyset.$$

- For $A \subseteq E(T)$, let $I(A, k)$ denote the set $\{(e, s, i) : e \in A, s \in S(e), i \in [k]\}$.
- An *extender* for \mathcal{Q} is a family $\mathcal{E} = (E_{e,s,i} : (e, s, i) \in I(\partial_\epsilon(S), k))$ of ϵ -rays in Γ such that the graphs in $\mathcal{E}^- \cup \mathcal{Q}$ are pairwise disjoint and such that $\text{init}(E_{e,s,i}) \in Q_i(s)$ for every $(e, s, i) \in I(\partial_\epsilon(S), k)$ (using the notation as in Definition 3.3).
- Given an extender \mathcal{E} , an edge $e \in \partial_\epsilon(S)$, and $i \in [k]$, we let

$$\mathcal{E}_{e,i} := (E_{e,s,i} : s \in S(e)).$$

Recall that, since ϵ is half-grid like, there is a partial order \leq_ϵ defined on the core rays of ϵ , see Lemma 6.13. Furthermore, if \mathcal{F} strongly agrees about $\partial(S)$ then, as in Definition 7.7, for each $e \in \partial_\epsilon(S)$, there is a linear order $\leq_{\mathcal{F},e}$ on $S(e)$.

Definition 8.2 (Extension scheme). Under the conditions above, we call a tuple (X, \mathcal{E}) an *extension scheme* for \mathcal{Q} if the following holds:

- (ES1) X is a bounder for \mathcal{Q} and \mathcal{E} is an extender for \mathcal{Q} ;
- (ES2) \mathcal{E} is a family of core rays;
- (ES3) the order \leq_ϵ on $\mathcal{E}_{e,i}$ (and thus on $\mathcal{E}_{e,i}^-$) agrees with the order induced by $\leq_{\mathcal{F},e}$ on $\mathcal{E}_{e,i}^-$ for all $e \in \partial_\epsilon(S)$ and $i \in [k]$;
- (ES4) the sets $\mathcal{E}_{e,i}^-$ are intervals with respect to \leq_ϵ on \mathcal{E}^- for all $e \in \partial_\epsilon(S)$ and $i \in [k]$.

We will in fact split our inductive construction into two types of extensions, which we will do on odd and even steps respectively.

In an even step $n = 2k$, starting with a G -tribe \mathcal{F}_k , k disjoint inflated copies $(Q_i^n : i \in [k])$ of $G(T_k^{-\epsilon})$, and an appropriate extension scheme, we will construct Q_{k+1}^n , a further disjoint inflated copy of $G(T_k^{-\epsilon})$, and an appropriate extension scheme for everything we built so far.

In an odd step $n = 2k - 1$ (for $k \geq 1$), starting with the same G -tribe \mathcal{F}_{k-1} from the previous step, k disjoint inflated copies of $G(T_{k-1}^{-\epsilon})$, and an appropriate extension scheme, we will refine to a new G -tribe \mathcal{F}_k , which strongly agrees on $\partial(T_k)$, extend each copy Q_i^n of $G(T_{k-1}^{-\epsilon})$ to a copy Q_i^{n+1} of $G(T_k^{-\epsilon})$ for $i \in [k]$, and construct an appropriate extension scheme for everything we built so far.

So, we will assume inductively that for some $n \in \mathbb{N}_0$, with $\rho := \lfloor n/2 \rfloor$ and $\sigma := \lceil n/2 \rceil$ we have:

(I1) a thick G -tribe \mathcal{F}_ρ in Γ which

- is concentrated at ϵ ;
- strongly agrees about $\partial(T_\rho)$;
- is well-separated from ϵ at T_ρ ;
- whenever $k < l \leq \rho$, the tree $T_k^{-\epsilon}$ with respect to \mathcal{F}_k is a subtree of $T_l^{-\epsilon}$ with respect to \mathcal{F}_l .

(I2) a family $\mathcal{Q}_n = (Q_i^n : i \in [\sigma])$ of σ pairwise disjoint inflated copies of $G(T_\rho^{-\epsilon})$ (where $T_\rho^{-\epsilon}$ is considered with respect to \mathcal{F}_ρ) in Γ ;

if $n \geq 1$, we additionally require that Q_i^n extends Q_i^{n-1} for all $i \leq \sigma - 1$;

(I3) an extension scheme (X_n, \mathcal{E}_n) for \mathcal{Q}_n ;

(I4) if n is even and $\partial_\epsilon(T_\rho) \neq \emptyset$, we require that there is a set \mathcal{J}_ρ of disjoint core ϵ -rays disjoint to \mathcal{E}_n , with $|\mathcal{J}_\rho| \geq (|\partial_\epsilon(T_\rho)| + 1) \cdot |\mathcal{E}_n|$.

We illustrate the combination of the construction steps in Figure 8.1

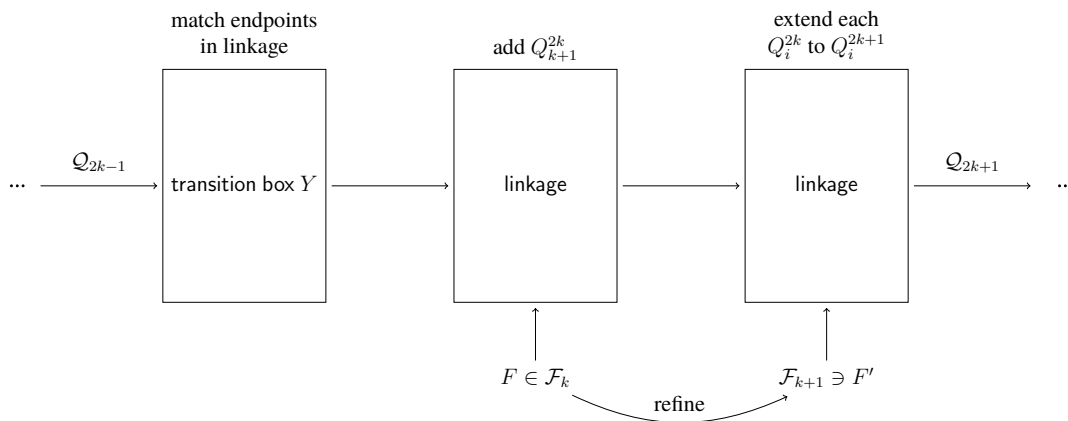


Figure 8.1: The combination of the steps in our construction.

Suppose we have inductively constructed \mathcal{Q}_n for all $n \in \mathbb{N}$. Let us define $H_i := \bigcup_{n \geq 2i-1} \mathcal{Q}_i^n$. Since $T_k^{-\epsilon}$ with respect to \mathcal{F}_k is a subtree of $T_l^{-\epsilon}$ with respect to \mathcal{F}_l for all $k < l$, we have that $\bigcup_{n \in \mathbb{N}} T_n^{-\epsilon} = T$ (where we considered $T_n^{-\epsilon}$ w.r.t. \mathcal{F}_n), and due to the extension property (I2), the collection $(H_i : i \in \mathbb{N})$ is an infinite family of disjoint G -minors, as required.

So let us start the construction. To see that our assumptions can be fulfilled for the case $n = 0$, we first note that since $T_0 = t_0$, by Lemma 7.14 there is a thick subtribe \mathcal{F}_0 of \mathcal{F} which satisfies (II). Let us further take $\mathcal{Q}_0 = \mathcal{E}_0 = X_0 = \mathcal{J}_0 = \emptyset$.

The following notation will be useful throughout the construction. Given $e \in E(T)$ and some inflated copy H of G , recall that $H^\downarrow(\mathcal{R}_e)$ denotes the family $(H^\downarrow(R_{e,s}) : s \in S(e))$. Given a G -tribe \mathcal{F} , a layer $F \in \mathcal{F}$ and a family of disjoint rays \mathcal{R} in G we will write $F^\downarrow(\mathcal{R}) = (H^\downarrow(R) : H \in F, R \in \mathcal{R})$.

Construction part 1: $n = 2k$ is even

Case 1: $\partial_\epsilon(T_k) = \emptyset$.

In this case, $T_k^{-\epsilon} = T$ and so picking any member $H \in \mathcal{F}_k$ with $H \subseteq C(X_n, \epsilon)$ and setting $\mathcal{Q}_{k+1}^{n+1} = H(T_k^{-\epsilon})$ gives us a further inflated copy of $G(T_k^{-\epsilon})$ disjoint from all the previous ones. We set $\mathcal{Q}_i^{n+1} = \mathcal{Q}_i^n$ for all $i \in [k]$ and $\mathcal{Q}_{n+1} = (\mathcal{Q}_i^{n+1} : i \in [k+1])$. Since \mathcal{F}_k is well-separated from ϵ at T_k , there is a suitable bounder $X_{n+1} \supseteq X_n$ for \mathcal{Q}_{n+1} . Then (X_{n+1}, \emptyset) is an extension scheme for \mathcal{Q}_{n+1} while \mathcal{F}_k remains unchanged.

Case 2: $\partial_\epsilon(T_k) \neq \emptyset$. (See Figure 8.2)

Consider the family $\mathcal{R}^- := \bigcup \{\mathcal{R}_e^- : e \in \partial_\epsilon(T_k)\}$. Moreover, set $\mathcal{C} := \mathcal{E}_n^- \cup \mathcal{J}_k$ and consider $\bar{\mathcal{C}}$ as in Definition 6.16. Let $Y \subseteq C(X_n, \epsilon)$ be a finite subgraph, which is a transition box between $\bar{\mathcal{E}}_n^-$ and $\bar{\mathcal{C}}$ after X_n as in Lemma 3.17. Let \mathcal{F}' be a flat thick G -subtribe of \mathcal{F}_k , such that each member of \mathcal{F}' is contained in $C(X_n \cup V(Y), \epsilon)$, which exists, by Lemma 7.5, since both X_n and $V(Y)$ are finite.

Let $F \in \mathcal{F}'$ be large enough such that we may apply Lemma 3.16 to find a transitional linkage \mathcal{P} , such that $\bigcup \mathcal{P} \subseteq C(X_n \cup V(Y), \epsilon)$, from $\bar{\mathcal{C}}$ to $F^\downarrow(\mathcal{R}^-)$ after $X_n \cup V(Y)$ avoiding some member $H \in F$. Note that, since X_n is a bounder and $\bigcup \mathcal{P} \subseteq C(X_n \cup V(Y), \epsilon)$, we get that each element of \mathcal{P} is disjoint from all \mathcal{Q}_n and Y .

Let

$$\mathcal{Q}_{k+1}^{n+1} := H(T_k^{-\epsilon}).$$

Note that \mathcal{Q}_{k+1}^{n+1} is an inflated copy of $G(T_k^{-\epsilon})$. Moreover, let $\mathcal{Q}_i^{n+1} := \mathcal{Q}_i^n$ for all $i \in [k]$ and $\mathcal{Q}_{n+1} := (\mathcal{Q}_i^{n+1} : i \in [k+1])$, yielding property (I2).

Since \mathcal{F}_k is well-separated from ϵ at T_k , and $H \in \bigcup \mathcal{F}_k$, there is a finite set $X_{n+1} \subseteq V(\Gamma)$ containing $X_n \cup V(Y)$, such that $C(X_{n+1}, \epsilon) \cap \mathcal{Q}_{k+1}^{n+1} = \emptyset$. This set X_{n+1} is a bounder for \mathcal{Q}_{n+1} .

Since \mathcal{P} is transitional, Lemma 6.19 implies that the linkage is preserving on \mathcal{C} . Since all rays in $F^\downarrow(\mathcal{R}^-)$ are core rays, we have that \leq_ϵ is a linear order on $F^\downarrow(\mathcal{R}^-)$. Moreover, for each $e \in \partial_\epsilon(T_k)$, the rays in $H^\downarrow(\mathcal{R}_e^-)$ correspond to an interval in this order. Thus, deleting these intervals from $F^\downarrow(\mathcal{R}^-)$ leaves behind at most $|\partial_\epsilon(T_k)| + 1$ intervals in $F^\downarrow(\mathcal{R}^-)$ (with respect to \leq_ϵ) which do not contain any rays in $H^\downarrow(\mathcal{R}^-)$. Since $|\mathcal{J}_k| \geq (|\partial_\epsilon(T_k)| + 1) \cdot |\mathcal{E}_n|$, by the pigeonhole principle there is one such interval on $F^\downarrow(\mathcal{R}^-)$ that

- does not contain rays in $H^\downarrow(\mathcal{R})$;
- where a subset $\mathcal{P}' \subseteq \mathcal{P}$ of size $|\mathcal{E}_n^-|$ links a corresponding subset \mathcal{A} of \mathcal{C} to a set of rays \mathcal{B} in that interval.

By Lemmas 3.17, 6.23 and 6.19, and Remark 6.221, there is a linkage \mathcal{P}'' from $\overline{\mathcal{E}_n^-}$ to \mathcal{A} contained in Y which is preserving on \mathcal{E}_n^- .

For $e \in \partial_\epsilon(T_k)$ and $s \in S(e)$, define

$$E_{e,s,k+1}^{n+1} = H^\downarrow(R_{e,s}) \text{ for the corresponding ray } R_{e,s} \in \mathcal{R}_e.$$

Moreover for each $i \in [k]$, we define

$$E_{e,s,i}^{n+1} = (E_{e,s,i}^n \circ_{\mathcal{P}''} \mathcal{A}) \circ_{\mathcal{P}'} \mathcal{B},$$

noting that \mathcal{P}'' is also a linkage from \mathcal{E}_n to \mathcal{A} .

By construction, all these rays are, except for their first vertex, disjoint from \mathcal{Q}_{n+1} . Moreover, $\mathcal{E}_{n+1} := (E_{e,s,i}^{n+1} : (e, s, i) \in I(\partial_\epsilon(T_k), k+1))$ is an extender for \mathcal{Q}_{n+1} . Note that each ray in \mathcal{E}_{n+1} shares a tail with a ray in $F^\downarrow(\mathcal{R}^-)$.

We claim that $(X_{n+1}, \mathcal{E}_{n+1})$ is an extension scheme for \mathcal{Q}_{n+1} and hence property (I3) is satisfied. Since every ray in \mathcal{E}_{n+1} has a tail which is also a tail of a ray in $F^\downarrow(\mathcal{R}^-)$, property (ES2) is satisfied by Remark 6.141. Since \mathcal{P}' is preserving on \mathcal{A}' and \mathcal{P}'' is preserving on \mathcal{E}_n^- , Remark 6.222 implies that the linkage $\mathcal{P}'' + \mathcal{P}'$ is preserving on \mathcal{E}_n^- . Hence, property (ES3) holds for each $i \in [k]$. Furthermore, since $E_{e,s,k+1}^{n+1} = H^\downarrow(R_{e,s})$ for each $e \in \partial_\epsilon(T_k)$ and $s \in S(e)$ and \mathcal{F}_k strongly agrees about $\partial(T_k)$, it is clear that property (ES3) holds for $i = k+1$. Finally, property (ES4) holds for $i = k+1$ since for each $e \in \partial_\epsilon(T_k)$, the rays in $H^\downarrow(\mathcal{R}_e)$ are an interval with respect to \leq_ϵ on $F^\downarrow(\mathcal{R}^-)$, and it holds for $i \in [k]$ by the fact that $\mathcal{P}'' + \mathcal{P}'$ is preserving on \mathcal{E}_n^- together with the fact that $\mathcal{P}'' + \mathcal{P}'$ links \mathcal{E}_n^- to an interval of $F^\downarrow(\mathcal{R}^-)$ containing no ray in $H^\downarrow(\mathcal{R})$.

Finally, note that (II) is still satisfied by \mathcal{F}_k and T_k , and (I4) is vacuously satisfied.

Construction part 2: $n = 2k - 1$ is odd (for $k \geq 1$).

Let f denote the unique edge of T between T_{k-1} and $T_k \setminus T_{k-1}$.

Case 1: $f \notin \partial_\epsilon(T_{k-1})$.

Let $\mathcal{F}_k := \mathcal{F}_{k-1}$. Since \mathcal{F}_{k-1} is well-separated from ϵ at T_{k-1} , it follows that $e \in \partial_{-\epsilon}(T_k)$ for every $e \in \partial(T_k) \setminus \partial(T_{k-1})$. Hence $T_k^{-\epsilon} = T_{k-1}^{-\epsilon}$ and $\partial_\epsilon(T_{k-1}) = \partial_\epsilon(T_k)$, and so \mathcal{F}_k is well-separated from ϵ at T_k and we can simply take $\mathcal{Q}_{n+1} := \mathcal{Q}_n$, $\mathcal{E}_{n+1} := \mathcal{E}_n$, $\mathcal{J}_k := \mathcal{J}_{k-1}$ and $X_{n+1} := X_n$ to satisfy (II), (I2), (I3) and (I4).

Case 2: $f \in \partial_\epsilon(T_{k-1})$. (See Figure 8.3)

By (II) we can apply Lemma 7.16 to \mathcal{F}_{k-1} and T_{k-1} in order to find a thick G -tribe \mathcal{F}_k and a thick flat subtribe \mathcal{F}^* of \mathcal{F}_{k-1} , both concentrated at ϵ , satisfying properties (i)–(vi) from that lemma. It follows that \mathcal{F}_k satisfies (II) for the next step.

Let $F \in \mathcal{F}^*$ be a layer of \mathcal{F}^* such that

$$|F| \geq (\partial_\epsilon(T_k) + 2) \cdot |I(\partial_\epsilon(T_k), k)|$$

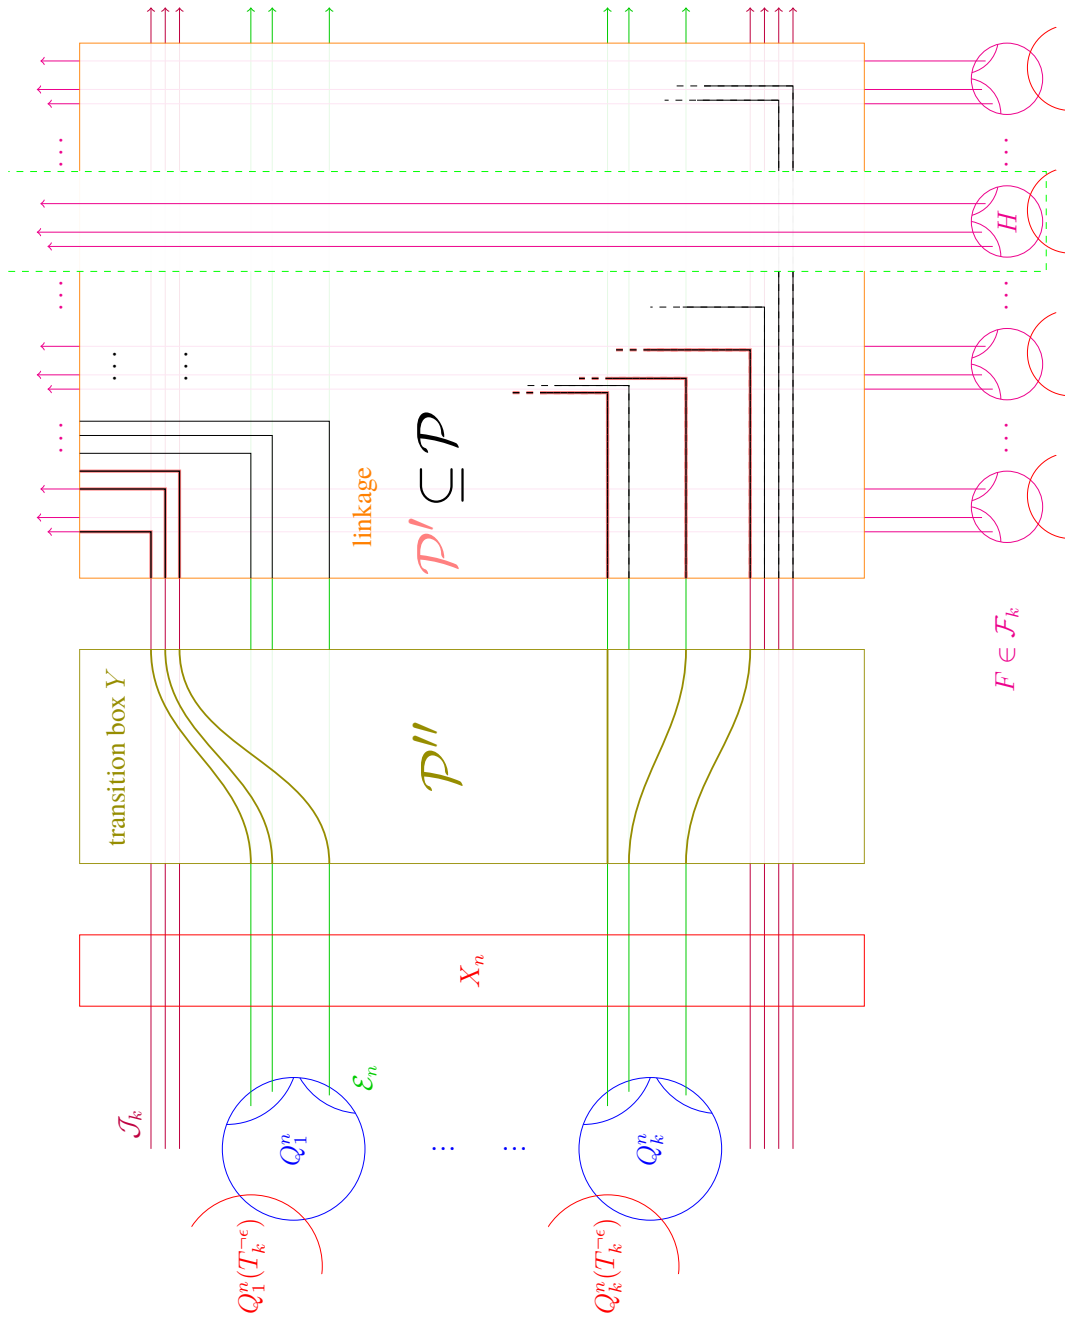


Figure 8.2: Adding a new copy when $n = 2k$ is even.

and consider the rays $F^\downarrow(\mathcal{R}_f)$. Consider the rays in the extender corresponding to the edge f , that is $\mathcal{E}_f := (E_{f,s,i}^n : i \in [k], s \in S(f))$. By Lemma 6.23, there is, for every subset \mathcal{S} of $F^\downarrow(\mathcal{R}_f)$ of size $|\mathcal{E}_f^-|$, a transitional linkage \mathcal{P} from $\mathcal{E}_f^- \subseteq \overline{\mathcal{E}_n^-}$ to $\mathcal{S} \subseteq \overline{F^\downarrow(\mathcal{R}_f)}$ after $X_n \cup \text{init}(\mathcal{E}_n)$, which is preserving on \mathcal{E}_f^- .

Let us choose $H_1, H_2, \dots, H_k \in F$ and let $\mathcal{S} = (H_i^\downarrow(R_{f,s}) : i \in [k], s \in S(f))$. Let \mathcal{P} be the linkage given by the previous paragraph, which we recall is preserving on \mathcal{E}_f^- . Since for every $i \leq k$, the family $(E_{f,s,i}^{n-} : s \in S(f))$ forms an interval in \mathcal{E}_n^- and the set $H_i^\downarrow(\mathcal{R}_f)$ forms an interval in $F^\downarrow(\mathcal{R}_f)$, and furthermore the order \leq_ϵ agrees with $\leq_{\mathcal{F}_k, f}$ on $S(f)$, it follows that, after perhaps relabelling the H_i , for every $i \in [k]$ and $s \in S(f)$, \mathcal{P} links $E_{f,s,i}^{n-}$ to $H_i^\downarrow(R_{f,s})$.

Let $Z \subseteq V(\Gamma)$ be a finite set such that $\top(\omega, R)$ and $\perp(\omega, R)$ are separated by Z in $\Gamma - V(R)$ for all $R \in F^\downarrow(\mathcal{R}_f)$ (cf. Lemma 6.15).

Since $|F|$ is finite and (T, \mathcal{V}) is an extensive tree-decomposition, there exists an $m \in \mathbb{N}$ such that if $e \in T_{f+}$ with $\text{dist}(f^-, e^-) = m$, then $H(B(e)) \cap (X_n \cup Z \cup V(\bigcup \mathcal{P})) = \emptyset$ for every $H \in F$. Let $F' \in \mathcal{F}_k$ be as in Lemma 7.16(vi) for F with such an m .

Hence, by definition, for each $H_i \in F$ there is some $H'_i \in F'$ which is a push-out of H_i to depth m along f , and so there is some edge $e \in T_{f+}$ with $\text{dist}(f^-, e^-) = m$ and some subgraph $W_i \subseteq H(B(e))$ which is an $\overline{IG[B(f)]}$ such that for each $s \in S(f)$, we have that $W_i(s)$ contains the first vertex of W_i on $H_i^\downarrow(R_{f,s})$.

For each $i \in [k]$ we construct Q_i^{n+1} from Q_i^n as follows. Consider the part of G that we want to add $G(T_{k-1}^{-\epsilon})$ to obtain $G(T_k^{-\epsilon})$, namely

$$D := \overline{G[B(f)]} \left[V_{f+} \cup \bigcup \{B(e) : e \in \partial_{-\epsilon}(T_k) \setminus \partial_{-\epsilon}(T_{k-1})\} \right].$$

Let $K_i := W_i(D)$. Note that this is an inflated copy of D , and for each $s \in S(f)$ and each $i \in [k]$ the branch set $K_i(s)$ contains the first vertex of K_i on $H_i^\downarrow(R_{f,s})$.

Note further that, by the choice of m , all the K_i are disjoint to \mathcal{Q}_n . Let $x_{f,s,i}$ denote the first vertex on the ray $H_i^\downarrow(R_{f,s})$ in K_i , and let

$$O_{s,i} := (E_{f,s,i}^n \circ_{\mathcal{P}} F^\downarrow(\mathcal{R}_f))x_{f,s,i},$$

where as before we note that \mathcal{P} is also a linkage from \mathcal{E}_n to $F^\downarrow(\mathcal{R}_f)$.

Then, if we let $\mathcal{O}_i := (O_{s,i} : s \in S(f))$ and $\mathcal{O} = (O_{s,i} : s \in S(f), i \in [k])$, we see that

$$Q_i^{n+1} := Q_i^n \oplus_{\mathcal{O}_i} K_i$$

(see Definition 4.7) is an inflated copy of $G(T_k^{-\epsilon})$ extending Q_i^n . Hence,

$$\mathcal{Q}^{n+1} := (Q_i^{n+1} : i \in [k])$$

is a family satisfying (I2).

Since \mathcal{F}_k is well-separated from ϵ at T_k , and each K_i is a subgraph of the restriction of $W_i \subseteq H'_i$ to D , for each K_i , there is a finite set \hat{X}_i separating K_i from ϵ , and hence the set

$$X_{n+1} := X_n \cup \bigcup_{i \in [k]} \hat{X}_i \cup V\left(\bigcup \mathcal{O}\right)$$

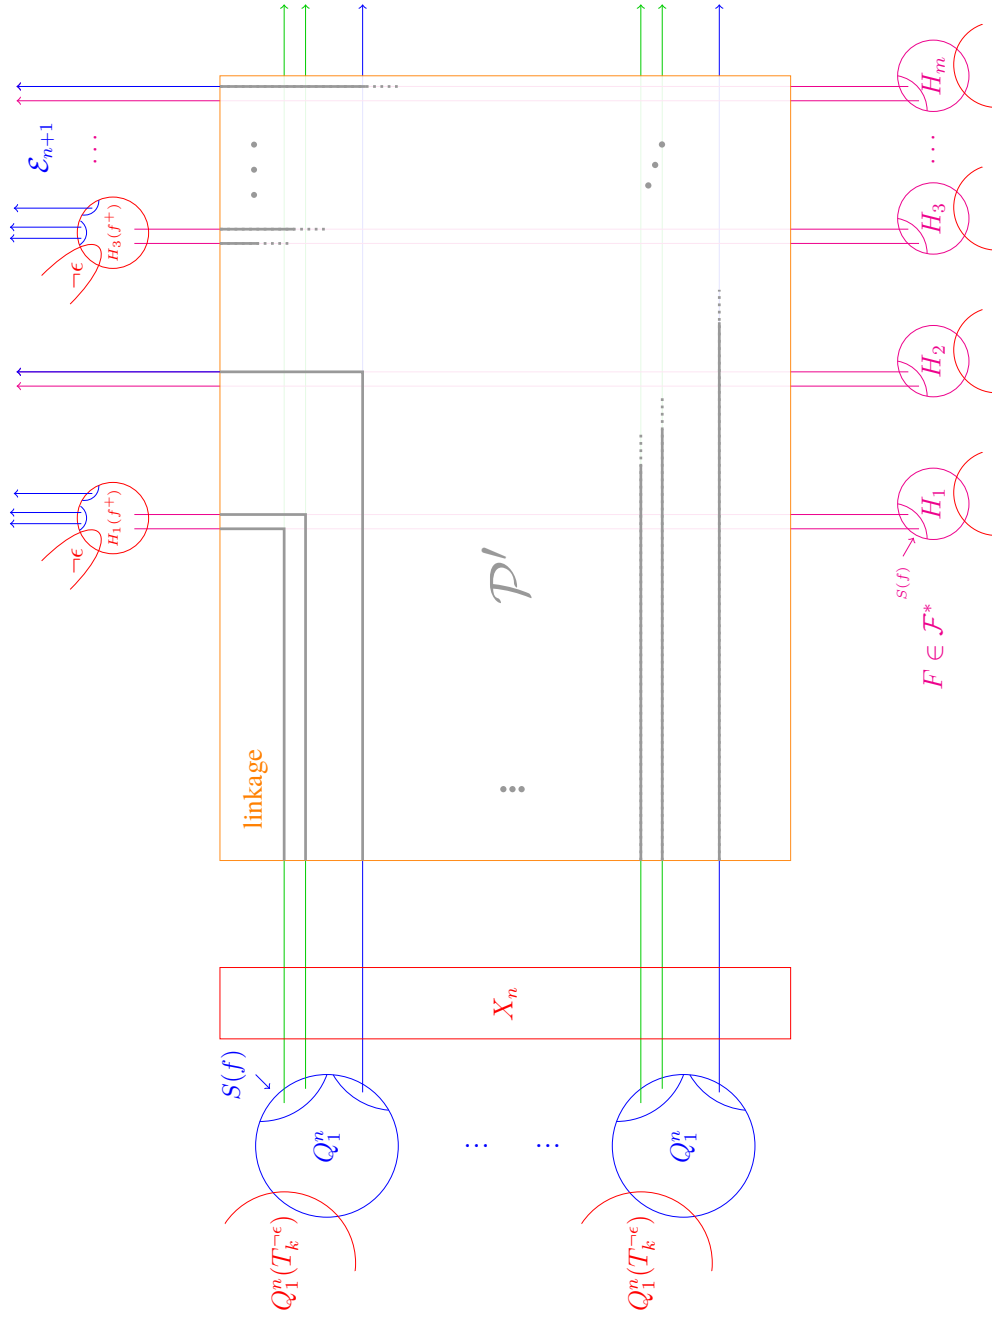


Figure 8.3: Extending the copies when $n = 2k - 1$ is odd.

is a bounder for \mathcal{Q}^{n+1} .

For $e \in \partial_\epsilon(T_{k-1}) \setminus \{f\}$, $s \in S(e)$, and $i \in [k]$, we set

$$E_{e,s,i}^{n+1} = E_{e,s,i}^n \circ_{\mathcal{P}} F^\downarrow(\mathcal{R}_f),$$

and set

$$\mathcal{E}' := (E_{e,s,i}^{n+1} : (e, s, i) \in I(\partial_\epsilon(T_{k-1}) \setminus \{f\}, k))$$

Moreover, for $e \in \partial_\epsilon(T_k) \setminus \partial_\epsilon(T_{k-1})$, $s \in S(e)$, and $i \in [k]$, we set

$$E_{e,s,i}^{n+1} = H_i^{\downarrow}(R_{e,s}),$$

and set

$$\mathcal{E}'' := (E_{e,s,i}^{n+1} : (e, s, i) \in I(\partial_\epsilon(T_k) \setminus \partial_\epsilon(T_{k-1}), k)).$$

Note that, by construction, any such ray $E_{e,s,i}^{n+1}$ has its initial vertex in the branch set $Q_i^{n+1}(s)$ and is otherwise disjoint to $\bigcup \mathcal{Q}_{n+1}$. We set $\mathcal{E}_{n+1} := \mathcal{E}' \cup \mathcal{E}''$. It is easy to check that this is an extender for \mathcal{Q}_{n+1} .

We claim that $(X_{n+1}, \mathcal{E}_{n+1})$ is an extension scheme. Property (ES1) is apparent. Since \mathcal{F}_k strongly agrees about $\partial(T_k)$, every ϵ -ray in an any member of \mathcal{F}_k is core. Then, since \mathcal{F}^* is a flat subtribe of \mathcal{F}_k and every ray in \mathcal{E}_{n+1} shares a tail with a ray in a member of \mathcal{F}_k or \mathcal{F}^* , it follows by Remark 6.14 1 that all rays in \mathcal{E}_{n+1} are core rays, and so (ES2) holds.

For any $e \in \partial_\epsilon(T_{k-1}) \setminus \{f\}$ and $i \in [k]$, the rays $\mathcal{E}_{n+1,e,i}$ are a subfamily of \mathcal{E}' , obtained by transitioning from the family $\mathcal{E}_{n,e,i}$ to $F^\downarrow(\mathcal{R}_f)$ along the linkage \mathcal{P} . By the induction hypothesis, \leq_ϵ agreed with the order induced by $\leq_{\mathcal{F}_{k-1},e}$ on $\mathcal{E}_{n,e,i}$, and, since $\mathcal{F}_k \cup \mathcal{F}_{k-1}$ strongly agrees about $\partial_\epsilon(T_{k-1}) \setminus \{f\}$, this is also the order induced by $\leq_{\mathcal{F}_k,e}$. Hence, since \mathcal{P} is preserving, by Lemma 6.19, it follows that the order induced by $\leq_{\mathcal{F}_k,e}$ on $\mathcal{E}_{n+1,e,i}$ agrees with \leq_ϵ .

For for $e \in \partial_\epsilon(T_k) \setminus \partial_\epsilon(T_{k-1})$ and $i \in [k]$, the rays $\mathcal{E}_{n+1,e,i}$ are $(H_i^{\downarrow}(R_{e,s}) : s \in S(e))$. Since $H_i' \in F' \in \mathcal{F}_k$ and \mathcal{F}_k strongly agrees about $\partial(T_k)$, it follows that the order induced by $\leq_{\mathcal{F}_k,e}$ on $\mathcal{E}_{n+1,e,i}$ agrees with \leq_ϵ . Hence Property (ES3) holds.

Finally, by Lemma 3.20 it is clear that for any $e \in \partial_\epsilon(T_{k-1}) \setminus \{f\}$ and $i \in [k]$, the rays $\mathcal{E}_{n+1,e,i}^-$ form an interval with respect to \leq_ϵ on \mathcal{E}_{n+1}^- , since they are each contained in a connected subgraph H_i' to which the tails of the rest of \mathcal{E}_{n+1}^- are disjoint. Furthermore, by choice of Z and Lemma 6.15, it is clear that, since \mathcal{P} is preserving on \mathcal{E}_n^- , for each $e \in \partial_\epsilon(T_k) \setminus \partial_\epsilon(T_{k-1})$ and $i \in [k]$, the rays $\mathcal{E}_{n+1,e,i}^-$ also form an interval with respect to \leq_ϵ on \mathcal{E}_{n+1}^- . Hence, property (ES4) holds and therefore (I3) is satisfied for the next step.

For property (I4), we note that every ray in \mathcal{E}_{n+1} has a tail in some $H \in F \in \mathcal{F}^*$ or some pushout H' of H in \mathcal{F}_k . Note that $V(H') \subseteq V(H)$. Since there is at least one core ϵ -ray in each $H \in F \in \mathcal{F}^*$, and the H in F are pairwise disjoint, we can find a family of at least $|F| - |\mathcal{E}_{n+1}|$ such rays disjoint from \mathcal{E}_{n+1} . However, since

$$|F| \geq (\partial_\epsilon(T_k) + 2) \cdot |\mathcal{E}_{n+1}|,$$

it follows that we can find a suitable family $|\mathcal{J}_k|$.

This concludes the induction step. □

8.2. The grid-like case

In this section we will give a brief sketch of how the argument differs in the case where the end ϵ , towards which we may assume our G -tribe \mathcal{F} is concentrated, is grid-like.

In the case where ϵ is half-grid-like we showed that the end ϵ had a roughly linear structure, in the sense that there is a global partial order \leq_ϵ which is defined on almost all of the ϵ -rays, namely the core ones, such that every pair of disjoint core rays are comparable, and that this order determines the relative structure of any finite family of disjoint core rays, since it determines the ray graph.

Since, by Corollary 7.11, $\text{RG}_G(\mathcal{R}_e)$ is a path whenever $e \in \partial_\epsilon(T_k)$, there are only two ways that \leq_ϵ can order $H^\downarrow(\mathcal{R}_e)$, and, since $\partial_\epsilon(T_k)$ is finite, by various pigeon-hole type arguments we can assume that it does so consistently for each $H \in \bigcup \mathcal{F}_k$ and each $\mathcal{E}_{e,i}$.

We use this fact crucially in part 2 of the construction, where we wish to extend the graphs $(Q_i^n : i \in [k])$ from inflated copies of $G(T_{k-1}^{-\epsilon})$ to inflated copies of $G(T_k^{-\epsilon})$ along an edge $e \in \partial(T_{k-1})$. We wish to do so by constructing a linkage from the extender \mathcal{E}_n to some layer $F \in \mathcal{F}_k$, using the self-similarity of G to find an inflated copy of $G[B(e)]$ which is ‘rooted’ on the rays $H^\downarrow(\mathcal{R}_e)$ and extending each Q_i^n by such a subgraph.

However, for this step to work it is necessary that the linkage from \mathcal{E}_n to $F^\downarrow(\mathcal{R}_e)$ is such that for each $i \in [k]$, there is some $H \in F$ such that ray $E_{e,s,i}$ is linked to $H^\downarrow(\mathcal{R}_{e,s})$ for each $s \in S(e)$. However, since any transitional linkage we construct between \mathcal{E}_n and a layer $F \in \mathcal{F}_n$ will respect \leq_ϵ , we can use a transition box to ‘re-route’ our linkage such that the above property holds.

In the case where ϵ is grid-like we would like to say that the end has a roughly cyclic structure, in the sense that there is a global ‘partial cyclic order’ C_ϵ , defined again on almost all of the ϵ -rays, which will again determine the relative structure of any finite family of disjoint ‘core’ rays.

As before, since $\text{RG}_G(\mathcal{R}_e)$ is a path whenever $e \in \partial_\epsilon(T_n)$, there are only two ways that C_ϵ can order $H^\downarrow(\mathcal{R}_e)$ (‘clockwise’ or ‘anti-clockwise’) and so we can use similar arguments to assume that it does so consistently for each $H \in \bigcup \mathcal{F}_k$ and each $\mathcal{E}_{e,i}$, which allows us as before to control the linkages we build.

To this end, suppose ϵ is a grid-like end, and that N is as in Lemma 6.3, so that the ray graph of any family of at least $N + 2$ disjoint rays is a cycle. We say that an ϵ -ray R is a *core ray* (of ϵ) if there is some finite family $(R_i : i \in [n])$ of $n \geq N + 3$ disjoint ϵ -rays such that $R = R_i$ for some $i \in [n]$ ⁷.

Every large enough ray graph is a cycle, which has a correct orientation by Lemma 6.3, and we would like to say that this orientation is induced by a global ‘partial cyclic order’ defined on the core rays of ϵ .

By a similar argument as in Section 6.1, one can show the following:

Lemma 8.3. *For every core ray R of a grid-like end ϵ there is a unique sub-end of ϵ in $G - V(R)$, which is linear (cf. Definition 7.9).*

It follows that if R and R' are disjoint core rays then ϵ splits into at most two ends in $G - (V(R) \cup V(R'))$.

⁷We note that it is possible to show that, if ϵ is grid-like, then in fact $N = 3$.

Definition 8.4. Let R and R' be disjoint core rays of ϵ . We denote by $\top(\epsilon, R, R')$ the end of $G - (V(R) \cup V(R'))$ containing rays which appear between R and R' according to the correct orientation of some ray graph of a family of at least $N + 3$ ϵ -rays and by $\perp(\epsilon, R, R')$ the end of $G - (V(R) \cup V(R'))$ containing rays which appear between R' and R in the correct orientation of some ray graph of a family of at least $N + 3$ ϵ -rays.

We will model our global ‘partial cyclic order’ as a ternary relation on the set of core rays of ϵ . That is, a *partial cyclic order* on a set X is a relation $C \subset X^3$ written $[a, b, c]$ satisfying the following axioms:

- If $[a, b, c]$ then $[b, c, a]$.
- If $[a, b, c]$ then not $[c, b, a]$.
- If $[a, b, c]$ and $[a, c, d]$ then $[a, b, d]$.

Lemma and Definition 8.5. Let $\text{core}(\epsilon)$ denote the set of core rays of ϵ . We define a *partial cyclic order* C_ϵ on $\text{core}(\epsilon)$ as follows:

$[R, S, T]$ if and only if R, S, T have disjoint tails xR, yS, zT and $yS \in \top(\epsilon, xR, zT)$.

Then, for any family $(R_i : i \in [n])$ of $n \geq N + 3$ disjoint ϵ -rays, the cyclic order induced on $(R_i : i \in [n])$ by C_ϵ agrees with the correct orientation.

Again, by a similar argument as in Section 6.1, one can show that this relation is in fact a partial cyclic order and that it always agrees with the correction orientation of large enough ray graphs. Furthermore, by Lemma 6.3, given two families \mathcal{R} and \mathcal{S} of at least $N + 3$ disjoint ϵ -rays, every transitional linkage between \mathcal{R} and \mathcal{S} preserves C_ϵ , for the obvious definition of preserving.

Given a family of disjoint ϵ -rays $\mathcal{R} = (R_i : i \in [n])$ with a linear order \leq on \mathcal{R} , we say that \leq agrees with C_ϵ if $[R_i, R_j, R_k]$ whenever $R_i < R_j < R_k$.

Given a family $F = (f_i : i \in I)$ and a linear order \leq on I , we denote by $F(\leq)$ the linear order on F induced by \leq , i.e. the order defined by $f_i F(\leq) f_j$ if and only if $i \leq j$.

As in Section 7 we say a thick G -tribe \mathcal{F} strongly agrees about $\partial(T_n)$ if

- it weakly agrees about $\partial(T_n)$;
- for each $H \in \bigcup \mathcal{F}$ every ϵ -ray $R \subseteq H$ is in $\text{core}(\epsilon)$;
- for every $e \in \partial_\epsilon(T_n)$ there is a linear order $\leq_{\mathcal{F}, e}$ on $S(e)$ such that $H^\downarrow(\mathcal{R}_e)(\leq_{\mathcal{F}, e})$ agrees with C_ϵ on $H^\downarrow(\mathcal{R}_e)$ for all $H \in \bigcup \mathcal{F}$.

Using this definition, the G -tribe refinement lemma (Lemma 7.16) can also be shown to hold in the case where ϵ is a grid-like-end.

Furthermore, we modify the definition of an extension scheme for a family of disjoint inflated copies of $G(T_n^{-\epsilon})$.

Definition 8.6 (Extension scheme). Let $\mathcal{Q} = (Q_i : i \in [k])$ be a family of disjoint inflated copies of $G(S^{-\epsilon})$ and \mathcal{F} be a G -tribe which strongly agrees about $\partial(S)$. We call a tuple (X, \mathcal{E}) an *extension scheme* for \mathcal{Q} if the following holds:

- (ES1) X is a bounder for \mathcal{Q} and \mathcal{E} is an extender for \mathcal{Q} ;
- (ES2) \mathcal{E} is a family of core rays;
- (ES3) the order C_ϵ agrees with $\mathcal{E}_{e,i}^-(\leq_{\mathcal{F},e})$ for every $e \in \partial_\epsilon(S)$;
- (ES4) the sets $\mathcal{E}_{s,i}^-$ are intervals of C_ϵ on \mathcal{E}^- for all $e \in \partial_\epsilon(S)$ and $i \in [k]$.

We can then proceed by induction as before, with the same induction hypotheses. For the most part the proof will follow verbatim, apart from one slight technical issue.

Recall that, in the case where n is even, we use the existence of the family of rays $\bar{\mathcal{C}}$ to find a linkage from \mathcal{C} to $F^\downarrow(\mathcal{R}^-)$ which is preserving on \mathcal{C} and similarly, in the case where n is odd, we do the same for $\bar{\mathcal{E}}_n^-$. In the grid-like case we do not have to be so careful, since every transitional linkage from \mathcal{C} to $F^\downarrow(\mathcal{R}^-)$ will preserve C_ϵ , as long as $|\mathcal{C}|$ is large enough.

However, in order to ensure that $|\mathcal{C}|$ and $|\mathcal{E}_n^-|$ are large enough in each step, we should start by building $N + 3$ inflated copies of $G(T_0^{-\epsilon})$ in the first step, which can be done relatively straightforwardly. Indeed, in the case $n = 0$ most of the argument in the construction is unnecessary, since a large part of the construction is constructing a new copy whilst re-routing the rays \mathcal{E}_n to avoid this new copy, but \mathcal{E}_0 is empty. Therefore, it is enough to choose a layer $F \in \mathcal{F}_0$ with $|F| \geq N + 3$, with say $H_1, \dots, H_{N+3} \in F$ and to take, for each $i \in [N + 3]$,

$$Q_i^1 := H_i(T_k^{-\epsilon})$$

and to take $E_{e,s,i}^1 = H_i^\downarrow(R_{e,s})$ for each $e \in \partial_\epsilon(T_0)$, $s \in S(e)$, and $i \in [N + 3]$. One can then proceed as before, extending the copies in odd steps and adding a new copy in even steps.

9. Outlook: connections with well-quasi-ordering and better-quasi-ordering

Our aim in this section is to sketch what we believe to be the limitations of the techniques of this paper. We will often omit or ignore technical details in order to give a simpler account of the relationship of the ideas involved.

Our strategy for proving ubiquity is heavily reliant on well-quasi-ordering results. The reason is that they are the only known tool for finding extensive tree-decompositions for broad classes of graphs.

To more fully understand this, let us recall how well-quasi-ordering was used in the proofs of Lemmas 5.7 and 5.12. Lemma 5.7 states that any locally finite connected graph with only finitely many ends, all of them thin, has an extensive tree-decomposition. The key idea of the proof was as follows: for each end, there is a sequence of separators converging towards that end. The graphs between these separators are finite, and so are well-quasi-ordered by the Graph Minor Theorem. This well-quasi-ordering guarantees the necessary self-similarity.

Lemma 5.12, where infinitely many ends are allowed but the graph must have finite tree-width, is similar: once more, for each end there is a sequence of separators converging towards that end. The graphs between these separators are not necessarily finite, but they have bounded tree-width and so they are again well-quasi-ordered.

Note that the Graph Minor Theorem is not needed for this latter result. Instead, the reason it works can be expressed in the following slogan, which will motivate the considerations in the rest of this section:

Trees of wombats are well-quasi-ordered precisely when wombats themselves are better-quasi-ordered.

Here better-quasi-ordering is a strengthening of well-quasi-ordering, introduced by Nash-Williams in [NW65] essentially in order to make this slogan be true. Since graphs of bounded tree-width can be encoded as trees of graphs of bounded size, what is used here is that graphs of bounded size are better-quasi-ordered.

What if we wanted to go a little further, for example by allowing infinite tree-width but requiring that all ends should be thin? In that case, all we would know about the graphs between the separators would be that all their ends are thin. Such graphs are essentially trees of finite graphs. So, by the slogan above, to show that such trees are well-quasi-ordered we would need the statement that finite graphs are better-quasi-ordered.

Indeed, this problem arises even if we restrict our attention to the following natural common strengthening of Theorems 1.1 and 1.2:

Conjecture 9.1. Any locally finite connected graph in which all blocks are finite is \preceq -ubiquitous.

In order to attack this conjecture with our current techniques we would need better-quasi-ordering of finite graphs.

Thomas has conjectured [Tho89] that countable graphs are well-quasi-ordered with respect to the minor relation. If this were true, it could allow us to resolve problems like those discussed above for countable graphs at least, since all the graphs appearing between the separators are countable. But this approach does not allow us to avoid the issue of better-quasi-ordering of finite graphs. Indeed, since countable trees of finite graphs can be coded as countable graphs, well-quasi-ordering of countable graphs would imply better-quasi-ordering of finite graphs.

Thus until better-quasi-ordering of finite graphs has been established, the best that we can hope for – using our current techniques – is to drop the condition of local finiteness from the main results of this paper. For countable graphs we hope to show this in a future paper, however for graphs of larger cardinalities further issues arise.

References

- [And77] Thomas Andreae. Bemerkung zu einem Problem aus der Theorie der unendlichen Graphen. *Abh. Math. Sem. Univ. Hamburg*, 46:91–95, 1977. doi:10.1007/BF02993015.

- [And79] Thomas Andreae. Über eine Eigenschaft lokalfiniten, unendlicher Bäume. *J. Combin. Theory Ser. B*, 27(2):202–215, 1979. doi:10.1016/0095-8956(79)90082-0.
- [And80] Thomas Andreae. On a problem of R. Halin concerning infinite graphs. II. *Discrete Math.*, 29(3):219–227, 1980. doi:10.1016/0012-365X(80)90149-1.
- [And02] Thomas Andreae. On disjoint configurations in infinite graphs. *J. Graph Theory*, 39(4):222–229, 2002. doi:10.1002/jgt.10016.
- [And13] Thomas Andreae. Classes of locally finite ubiquitous graphs. *J. Combin. Theory Ser. B*, 103(2):274–290, 2013. doi:10.1016/j.jctb.2012.11.003.
- [BEE⁺22] Nathan Bowler, Christian Elbracht, Joshua Erde, J. Pascal Gollin, Karl Heuer, Max Pitz, and Maximilian Teegen. Topological ubiquity of trees. *J. Combin. Theory Ser. B*, 157:70–95, 2022. doi:10.1016/j.jctb.2022.05.011.
- [BEE⁺23] Nathan Bowler, Christian Elbracht, Joshua Erde, J. Pascal Gollin, Karl Heuer, Max Pitz, and Maximilian Teegen. Ubiquity of graphs with nowhere-linear end structure. *J. Graph Theory*, 103(3):564–598, 2023. doi:10.1002/jgt.22936.
- [Die17] Reinhard Diestel. *Graph theory*, volume 173 of *Graduate Texts in Mathematics*. Springer, Berlin, fifth edition, 2017. doi:10.1007/978-3-662-53622-3.
- [Hal65] R. Halin. Über die Maximalzahl fremder unendlicher Wege in Graphen. *Math. Nachr.*, 30(1-2):63–85, 1965. doi:10.1002/mana.19650300106.
- [Hal70] R. Halin. Die Maximalzahl fremder zweiseitig unendlicher Wege in Graphen. *Math. Nachr.*, 44(1-6):119–127, 1970. doi:10.1002/mana.19700440109.
- [Kom95] Péter Komjáth. A note on minors of uncountable graphs. *Math. Proc. Cambridge Philos. Soc.*, 117(1):7–9, 1995. doi:10.1017/S0305004100072881.
- [KT91] Igor Kříž and Robin Thomas. The Menger-like property of the tree-width of infinite graphs. *J. Combin. Theory Ser. B*, 52(1):86–91, 1991. doi:10.1016/0095-8956(91)90093-Y.
- [Lak76] John Lake. A problem concerning infinite graphs. *Discrete Math.*, 14(4):343–345, 1976. doi:10.1016/0012-365X(76)90066-2.
- [NW65] C. St. J. A. Nash-Williams. On well-quasi-ordering infinite trees. *Proc. Cambridge Philos. Soc.*, 61(3):697–720, 1965. doi:10.1017/s0305004100039062.
- [Pit23] Max Pitz. A note on minor antichains of uncountable graphs. *J. Graph Theory*, 102(3):552–555, 2023. doi:10.1002/jgt.22886.
- [RS04] Neil Robertson and P. D. Seymour. Graph minors. XX. Wagner’s conjecture. *J. Combin. Theory Ser. B*, 92(2):325–357, 2004. doi:10.1016/j.jctb.2004.08.001.
- [RS10] Neil Robertson and Paul Seymour. Graph minors XXIII. Nash-Williams’ immersion conjecture. *J. Combin. Theory Ser. B*, 100(2):181–205, 2010. doi:10.1016/j.jctb.2009.07.003.
- [Tho88] Robin Thomas. A counterexample to “Wagner’s conjecture” for infinite graphs. *Math. Proc. Cambridge Philos. Soc.*, 103(1):55–57, 1988. doi:10.1017/S0305004100064616.

- [Tho89] Robin Thomas. Well-quasi-ordering infinite graphs with forbidden finite planar minor. *Trans. Amer. Math. Soc.*, 312(1):279–313, 1989. doi:10.2307/2001217.
- [Woo76] D. R. Woodall. A note on a problem of R. Halin’s: “A problem in infinite graph-theory” (Abh. Math. Sem. Univ. Hamburg **43** (1975), 79–84). *J. Combinatorial Theory Ser. B*, 21(2):132–134, 1976. doi:10.1016/0095-8956(76)90052-6.