

Ubiquity, Hamiltonicity and dijoins in graphs

Habilitationschrift

vorgelegt im Fachbereich Mathematik
der Universität Hamburg

von

Karl Magnus Heuer

Hamburg, Juli 2022

To Meike

Contents

General introduction	7
Overview about Chapter I: Ubiquity of graphs	7
Overview about Chapter II: Hamiltonicity of locally finite graphs	9
Overview about Chapter III: Dijoins of digraphs	11
Acknowledgement	13
I. Ubiquity of graphs	15
A. Topological ubiquity of trees	17
A.1. Introduction	17
A.2. Preliminaries	20
A.3. Well-quasi-orders and κ -embeddability	21
A.4. Linkages between rays	23
A.5. G -tribes and concentration of G -tribes towards an end	28
A.6. Countable subtrees	32
A.6.1. Preprocessing	33
A.6.2. Only finitely many vertices dominate ϵ	37
A.6.3. Infinitely many vertices dominate ϵ	40
A.7. The induction argument	42
B. Ubiquity of graphs with nowhere-linear end structure	47
B.1. Introduction	47
B.2. Preliminaries	51
B.3. The Ray Graph	54
B.4. A pebble-pushing game	56
B.5. Pebbly and non-pebbly ends	61
B.6. The structure of non-pebbly ends	65
B.6.1. Polypods	65

B.6.2. Frames	69
B.7. Grid-like and half-grid-like ends	73
B.7.1. Grid-like ends	74
B.7.2. Half-grid-like ends	76
B.8. G -tribes and concentration of G -tribes towards an end	80
B.9. Ubiquity of minors of the half-grid	83
B.10. Proof of main results	85
C. Ubiquity of locally finite graphs with extensive tree-decompositions	91
C.1. Introduction	91
C.2. Proof sketch	93
C.3. Preliminaries	101
C.3.1. Rays and ends	101
C.3.2. Inflated copies of graphs	103
C.3.3. Transitional linkages and the strong linking lemma	104
C.3.4. Separations and tree-decompositions of graphs	106
C.4. Extensive tree-decompositions and self minors	107
C.4.1. Extensive tree-decompositions	107
C.4.2. Self minors in extensive tree-decompositions	109
C.5. Existence of extensive tree-decompositions	110
C.5.1. Finitely many thin ends	111
C.5.2. Finite tree-width	114
C.5.3. Sporadic examples	118
C.6. The structure of non-pebbly ends	119
C.6.1. Core rays in the half-grid-like case	120
C.7. G -tribes and concentration of G -tribes towards an end	128
C.8. The inductive argument	138
C.8.1. The half-grid-like case	139
C.8.2. The grid-like case	150
C.9. Outlook: connections with well-quasi-ordering and better-quasi-ordering	153

II. Hamiltonicity of locally finite graphs	157
D. Forcing Hamiltonicity in locally finite graphs via forbidden induced subgraphs I: nets and bulls	159
D.1. Introduction	159
D.2. Preliminaries	163
D.2.1. Basic notions	163
D.2.2. Topological notions	165
D.2.3. Tools	166
D.3. About the structure of the graphs considered in this paper	169
D.4. Hamiltonicity results	179
E. Forcing Hamiltonicity in locally finite graphs via forbidden induced subgraphs II: paws	191
E.1. Introduction	191
E.2. Preliminaries	193
E.2.1. Basic notions	193
E.2.2. Topological notions	194
E.2.3. Tools	195
E.3. Examples of graph meeting the criteria of Theorem E.1.3	197
E.4. Proof of Theorem E.1.3	200
F. Hamilton-laceable bi-powers of locally finite bipartite graphs	213
F.1. Introduction	213
F.2. Preliminaries	215
F.2.1. Basic notions and tools	215
F.2.2. Topological notions and tools	218
F.3. Counterexamples for potential strengthenings of Theorem F.1.5	219
F.3.1. No perfect matching in G	219
F.3.2. Perfect matching in G_B^3 instead of G	219
F.3.3. Hamilton-connectedness	220
F.4. Proof of the main result	221

III. Dijoins of digraphs 231

G. Disjoint dijoin for classes of dicuts in finite and infinite digraphs 233

G.1. Introduction	233
G.2. Preliminaries	236
G.2.1. Digraphs	236
G.2.2. Hypergraphs	238
G.3. On the capacitated version	238
G.4. The dicut hypergraph	241
G.4.1. Transversal packings of hypergraphs	242
G.4.2. Applications to the dicut hypergraph	244
G.5. Minimum size dicuts and disjoint mini-dijoins	245
G.6. Dijoin in infinite digraphs	248
G.6.1. Classes of finite dicuts in infinite digraphs	249
G.6.2. Classes of infinite dibonds	254

H. Even circuits in oriented matroids 259

H.1. Introduction	259
H.1.1. The Even Directed Circuit Problem in Oriented Matroids	260
H.2. Background	264
H.2.1. Non-Evenness and GB-minors	269
H.3. On the Complexity of the Even Directed Circuit Problem	271
H.3.1. Proof of Theorem H.1.6	275
H.4. Digraphs Admitting an Odd Dijoin	284
H.4.1. Forbidden cut minors for digraphs with an odd dijoin	286
H.5. Concluding remarks	296

General introduction

This habilitation thesis consists of three chapters. Each of these chapters naturally stands for its own research topic, in which several publications emerged. The first chapter falls in the area of Infinite Graph Theory and focusses on questions about the so-called ‘ubiquity’ of certain graphs. The second chapter also contains results about infinite graphs. More precisely, it presents research regarding the Hamiltonicity of locally finite infinite graphs, i.e. infinite graphs where every vertex has finite degree. In contrast to the first two chapters of this thesis, the third chapter deals with directed graphs. Furthermore, the research presented in that chapter concerns mostly finite directed graphs, but to some extent also infinite directed graphs as well as finite oriented matroids.

In total this thesis contains the results of eight papers [13–15, 41, 53, 57–59] and one addendum to [53], which is included in an extended preprint [52] on ArXiv. I have written these articles together with several co-authors since 2018 during my times at the University of Hamburg, at the Technical University of Denmark and at the Technical University of Berlin.

Overview about Chapter I: Ubiquity of graphs

This chapter focusses on questions regarding ubiquity of graphs. We call a graph G *ubiquitous* if whenever some graph Γ contains k many disjoint copies of G as a subgraph for every $k \in \mathbb{N}$, then Γ also contains infinitely many disjoint copies of G as a subgraph. Analogously we can define being ubiquitous also with respect to other containment relations than the subgraph relation, for example for the topological minor relation or the minor relation. In general we shall therefore speak about being ubiquitous with respect to some containment relation.

Questions regarding the ubiquity of graphs are very fundamental and deep ones within Infinite Graph Theory. The first result in this area is due to Halin [49], which proves that *rays*, i.e. one-way infinite paths, are ubiquitous with respect

to the subgraph relation. While all finite graphs are trivially ubiquitous with respect to the subgraph relation, not every infinite graph is so as well. See for example [6] for infinite graphs of relatively simple structure that are not ubiquitous with respect to the subgraph relation or the topological minor relation.

The most important conjecture in this field is the so-called *Ubiquity Conjecture* due to Andreae [3].

The Ubiquity Conjecture. *Every locally finite connected graph is ubiquitous with respect to the minor relation.*

Although this conjecture has been verified in some special cases, it is still widely open.

The results presented in the first chapter of this thesis are centred around the Ubiquity Conjecture and are spread over three articles [13–15], of which the first has already been published in the Journal of *Combinatorial Theory, Series B*. These three articles are joint work together with Nathan Bowler, Christian Elbracht, Joshua Erde, Pascal Gollin, Max Pitz and Maximilian Teegen.

The main result of the first article [13] is that all trees, irrespective of their cardinality, are ubiquitous with respect to the topological minor relation.

The second article [14] provides a sufficient condition for a graph to be ubiquitous with respect to the minor relation. This condition is encoded as a structural property of the graph G in terms of its *ends*, which are equivalence classes of rays in G where two rays in G are declared to be *equivalent* if they are joined by infinitely many disjoint paths in G . Probably the most prominent result deduced from this work is the verification of the full grid (the graph which is sometimes also referred to as $\mathbb{Z} \square \mathbb{Z}$) being ubiquitous with respect to the minor relation.

The third paper [15] shows that all locally finite graphs that admit a certain type of tree-decomposition are ubiquitous with respect to the minor relation. The two most relevant special cases that are deduced from this are the following ones. First it is shown that all graphs of finite treewidth are ubiquitous with respect to the minor relation. Note that here, *finite treewidth* means that all bags of the decomposition are bounded in size by some common integer. The second deduced result states that all graphs with only finitely many ends, all of which are *thin*, i.e. there do not exist infinitely many disjoint rays in that end, are ubiquitous with respect to the minor relation.

Overview about Chapter II: Hamiltonicity of locally finite graphs

The problem to decide whether a finite graph contains a cycle that covers all vertices of that graph is a very prominent one within Graph Theory. Such a cycle is referred to as a *Hamilton cycle*. While the general decision problem for the existence of a Hamilton cycle is hard, which can formally be encoded in terms of computational complexity, there are many special cases, e.g. in terms of restricting the problem to certain graph classes, for which a Hamilton cycle can (even efficiently) be found. In other words, the research field of finding sufficient conditions to impose on a graph in order to guarantee the existence of a Hamilton cycle is a very active one.

The majority of Hamiltonicity results consider only finite graphs. The reason for this is that there is no obvious answer to the question what a Hamilton cycle of an infinite graph should be. A quite successful definition for locally finite graphs, measured by the number of Hamiltonicity results for finite graphs that could be generalised to locally finite ones, is the following. Take as cycles of a locally finite connected graph G the homeomorphic images of the unit circle $S^1 \subseteq \mathbb{R}^2$ in the Freudenthal compactification $|G|$ of G . While all finite cycles of G are encompassed by this definition, it also provides a variety of possible infinite cycles the graph G might have, of course depending on its structure. Note that, roughly speaking, the topological space $|G|$ can be seen as the graph G with additional points at infinity. These additional points correspond to the ends of G . Note that within $|G|$ every ray of G converges to the end of G it belongs to. This topological approach for defining infinite cycles is due to Diestel and Kühn [27, 28] and it forms the foundation for the studies that have been conducted in this chapter. Additionally, it enables us to easily lift the definition of a Hamilton cycle from finite graphs to locally finite ones.

The results of this second chapter are covered by three papers [53, 57, 58] and an additional addendum to paper [53], which has been incorporated in an extended version of [53] on ArXiv, see [52]. The first two articles [57, 58] are joined work with Deniz Sarikaya. The results of these two articles concern lifting sufficient conditions for the existence of a Hamilton cycle in finite graphs to locally finite graphs. In both articles the considered sufficient conditions are formulated via

forbidden induced subgraphs (or slight relaxations of this). The considered induced subgraphs are the claw, the net, the bull and the paw (see Figure 1).




Figure 1.: The subgraphs focused on in Sections D and E.

The main results of paper [57] are a structural description of locally finite connected graphs without induced claws or nets, which generalises a corresponding result about finite graphs due to Shepherd [98] to locally finite ones. One consequence of this result is that such graphs have at most two ends. Furthermore, this result is utilised to prove, beside other results, that every locally finite 2-connected graph without induced claws or nets admits a Hamilton cycle.

Regarding bulls, a slight relaxation to completely forbidding these graphs as induced subgraphs is considered. The relaxed condition allows induced bulls within the graph, but only under the condition that its *horns*, i.e. the two vertices of degree 1 within a bull (cf. Figure 1), have a common neighbour outside the bull. The main result of [57] with respect to bulls states that for locally finite infinite 2-connected graphs without induced claws, demanding the relaxed condition for bulls is actually equivalent to demanding the graph to no contain any induced bull.

The second paper [58] whose results are included in this chapter focusses on forbidding claws as induced subgraphs and imposing a similar condition for induced paws as for bulls before in paper [57]. The main result states that every locally finite 2-connected graph without induced claws and all whose induced paws admit a corresponding relaxed condition has a Hamilton cycle.

The third paper [53] belonging to this chapter is single authored and has already been published in the journal *Discrete Mathematics*. The article considers the *bi-cube* G_B^3 of a bipartite graph G , which is similar to forming the cube of G but additional edges due to the third power are only added among vertices whose distance from each other in G is odd. The main result proves that for every

locally finite connected bipartite graph G that admits a perfect matching, its bi-cube is *Hamilton-laceable*, i.e. for any two vertices u, v of G that lie in different bipartition classes of G there exists a homeomorphic image of $[0, 1] \subseteq \mathbb{R}$ within $|G_B^3|$ where 0 is mapped to u and 1 to v that contains all vertices of G . Note that Hamilton-laceability is a relaxation of Hamilton-connectedness where only pairs of vertices are considered whose distance from each other is odd.

Note that two questions have been raised in paper [53]. Both of them have been answered negatively after the paper has been published. The short counterexample, which provides the negative answers, is included in an extended version of [53] on ArXiv, see [52]. In Section F the counterexample is incorporated directly after the two raised questions.

Overview about Chapter III: Dijoins of digraphs

This chapter contains the results of two articles [41, 59]. In contrast to the two previous chapters, the content of this chapter has a focus on directed graphs, briefly called *digraphs*. Furthermore, although infinite digraphs are studied as well, the focus of the work rather lies on finite digraphs than on infinite ones. The common objects which are studied in both papers, although under different aspects, are dijoins of digraphs. For a digraph D , a set of edges F is called a *dijoin* of D if F intersects every dicut of D , where a *dicut* of D is a cut all whose edges are directed towards the same common side of the cut. Note that dijoins are a special type of transversal edge sets, namely for the set of all dicuts of a digraph. Furthermore, note that the dual objects to dijoins, with respect to planar duality, are known as *feedback arc sets*. They form transversal edge sets for the set of all directed cycles of a digraph.

The first article [41] is joint work with Pascal Gollin and Konstantinos Stavropoulos. The work is motivated by a very deep and still open conjecture regarding dijoins of finite digraphs, called Woodall's Conjecture:

Woodall's Conjecture (Woodall 1976 [111]). The size of a smallest non-empty dicut in a finite digraph D is equal to the size of the largest set of disjoint dijoins of D .

Paper [41] considers relaxations of Woodall's Conjecture by restricting the

attention to subclasses of all dicuts and accordingly adapting the definition of dijoins to only intersect the members of such subclasses. The results encompass several classes of dicuts for which a corresponding version of Woodall's Conjecture holds. Two main examples are arbitrary classes of nested dicuts and the set of dicuts of smallest size within a digraph. Additionally, several existing results about Woodall's Conjecture, but also the new results of [41] regarding relaxations of Woodall's Conjecture for finite digraphs are lifted to infinite digraphs.

The second article [59] is joint work with Raphael Steiner and Sebastian Wiederrecht and has already been published in the journal *Combinatorial Theory*. This work studies a special type of dijoins of finite digraphs, namely so-called odd dijoins. A dijoin of a digraph is called *odd* if it intersects every di-cut of D in an odd number of edges. While the existence of a dijoin of a digraph is trivial, just take all edges of the digraph, the existence of odd dijoins is not. For the analogously defined (planar) dual objects, call them *odd feedback arc sets*, their existence has been characterised in terms of forbidden butterfly minors by Seymour and Thomassen [96]. The main result of [59] proves a dual version of the characterisation of Seymour and Thomassen for the existence of odd dijoins. For this, a dual notion of butterfly minors has been introduced, which is called *cut minors*. Interestingly, the obstructions that occur as forbidden cut minors are not just the planar dual digraphs to those from the characterisation by Seymour and Thomassen for the existence of odd feedback arc sets. A new class of non-planar digraphs has to be included in the list of obstruction.

Due to the duality aspect of the newly introduced definition of cut minors, a lot of the work has been lifted even further into the setting of regular oriented matroids, which generalise digraphs in a matroidal way. A matroidal notion of butterfly minors is introduced, called *generalised butterfly minors*. For graphic oriented matroids this notion corresponds to usual butterfly minors and, after dualising, to cut minors. Motivated by the results about characterising the existence of odd dijoins, a unifying conjecture for regular oriented matroids is raised, where an odd dijoin, or rather an odd feedback arc set now translates to a set of elements intersecting every directed circuit in an odd number of elements.

Beside this work, also questions regarding computational complexity are studied in paper [59]. The decision problem whether a finite digraph contains an even cycle turned out to be a very complicated one, and became known as the *Even*

Directed Cycle Problem. This problem was resolved by Robertson, Seymour and Thomas [90] and independently by McCuaig [80] who provided polynomial time algorithms to answer this decision problem. Motivated by this a corresponding question is raised in paper [59] which asks whether the decision problem for the existence of an even directed circuit in a regular oriented matroid can be answered in polynomial time with respect to the size of the representative matrix for the oriented matroid. Note that for a finite digraph D , deciding whether D contains a directed cycle of even size is polynomial time equivalent to deciding whether D admits an odd feedback arc set. This result is due to Seymour and Thomassen [96]. A lifted version of this into the setting of regular oriented matroids forms another result of paper [59].

Acknowledgement

I would like to thank my co-authors Nathan Bowler, Christian Elbracht, Joshua Erde, Pascal Gollin, Max Pitz, Deniz Sarikaya, Konstantinos Stavropoulos, Raphael Steiner, Maximilian Teegen and Sebastian Wiederrecht for the fruitful and interesting collaborations. In particular, I want to additionally thank Pascal Gollin for the many nice hours of thinking together about problems.

Furthermore, I would like to thank Carsten Thomassen, Reinhard Diestel and Stephan Kreutzer for the great environment they have provided during my times in their research groups as well as for many interesting and inspiring discussions.

I also want to thank my mother, Christine, for all her love and everlasting support.

Finally, I want to thank my wife, Meike, for coming with me when being abroad and all the wonderful moments together. I am even more grateful for the family and home full of love, security and joy that you, Meike, have created with me. I love you!

Chapter I.

Ubiquity of graphs

A. Topological ubiquity of trees

A.1. Introduction

Let \triangleleft be a relation between graphs, for example the subgraph relation \subseteq , the topological minor relation \leq or the minor relation \preceq . We say that a graph G is \triangleleft -ubiquitous if whenever Γ is a graph with $nG \triangleleft \Gamma$ for all $n \in \mathbb{N}$, then one also has $\aleph_0 G \triangleleft \Gamma$, where αG is the disjoint union of α many copies of G .

Two classic results of Halin [48, 49] say that both the ray and the double ray are \subseteq -ubiquitous, i.e. any graph which contains arbitrarily large collections of disjoint (double) rays must contain an infinite collection of disjoint (double) rays. However, even quite simple graphs can fail to be \subseteq or \leq -ubiquitous, see e.g. [6, 69, 112], examples of which, due to Andreae [2], are depicted in Figures A.1 and A.2 below.




Figure A.1.: A graph which is not \subseteq -ubiquitous.




Figure A.2.: A graph which is not \leq -ubiquitous.

However, for the minor relation, no such simple examples of non-ubiquitous graphs are known. Indeed, one of the most important problems in the theory of infinite graphs is the so-called *Ubiquity Conjecture* due to Andreae [3].

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

In [3], Andreae constructed a graph that is not \preceq -ubiquitous. However, this construction relies on the existence of a counterexample to the well-quasi-ordering

of infinite graphs under the minor relation, for which counterexamples are only known with very large cardinality [103]. In particular, it is still an open question whether or not there exists a countable connected graph which is not \preceq -ubiquitous.

Whilst there are examples, see Figure A.2, of quite simple graphs which are not \preceq -ubiquitous, it was shown by Halin [46] that all trees of maximum degree 3 are \preceq -ubiquitous. Andreae improved this result to show that all *locally finite* trees are \preceq -ubiquitous [5], and asked if his result could be extended to arbitrary trees [5, p. 214]. Our main result of this paper answers this question in the affirmative. Moreover, it seems that the methods we develop to tackle this question will be useful in other contexts. In forthcoming papers [14, 15] we will use similar ideas to show that the ubiquity conjecture holds for all locally finite graphs of finite tree-width.

Theorem A.1.1. *Every tree is ubiquitous with respect to the topological minor relation.*

The proof will use some results of Nash-Williams [83] and Laver [70] about the well-quasi-ordering of trees under the topological minor relation, as well as some notions about the topological structure of infinite graphs [29]. Interestingly, most of the work in proving Theorem A.1.1 lies in dealing with the countable case, where several new ideas are needed. In fact, we will prove a slightly stronger statement in the countable case, which will allow us to derive the general result via transfinite induction on the cardinality of the tree, using some ideas from Shelah’s singular compactness theorem [97].

To explain our strategy, let us fix some notation. When H is a subdivision of G we write $G \preceq^* H$. Then, $G \preceq \Gamma$ means that there is a subgraph $H \subseteq \Gamma$ which is a subdivision of G , that is, $G \preceq^* H$. If H is a subdivision of G and v a vertex of G , then we denote by $H(v)$ the corresponding vertex in H . More generally, given a subgraph $G' \subseteq G$, we denote by $H(G')$ the corresponding subdivision of G' in H .

Now, suppose we have a rooted tree T and a graph Γ . Given a vertex $t \in T$, let T_t denote the rooted subtree of T rooted in t . We say that a vertex $v \in \Gamma$ is *t -suitable* if there is some subdivision H of T_t in Γ with $H(t) = v$. For a subtree $S \subseteq T$ we say that a subdivision H of S in Γ is *T -suitable* if for each vertex $s \in V(S)$ the vertex $H(s)$ is *s -suitable*, i.e. for every $s \in V(S)$ there is a subdivision H' of T_s such that $H'(s) = H(s)$.

An S -horde is a sequence $(H_i: i \in \mathbb{N})$ of disjoint suitable subdivisions of S in Γ . If S' is a subtree of S , then we say that an S -horde $(H_i: i \in \mathbb{N})$ *extends* an S' -horde $(H'_i: i \in \mathbb{N})$ if for every $i \in \mathbb{N}$ we have $H_i(S') = H'_i$.

In order to show that an arbitrary tree T is \leq -ubiquitous, our rough strategy will be to build, by transfinite recursion, S -hordes for larger and larger subtrees S of T , each extending all the previous ones, until we have built a T -horde. However, to start the induction it will be necessary to show that we can build S -hordes for countable subtrees S of T . This will be done in the following key result of this paper:

Theorem A.1.2. *Let T be a tree, S a countable subtree of T and Γ a graph such that $nT \leq \Gamma$ for every $n \in \mathbb{N}$. Then there is an S -horde in Γ .*

Note that Theorem A.1.2 in particular implies \leq -ubiquity of countable trees.

We remark that, whilst the relation \preceq is a relaxation of the relation \leq , which is itself a relaxation of the relation \subseteq , it is not clear whether \subseteq -ubiquity implies \leq -ubiquity, or whether \leq -ubiquity implies \preceq -ubiquity. In the case of Theorem A.1.1 however, we believe that similar methods will also show that arbitrary trees are \preceq -ubiquitous, although such a proof would involve quite a bit of extra technical detail to deal with the fact that the branch sets of vertices may themselves be infinite trees. We note, however, that it is surprisingly easy to show that countable trees are \preceq -ubiquitous, since it can be derived relatively straightforwardly from Halin's grid theorem, see [14, Theorem 1.5].

This paper is structured as follows. In Section A.2 we provide background on rooted trees, rooted topological embeddings of rooted trees (in the sense of Kruskal and Nash-Williams), and ends of graphs. In our graph theoretic notation we generally follow the textbook of Diestel [24]. Next, Sections A.3–A.5 introduce the key ingredients for our main ubiquity result. Section A.3, extending ideas of Andreae [5], lists three useful corollaries of Nash-Williams' and Laver's result that (labelled) trees are well-quasi-ordered under the topological minor relation, Section A.4 investigates under which conditions a given family of disjoint rays can be rerouted onto another family of disjoint rays, and Section A.5 shows that without loss of generality, we already have quite a lot of information about how exactly our copies of nG are placed in the host graph Γ .

Using these ingredients, we give a proof of the countable case, i.e. of Theo-

rem A.1.2, in Section A.6. Finally, Section A.7 contains the induction argument establishing our main result, Theorem A.1.1.

A.2. Preliminaries

Definition A.2.1. A *rooted graph* is a pair (G, v) where G is a graph and $v \in V(G)$ is a vertex of G which we call the *root*. Often, when it is clear from the context which vertex is the root of the graph, we will refer to a rooted graph (G, v) as simply G .

Given a rooted tree (T, v) , we define a partial order \leq , which we call the *tree-order*, on $V(T)$ by letting $x \leq y$ if the unique path between y and v in T passes through x . See [24, Section 1.5] for more background. For any edge $e \in E(T)$ we denote by e^- the endpoint closer to the root and by e^+ the endpoint further from the root. For any vertex t we denote by $N^+(t)$ the set of *children of t* in T , the neighbours s of t satisfying $t \leq s$. The subtree of T rooted at t is denoted by (T_t, t) , that is, the induced subgraph of T on the set of vertices $\{s \in V(T) : t \leq s\}$.

We say that a rooted tree (S, w) is a *rooted subtree* of a rooted tree (T, v) if S is a subgraph of T such that the tree order on (S, w) agrees with the induced tree order from (T, v) . In this case we write $(S, w) \subseteq_r (T, v)$.

We say that a rooted tree (S, w) is a *rooted topological minor* of a rooted tree (T, v) if there is a subgraph S' of T which is a subdivision of S such that for any $x \leq y \in V(S)$, $S'(x) \leq S'(y)$ in the tree-order on T . We call such an S' a *rooted subdivision of S* . In this case we write $(S, w) \leq_r (T, v)$, cf. [24, Section 12.2].

Definition A.2.2 (Ends of a graph, cf. [24, Chapter 8]). An *end* in an infinite graph Γ is an equivalence class of rays, where two rays R and S 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 in Γ . Given any 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 $C(X, \epsilon)$.

A vertex $v \in V(\Gamma)$ *dominates* an end ω if there is a ray $R \in \omega$ such that there are infinitely many $v-R$ -paths in Γ which but for v are pairwise vertex disjoint.

Definition A.2.3. 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 . 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 .

Given two paths or rays P and Q which are disjoint but for one of their endvertices, we write PQ for the *concatenation of P and Q* , that is the path, ray or double ray $P \cup Q$. Since concatenation of paths is associative, we will not use parentheses. Moreover, if we concatenate paths of the form vPw and wQx , then we omit writing w twice and denote the concatenation by $vPwQx$.

For an integer n , we denote by $[n]$ the set of positive integers up to n .

A.3. Well-quasi-orders and κ -embeddability

Definition A.3.1. Let X be a set and let \triangleleft be a binary relation on X . Given an infinite cardinal κ we say that an element $x \in X$ is κ -*embeddable* (with respect to \triangleleft) in X if there are at least κ many elements $x' \in X$ such that $x \triangleleft x'$.

Definition A.3.2 (well-quasi-order). A binary relation \triangleleft on a set X is a *well-quasi-order* if it is reflexive and transitive, and for every sequence $x_1, x_2, \dots \in X$ there is some $i < j$ such that $x_i \triangleleft x_j$.

Lemma A.3.3. *Let X be a set and let \triangleleft be a well-quasi-order on X . For any infinite cardinal κ the number of elements of X which are not κ -embeddable with respect to \triangleleft in X is less than κ .*

Proof. For $x \in X$ let $U_x = \{y \in X : x \triangleleft y\}$. Now suppose for a contradiction that the set $A \subseteq X$ of elements which are not κ -embeddable with respect to \triangleleft in X has size at least κ . Then, we can recursively pick a sequence $(x_n \in A : n \in \mathbb{N})$ such that $x_m \not\triangleleft x_n$ for $m < n$. Indeed, having chosen all x_m with $m < n$ it suffices to choose x_n to be any element of the set $A \setminus \bigcup_{m < n} U_{x_m}$, which is nonempty since A has size κ but each U_{x_m} has size $< \kappa$.

By construction we have $x_m \not\triangleleft x_n$ for $m < n$, contradicting the assumption that \triangleleft is a well-quasi-order on X . \square

We will use the following theorem of Nash-Williams on well-quasi-ordering of rooted trees, and its extension by Laver to labelled rooted trees.

Theorem A.3.4 (Nash-Williams [83]). *The relation \leq_r is a well-quasi order on the set of rooted trees.*

Theorem A.3.5 (Laver [70]). *The relation \leq_r is a well-quasi order on the set of rooted trees with finitely many labels, i.e. for every finite number $k \in \mathbb{N}$, whenever T_1, T_2, \dots is a sequence of rooted trees and c_1, c_2, \dots is a sequence of k -colourings $c_i: T_i \rightarrow [k]$, there is some $i < j$ such that there exists a subdivision H of T_i with $H \subseteq_r T_j$ and $c_i(t) = c_j(H(t))$ for all $t \in T_i$.**

Together with Lemma A.3.3 these results give us the following three corollaries:

Definition A.3.6. Let (T, v) be an infinite rooted tree. For any vertex t of T and any infinite cardinal κ , we say that a child t' of t is κ -embeddable if there are at least κ children t'' of t such that $T_{t'}$ is a rooted topological minor of $T_{t''}$.

Corollary A.3.7. *Let (T, v) be an infinite rooted tree, $t \in V(T)$ and $\mathcal{T} = \{T_{t'}: t' \in N^+(t)\}$. Then for any infinite cardinal κ , the number of children of t which are not κ -embeddable is less than κ .*

Proof. By Theorem A.3.4 the set $\mathcal{T} = \{T_{t'}: t' \in N^+(t)\}$ is well-quasi-ordered by \leq_r and so the claim follows by Lemma A.3.3 applied to \mathcal{T} , \leq_r , and κ . \square

Corollary A.3.8. *Let (T, v) be an infinite rooted tree, $t \in V(T)$ a vertex of infinite degree and $(t_i \in N^+(t): i \in \mathbb{N})$ a sequence of countably many of its children. Then there exists $N_t \in \mathbb{N}$ such that for all $n \geq N_t$,*

$$\{t\} \cup \bigcup_{i > N_t} T_{t_i} \leq_r \{t\} \cup \bigcup_{i > n} T_{t_i}$$

(considered as trees rooted at t) fixing the root t .

Proof. Consider a labelling $c: T_t \rightarrow [2]$ mapping t to 1, and all remaining vertices of T_t to 2. By Theorem A.3.5, the set $\mathcal{T} = \{\{t\} \cup \bigcup_{i > n} T_{t_i}: n \in \mathbb{N}\}$ is well-quasi-ordered by \leq_r respecting the labelling, and so the claim follows by applying Lemma A.3.3 to \mathcal{T} and \leq_r with $\kappa = \aleph_0$. \square

Definition A.3.9 (Self-similarity). A ray $R = r_0 r_1 r_2 \dots$ in a rooted tree (T, v) which is upwards with respect to the tree order *displays self-similarity of T* if there are infinitely many n such that there exists a subdivision H of T_{r_0} with $H \subseteq_r T_{r_n}$ and $H(R) \subseteq R$.

*In fact, Laver showed that rooted trees labelled by a *better-quasi-order* are again better-quasi-ordered under \leq_r respecting the labelling, but we shall not need this stronger result.

Corollary A.3.10. *Let (T, v) be an infinite rooted tree and let $R = r_0 r_1 r_2 \dots$ be a ray which is upwards with respect to the tree order. Then there is a $k \in \mathbb{N}$ such that $r_k R$ displays self-similarity of T .[†]*

Proof. Consider a labelling $c: T \rightarrow [2]$ mapping the vertices on the ray R to 1, and labelling all remaining vertices of T with 2. By Theorem A.3.5, the set $\mathcal{T} = \{(T_{r_i}, c_i): i \in \mathbb{N}\}$, where c_i is the natural restriction of c to T_{r_i} , is well-quasi-ordered by \leq_r respecting the labellings. Hence, by Lemma A.3.3, the number of indices i such that T_{r_i} is not \aleph_0 -embeddable in \mathcal{T} is finite. Let k be larger than any such i . Then, since T_{r_k} is \aleph_0 -embeddable in \mathcal{T} , there are infinitely many $r_j \in r_k R$ such that $T_{r_k} \leq_r T_{r_j}$ respecting the labelling, i.e. mapping the ray to the ray, and hence $r_k R$ displays the self-similarity of T . \square

A.4. Linkages between rays

In this section we will establish a toolkit for constructing a disjoint system of paths from one family of disjoint rays to another.

Definition A.4.1 (Tail of a ray). Given a ray R in a graph Γ and a finite set $X \subseteq V(\Gamma)$ the *tail of R after X* , denoted by $T(R, X)$, is the unique infinite component of R in $\Gamma - X$.

Definition A.4.2 (Linkage of families of rays). Let $\mathcal{R} = (R_i: i \in I)$ and $\mathcal{S} = (S_j: j \in J)$ be families of vertex 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 joins 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.

We say that \mathcal{T} is obtained by *transitioning* from \mathcal{R} to \mathcal{S} along the linkage \mathcal{P} . Given a finite set of vertices $X \subseteq V(\Gamma)$, we say that \mathcal{P} is *after X* if $x'_i \in T(R_i, X)$ and $x'_i P_i y_{\sigma(i)} S_{\sigma(i)}$ avoids X for all $i \in I$.

[†]A slightly weaker statement, without the additional condition that $H(R) \subseteq R$ appeared in [5, Lemma 1].

Lemma A.4.3 (Weak linking lemma). *Let Γ be a graph, $\epsilon \in \Omega(\Gamma)$ and let $n \in \mathbb{N}$. Then for any families $\mathcal{R} = (R_i: i \in [n])$ and $\mathcal{S} = (S_j: j \in [n])$ of vertex disjoint rays in ϵ and any finite set X of vertices, there is a linkage from \mathcal{R} to \mathcal{S} after X .*

Proof. Let us write x_i for the initial vertex of each R_i and let x'_i be the initial vertex of the tail $T(R_i, X)$. Furthermore, let $X' = X \cup \bigcup_{i \in [n]} R_i x'_i$. For $i \in [n]$ we will construct inductively finite disjoint connected subgraphs $K_i \subseteq \Gamma$ for each $i \in [n]$ such that

- K_i meets $T(S_j, X')$ and $T(R_j, X')$ for every $j \in [n]$;
- K_i avoids X' .

Suppose that we have recursively constructed K_1, \dots, K_{m-1} for some $m \leq n$. Let us write $X_m = X' \cup \bigcup_{i < m} V(K_i)$. Since R_1, \dots, R_n and S_1, \dots, S_n lie in the same end ϵ , there exist paths $Q_{i,j}$ between $T(R_i, X_m)$ and $T(S_j, X_m)$ avoiding X_m for all $i \neq j \in [n]$. Let $K_m = F \cup \bigcup_{i \neq j \in [n]} Q_{i,j}$, where F consists of an initial segment of each $T(R_i, X_m)$ sufficiently large to make K_m connected. Then it is clear that K_m is disjoint from all previous K_i and satisfies the claimed properties.

Let $K = \bigcup_{i=1}^n K_i$ and for each $j \in [n]$ let y_j be the initial vertex of $T(S_j, V(K))$. Note that by construction $T(S_j, V(K))$ avoids X for each j , since K_1 meets $T(S_j, X)$ and so $T(S_j, V(K)) \subseteq T(S_j, X)$.

We claim that there is no separator of size $< n$ between $\{x'_1, \dots, x'_n\}$ and $\{y_1, \dots, y_n\}$ in the subgraph $\Gamma' \subseteq \Gamma$ where $\Gamma' = K \cup \bigcup_{j=1}^n T(R_j, X') \cup T(S_j, X')$. Indeed, any set of $< n$ vertices must avoid at least one ray R_i , at least one graph K_m and one ray S_j . However, since K_m is connected and meets R_i and S_j , the separator does not separate x'_i from y_j .

Hence, by a version of Menger's theorem for infinite graphs [24, Proposition 8.4.1], there is a collection of n disjoint paths P_i from x'_i to $y_{\sigma(i)}$ in Γ' . Since Γ' is disjoint from X and meets each $R_i x'_i$ in x'_i only, it is clear that $\mathcal{P} = (P_i: i \in [n])$ is as desired. \square

In some cases we will need to find linkages between families of rays which avoid more than just a finite subset X . For this we will use the following lemma. Broadly the idea is that if we have a family of disjoint rays $(R_i: i \in [n])$ tending to an end ϵ , then there is some fixed number $N = N(n)$ such that if we have N disjoint graphs H_j , each with a specified ray S_j tending to ϵ , then we can 're-route' the

rays $(R_i: i \in [n])$ to some of the rays $(S_j: j \in [N])$, in such a way that we totally avoid one of the graphs H_j .

Lemma A.4.4 (Strong linking lemma). *Let Γ be a graph and $\epsilon \in \Omega(\Gamma)$. Let X be a finite set of vertices, $n \in \mathbb{N}$, and $\mathcal{R} = (R_i: i \in [n])$ a family of vertex disjoint rays in ϵ . Let x_i be the initial vertex of R_i and let x'_i the initial vertex of the tail $T(R_i, X)$.*

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 linkage $\mathcal{P} = (P_i: i \in [n])$ from \mathcal{R} to $(S_j: j \in [N])$ which is after X and such that the family

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

avoids H_ℓ .

Proof. Let $X' := X \cup \bigcup_{i \in [n]} R_i x'_i$ and let $N_0 := |X'|$. We claim that the lemma holds with $N := N_0 + n^3 + 1$.

Indeed suppose that $(H_j: j \in [N])$ is a collection of vertex disjoint subgraphs as in the statement of the lemma. Since the H_j are vertex disjoint, we may assume without loss of generality that the graphs in the family $(H_j: j \in [n^3 + 1])$ are disjoint from X' .

For each $i \in [n^2]$ we will build inductively finite, connected, vertex disjoint subgraphs \hat{K}_i such that

- \hat{K}_i meets $T(R_{i'}, X')$ for the $i' \in [n]$ with $i \equiv i' \pmod{n}$;
- \hat{K}_i meets exactly n of the H_j , that is $|\{j \in [n^3 + 1] : \hat{K}_i \cap H_j \neq \emptyset\}| = n$,
and
- \hat{K}_i avoids X' .

Suppose we have done so for all $i < m$. Let $X_m := X' \cup \bigcup_{i < m} V(\hat{K}_i)$ and let $m' \in [n]$ with $m' \equiv m \pmod{n}$. We will build inductively for $t = 0, \dots, n$ increasing connected subgraphs \hat{K}_m^t that meet $R_{m'}$, meet exactly t of the H_j , and avoid X_m .

We start with $\hat{K}_m^0 := \emptyset$. For each $t = 0, \dots, n-1$, if $T(R_{m'}, X_m)$ meets some H_j not met by \hat{K}_m^t then there is some initial vertex $z_t \in T(R_{m'}, X_m)$ where it does so

and we set $\hat{K}_m^{t+1} := \hat{K}_m^t \cup T(R_{m'}, X_m)z_t$. Otherwise we may assume $T(R_{m'}, X_m)$ does not meet any such H_j . In this case, let $j \in [n^3 + 1]$ be such that $\hat{K}_m^t \cap H_j = \emptyset$. Since $R_{m'}$ and S_j belong to the same end ϵ , there is some path P between $T(R_{m'}, X_m)$ and $T(S_j, X_m)$ which avoids X_m . Since this path meets some H_k with $k \in [n^3 + 1]$ which \hat{K}_m^t does not, there is some initial segment P' which meets exactly one such H_k . To form \hat{K}_m^{t+1} we add this path to \hat{K}_m^t together with an appropriately large initial segment of $T(R_{m'}, X_m)$ such that \hat{K}_m^{t+1} is connected. Finally we let $\hat{K}_m := \hat{K}_m^n$.

Let $K = \bigcup_{i \in [n^2]} \hat{K}_i$. Since each \hat{K}_i meets exactly n of the H_j , the set

$$J := \{j \in [n^3 + 1] : H_j \cap K \neq \emptyset\}$$

satisfies $|J| \leq n^3$. For each $j \in J$ let y_j be the initial vertex of $T(S_j, V(K))$.

We claim that there is no separator of size $< n$ between $\{x'_1, \dots, x'_n\}$ and $\{y_j : j \in J\}$ in the subgraph $\Gamma' \subseteq \Gamma$ where $\Gamma' := K \cup \bigcup_{j \in [n]} T(R_j, X') \cup \bigcup_{j \in J} H_j$. Suppose for a contradiction that there is such a separator S . Then S cannot meet every R_i for $i \in [n]$, and hence avoids R_q for some $q \in [n]$. Furthermore, there are n distinct \hat{K}_i such that $i \equiv q \pmod{n}$, all of which are disjoint. Hence there is some \hat{K}_r with $r \equiv q \pmod{n}$ disjoint from S . Finally, $|\{j \in J : \hat{K}_r \cap H_j \neq \emptyset\}| = n$ and so there is some H_s disjoint from S such that $\hat{K}_r \cap H_s \neq \emptyset$. Since \hat{K}_r meets $T(R_q, X')$ and H_s , there is a path from x'_q to y_s in Γ' , contradicting our assumption.

Hence, by a version of Menger's theorem for infinite graphs [24, Proposition 8.4.1], there is a family of disjoint paths $\mathcal{P} = (P_i : i \in [n])$ in Γ' from x'_i to $y_{\sigma(i)}$ for some suitable injective function $\sigma : [n] \rightarrow J$. Furthermore, since $|J| \leq n^3$ there is some $\ell \in [n^3 + 1]$ such that H_ℓ is disjoint from K .

Therefore, since Γ' is disjoint from X' and meets each $R_i x'_i$ in x'_i only, the family \mathcal{P} is a linkage from \mathcal{R} to $(S_j : j \in [n^3 + 1])$ which is after X such that

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

avoids H_ℓ . □

We will also need the following result, which allows us to work with paths instead of rays if the end ϵ is dominated by infinitely many vertices.

Lemma A.4.5. *Let Γ be a graph and ϵ an end of Γ which is dominated by infinitely many vertices. Let x_1, x_2, \dots, x_k be distinct vertices. If there are disjoint rays from the x_i to ϵ then there are disjoint paths from the x_i to distinct vertices y_i which dominate ϵ .*

Proof. We argue by induction on k . The base case $k = 0$ is trivial, so let us assume $k > 0$.

Consider any family of disjoint rays R_i , each from x_i to ϵ . Let y_k be any vertex dominating ϵ . Let P be a $y_k - \bigcup_{i=1}^k R_i$ -path. Without loss of generality the endvertex u of P in $\bigcup_{i=1}^k R_i$ lies on R_k . Then by the induction hypothesis applied to the graph $\Gamma - R_k u P$ we can find disjoint paths in that graph from the x_i with $i < k$ to vertices y_i which dominate ϵ . These paths together with $R_k u P$ then form the desired collection of paths. \square

To go back from paths to rays we will use the following lemma.

Lemma A.4.6. *Let Γ be a graph and ϵ an end of Γ which is dominated by infinitely many vertices. Let y_1, y_2, \dots, y_k be vertices, not necessarily distinct, dominating Γ . Then there are rays R_i from the respective y_i to ϵ which are disjoint except at their initial vertices.*

Proof. We recursively build for each $n \in \mathbb{N}$ paths P_1^n, \dots, P_k^n , each P_i^n from y_i to a vertex y_i^n dominating ϵ , disjoint except at their initial vertices, such that for $m < n$ each P_i^n properly extends P_i^m . We take P_i^0 to be a trivial path. For $n > 0$, build the P_i^n recursively in i : To construct P_i^n , we start by taking X_i^n to be the finite set of all the vertices of the P_j^n with $j < i$ or P_j^{n-1} with $j \geq i$. We then choose a vertex y_i^n outside of X_i^n which dominates ϵ and a path Q_i^n from y_i^{n-1} to y_i^n internally disjoint from X_i^n . Finally we let $P_i^n := P_i^{n-1} y_{n-1} Q_i^n$.

Finally, for each $i \leq k$, we let R_i be the ray $\bigcup_{n \in \mathbb{N}} P_i^n$. Then the R_i are disjoint except at their initial vertices, and they are in ϵ , since each of them contains infinitely many dominating vertices of ϵ . \square

A.5. G -tribes and concentration of G -tribes towards an end

In order to show that a given graph G is ubiquitous with respect to a fixed relation \triangleleft , we shall assume that $nG \triangleleft \Gamma$ for every $n \in \mathbb{N}$ and show that this implies that $\aleph_0 G \triangleleft \Gamma$. Since each subgraph witnessing that $nG \triangleleft \Gamma$ will be a collection of n disjoint subgraphs each being a witness for $G \triangleleft \Gamma$, it will be useful to introduce some notation for talking about these families of collections of n disjoint witnesses for each n .

To do this formally, we need to distinguish between a relation like the topological minor relation and the subdivision relation. Recall that we write $G \leq^* H$ if H is a subdivision of G and $G \leq \Gamma$ if G is a topological minor of Γ . We can interpret the topological minor relation as the composition of the subdivision relation and the subgraph relation.

Given two relations R and S , let their *composition* $S \circ R$ be the relation defined by $x(S \circ R)z$ if and only if there is a y such that xRy and ySz .

Hence we have that $G(\subseteq \circ \leq^*)\Gamma$ if and only if there exists H such that $G \leq^* H \subseteq \Gamma$, that is, if and only if $G \leq \Gamma$.

While in this paper we will only work with the topological minor relation, we will state the following definition and lemmas in greater generality, so that we may apply them in later papers [14, 15].

In general, we want to consider a pair $(\triangleleft, \blacktriangleleft)$ of binary relations of graphs with the following properties.

$$(R1) \quad \triangleleft = (\subseteq \circ \blacktriangleleft);$$

$$(R2) \quad \text{Given a set } I \text{ and a family } (H_i : i \in I) \text{ of pairwise disjoint graphs with } G \blacktriangleleft H_i \text{ for all } i \in I, \text{ then } |I| \cdot G \blacktriangleleft \bigcup \{H_i : i \in I\}.$$

We call a pair $(\triangleleft, \blacktriangleleft)$ with these properties *compatible*.

Other examples of compatible pairs are (\subseteq, \cong) , where \cong denotes the isomorphism relation, as well as (\preceq, \preceq^*) , where $G \preceq^* H$ if H is an inflated copy of G .

Next, we introduce the notion of a G -tribe \mathcal{F} for a compatible pair $(\triangleleft, \blacktriangleleft)$. In the setting of the topological minor relation (\leq, \leq^*) , a G -tribe \mathcal{F} in a graph Γ is a collection $\mathcal{F} = \{F : F \in \mathcal{F}\}$ where each $F \in \mathcal{F}$ is a finite set $F = \{H : H \in F\}$ of

pairwise disjoint subgraphs H of Γ such that each H is a subdivision of G , i.e. we have $G \leq^* H$, or more generally $G \blacktriangleleft H$. The sets F are the *layers* of \mathcal{F} , and the graphs H are the *members* of \mathcal{F} . Note that two members may intersect as long as they come from different layers.

We will also want to talk about *G-subtribes*, *G-tribes* which are in some way ‘contained’ in other *G-tribes*. Given a *G-tribe* \mathcal{F} , a simple notion of containment would be a *G-tribe* \mathcal{F}' such that each layer of \mathcal{F}' is a subset of a layer of \mathcal{F} . However, in our proof it will be necessary to consider a slightly more general notion of containment for tribes. In our running example, each layer F of \mathcal{F} consists of a collection $\{H : H \in F\}$ of disjoint subdivisions of G . However, each member $H \in F$ might itself contain non-trivial subgraphs which are also subdivisions of G . Our more general notion of containment would allow the members H' of a layer $F' \in \mathcal{F}'$ to be subgraphs of the members H of the corresponding layer $F \in \mathcal{F}$. The following definition makes this precise and introduces some related terminology.

Definition A.5.1 (*G-tribes*). Let G and Γ be graphs, and let $(\triangleleft, \blacktriangleleft)$ be a compatible pair of relations between graphs.

- A *G-tribe* in Γ (with respect to $(\triangleleft, \blacktriangleleft)$) is a collection \mathcal{F} of finite sets F of disjoint subgraphs H of Γ such that $G \blacktriangleleft H$ for each *member* $H \in \bigcup \mathcal{F}$ of \mathcal{F} .
- A *G-tribe* \mathcal{F} in Γ is called *thick*, if for each $n \in \mathbb{N}$ there exists a *layer* $F \in \mathcal{F}$ with $|F| \geq n$; otherwise, it is called *thin*.[‡]
- A *G-tribe* \mathcal{F}' in Γ is a *G-subtribe* of a *G-tribe* \mathcal{F} in Γ , denoted by $\mathcal{F}' \triangleleft \mathcal{F}$, if there is an injection $\Psi: \mathcal{F}' \rightarrow \mathcal{F}$ such that for each $F' \in \mathcal{F}'$ there is an injection $\varphi_{F'}: F' \rightarrow \Psi(F')$ such that $V(H') \subseteq V(\varphi_{F'}(H'))$ for each $H' \in F'$. The *G-subtribe* \mathcal{F}' is called *flat*, denoted by $\mathcal{F}' \subseteq \mathcal{F}$, if there is such an injection Ψ satisfying $F' \subseteq \Psi(F')$.
- For a *G-tribe* \mathcal{F} and a *G-subtribe* \mathcal{F}' of \mathcal{F} , we denote by $\mathcal{F} - \mathcal{F}'$ the *G-tribe* obtained by *removing* \mathcal{F}' from \mathcal{F} , that is the flat *G-subtribe* of \mathcal{F} consisting of the subsets of the layers $F \in \mathcal{F}$ obtained by deleting the image of $\varphi_{F'}$ from F whenever there is some $F' \in \mathcal{F}'$ with $\Psi(F') = F$.

[‡]A similar notion of *thick* and *thin families* was also introduced by Andreae in [5] (in German) and in [2]. The remaining notions, and in particular the concept of a *concentrated G-tribe*, which will be the backbone of essentially all our results in this and forthcoming papers, is new.

- 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} .

Hence, for a given compatible pair $(\triangleleft, \blacktriangleleft)$, if we wish to show that G is \triangleleft -ubiquitous, we will need to show that the existence of a thick G -tribe in Γ with respect to $(\triangleleft, \blacktriangleleft)$ implies $\aleph_0 G \triangleleft \Gamma$. We first observe that removing a thin G -tribe from a thick G -tribe always leaves a thick G -tribe.

Lemma A.5.2 (cf. [5, Lemma 3] or [2, Lemma 2]). *If \mathcal{F} is a thick G -tribe and \mathcal{F}' is a thin G -subtribe of \mathcal{F} , then $\mathcal{F} - \mathcal{F}'$ is a thick flat G -subtribe of \mathcal{F} .*

Proof. $\mathcal{F}'' := \mathcal{F} - \mathcal{F}'$ is obviously a flat subtribe of \mathcal{F} . As \mathcal{F}' is thin, there is a $k \in \mathbb{N}$ such that $|F'| \leq k$ for every $F' \in \mathcal{F}'$. Thus for each $F \in \mathcal{F}$, the set we delete from F , that is the image of $\phi(F')$ for some $F' \in \mathcal{F}'$ with $\Psi(F') = F$, has size at most k . As \mathcal{F} is thick, for every $n \in \mathbb{N}$ there is a layer $F \in \mathcal{F}$ satisfying $|F| \geq n + k$, and hence \mathcal{F}'' contains a layer of size at least n . \square

Given a thick G -tribe, the members of this tribe may have different properties, for example, some of them contain a ray belonging to a specific end ϵ of Γ whereas some of them do not. The next lemma allows us to restrict onto a thick subtribe, in which all members have the same properties, as long as we consider only finitely many properties. For example, we find a subtribe in which either all members contain an ϵ -ray, or none of them contain such a ray.

Lemma A.5.3 (Pigeon hole principle for thick G -tribes). *Suppose for some $k \in \mathbb{N}$, we have a k -colouring $c: \bigcup \mathcal{F} \rightarrow [k]$ 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} .*

Proof. Since \mathcal{F} is a thick G -tribe, there is a sequence $(n_i : i \in \mathbb{N})$ of natural numbers and a sequence $(F_i \in \mathcal{F} : i \in \mathbb{N})$ such that

$$n_1 \leq |F_1| < n_2 \leq |F_2| < n_3 \leq |F_3| < \dots .$$

Now for each i , by pigeon hole principle, there is one colour $c_i \in [k]$ such that the subset $F'_i \subseteq F_i$ of elements of colour c_i has size at least n_i/k . Moreover, since $[k]$

is finite, there is one colour $c^* \in [k]$ and an infinite subset $I \subseteq \mathbb{N}$ such that $c_i = c^*$ for all $i \in I$. But this means that $\mathcal{F}' := \{F'_i : i \in I\}$ is a monochromatic, thick, flat G -subtribe. \square

Given a connected graph G and a compatible pair of relations $(\triangleleft, \blacktriangleleft)$ we say that a G -tribe \mathcal{F} w.r.t. $(\triangleleft, \blacktriangleleft)$ is *connected* if every member H of \mathcal{F} is connected. Note that for relations \blacktriangleleft like \cong, \leq^*, \preceq^* , if G is connected and $G \blacktriangleleft H$, then H is connected. In this case, any G -tribe will be connected.

Lemma A.5.4. *Let G be a connected graph (of arbitrary cardinality), $(\triangleleft, \blacktriangleleft)$ a compatible pair of relations of graphs and Γ a graph containing a thick connected G -tribe \mathcal{F} w.r.t. $(\triangleleft, \blacktriangleleft)$. Then either $\aleph_0 G \triangleleft \Gamma$, 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 .*

Proof. For every finite vertex set $X \subseteq V(\Gamma)$, only a thin subtribe of \mathcal{F} can meet X . Hence, by Lemma A.5.2, the tribe \mathcal{F}'' obtained by removing this thin subtribe from \mathcal{F} is a thick flat subtribe of \mathcal{F} which is contained in the graph $\Gamma - X$. Since each member of \mathcal{F}'' is connected, any member H of \mathcal{F}'' is contained in a unique component of $\Gamma - X$. If for any X , infinitely many components of $\Gamma - X$ contain a \blacktriangleleft -copy of G , the union of all these copies is a \blacktriangleleft -copy of $\aleph_0 G$ in Γ by (R2), hence $\aleph_0 G \triangleleft \Gamma$. Thus, we may assume that for each X , only finitely many components contain elements from \mathcal{F}'' , and hence, by colouring each H with a colour corresponding to the component of $\Gamma - X$ containing it, we may assume by the pigeon hole principle for G -tribes, Lemma A.5.3, that at least one component of $\Gamma - X$ contains a thick flat subtribe of \mathcal{F} .

Let $C_0 = \Gamma$ and $\mathcal{F}_0 = \mathcal{F}$ and consider the following recursive process: If possible, we choose a finite vertex set X_n in C_n such that there are two components $C_{n+1} \neq D_{n+1}$ of $C_n - X_n$ where C_{n+1} contains a thick flat subtribe \mathcal{F}_{n+1} of \mathcal{F}_n and D_{n+1} contains at least one \blacktriangleleft -copy H_{n+1} of G . Since by construction all H_n are pairwise disjoint, we either find infinitely many such H_n and thus, again by (R2), an $\aleph_0 G \triangleleft \Gamma$, or our process terminates at step N say. That is, we have a thick flat subtribe \mathcal{F}_N contained in a subgraph C_N such that there is no finite vertex set X_N satisfying the above conditions.

Let $\mathcal{F}' := \mathcal{F}_N$. We claim that for every finite vertex set X of Γ , there is a unique component C_X of $\Gamma - X$ that contains a thick flat G -subtribe of \mathcal{F}' . Indeed, note that if for some finite $X \subseteq \Gamma$ there are two components C and C'

of $\Gamma - X$ both containing thick flat G -subtribes of \mathcal{F}' , then since every G -copy in \mathcal{F}' is contained in C_N , it must be the case that $C \cap C_N \neq \emptyset \neq C' \cap C_N$. But then $X_N = X \cap C_N \neq \emptyset$ is a witness that our process could not have terminated at step N .

Next, observe that whenever $X' \supseteq X$, then $C_{X'} \subseteq C_X$. By a theorem of Diestel and Kühn, [29], it follows that there is a unique end ϵ of Γ such that $C(X, \epsilon) = C_X$ for all finite $X \subseteq \Gamma$. It now follows easily from the uniqueness of $C_X = C(X, \epsilon)$ that \mathcal{F}' is concentrated at this ϵ . \square

We note that concentration towards an end ϵ is a robust property in the following sense:

Lemma A.5.5. *Let G be a connected graph (of arbitrary cardinality), $(\triangleleft, \blacktriangleleft)$ a compatible pair of relations of graphs and Γ a graph containing a thick connected G -tribe \mathcal{F} w.r.t. $(\triangleleft, \blacktriangleleft)$ 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 , too.*

Proof. Let X be a finite vertex set. By definition, if the G -tribe \mathcal{F} is concentrated at ϵ , then \mathcal{F} is thick, and the subtribe \mathcal{F}_X consisting of for each $F \in \mathcal{F}$ of the sets

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

is a thin subtribe of \mathcal{F} , i.e. there exists $k \in \mathbb{N}$ such that $|F_X| \leq k$ for all $F_X \in \mathcal{F}_X$.

For (1), observe that the tribe \mathcal{F}' obtained by removing this thin subtribe from \mathcal{F} is a thick flat subtribe of \mathcal{F} by Lemma A.5.2, and all its members are contained in $C(X, \epsilon)$ by construction.

For (2), observe that if \mathcal{F}' is a subtribe of \mathcal{F} , then for every $F' \in \mathcal{F}'$ there is an injection $\varphi_{F'} : F' \rightarrow F$ for some $F \in \mathcal{F}$. Therefore, $|\varphi_{F'}^{-1}(F_X)| \leq k$ for $F_X \subseteq F$ as defined above, and so only a thin subtribe of \mathcal{F}' is not contained in $C(X, \epsilon)$. \square

A.6. Countable subtrees

In this section we prove Theorem A.1.2. Let S be a countable subtree of T . Our aim is to construct an S -horde $(Q_i : i \in \mathbb{N})$ of disjoint suitable subdivisions of S in

Γ inductively. By Lemma A.5.4, we may assume without loss of generality that there are an end ϵ of Γ and a thick T -tribe \mathcal{F} concentrated at ϵ .

In order to ensure that we can continue the construction at each stage, we will require the existence of additional structure for each n . But the details of what additional structure we use will vary depending on how many vertices dominate ϵ . So, after a common step of preprocessing, in Section A.6.1, the proof of Theorem A.1.2 splits into two cases according to whether the number of ϵ -dominating vertices in Γ is finite (Section A.6.2) or infinite (Section A.6.3).

A.6.1. Preprocessing

We begin by picking a root v for S , and also consider T as a rooted tree with root v . Let $V_\infty(S)$ be the set of vertices of infinite degree in S .

Definition A.6.1. Given S and T as above, define a spanning locally finite forest $S^* \subseteq S$ by

$$S^* := S \setminus \bigcup_{t \in V_\infty(S)} \{tt_i : t_i \in N^+(t), i > N_t\},$$

where N_t is as in Corollary A.3.8. We will also consider every component of S^* as a rooted tree given by the induced tree order from T .

Definition A.6.2. An edge e of S^* is an *extension edge* if there is a ray in S^* starting at e^+ which displays self-similarity of T . For each extension edge e we fix one such a ray R_e . Write $\text{Ext}(S^*) \subseteq E(S^*)$ for the set of extension edges.

Consider the forest $S^* - \text{Ext}(S^*)$ obtained from S^* by removing all extension edges. Since every ray in S^* must contain an extension edge by Corollary A.3.10, each component of $S^* - \text{Ext}(S^*)$ is a locally finite rayless tree and so is finite (this argument is inspired by [5, Lemma 2]). We enumerate the components of $S^* - \text{Ext}(S^*)$ as S_0^*, S_1^*, \dots in such a way that for every n , the set

$$S_n := S \left[\bigcup_{i \leq n} V(S_i^*) \right]$$

is a finite subtree of S containing the root r . Let us write $\partial(S_n) = E_{S^*}(S_n, S^* \setminus S_n)$, and note that $\partial(S_n) \subseteq \text{Ext}(S^*)$. We make the following definitions:

- For a given T -tribe \mathcal{F} and ray R of T , we say that R *converges to ϵ according to \mathcal{F}* if for all members H of \mathcal{F} the ray $H(R)$ is in ϵ . We say that R is *cut from ϵ according to \mathcal{F}* if for all members H of \mathcal{F} the ray $H(R)$ is not in ϵ . Finally we say that \mathcal{F} *determines whether R converges to ϵ* if either R converges to ϵ according to \mathcal{F} or R is cut from ϵ according to \mathcal{F} .
- Similarly, for a given T -tribe \mathcal{F} and vertex t of T , we say that t *dominates ϵ according to \mathcal{F}* if for all members H of \mathcal{F} the vertex $H(t)$ dominates ϵ . We say that t is *cut from ϵ according to \mathcal{F}* if for all members H of \mathcal{F} the vertex $H(t)$ does not dominate ϵ . Finally we say that \mathcal{F} *determines whether t dominates ϵ* if either t dominates ϵ according to \mathcal{F} or t is cut from ϵ according to \mathcal{F} .
- Given $n \in \mathbb{N}$, we say a thick T -tribe \mathcal{F} *agrees about $\partial(S_n)$* if for each extension edge $e \in \partial(S_n)$, it determines whether R_e converges to ϵ . We say that it *agrees about $V(S_n)$* if for each vertex t of S_n , it determines whether t dominates ϵ .
- Since $\partial(S_n)$ and $V(S_n)$ are finite for all n , it follows from Lemma A.5.3 that given some $n \in \mathbb{N}$, any thick T -tribe has a flat thick T -subtribe \mathcal{F} such that \mathcal{F} agrees about $\partial(S_n)$ and $V(S_n)$. Under these circumstances we set

$$\begin{aligned} \partial_\epsilon(S_n) &:= \{e \in \partial(S_n) : R_e \text{ converges to } \epsilon \text{ according to } \mathcal{F}\}, \\ \partial_{-\epsilon}(S_n) &:= \{e \in \partial(S_n) : R_e \text{ is cut from } \epsilon \text{ according to } \mathcal{F}\}, \\ V_\epsilon(S_n) &:= \{t \in V(S_n) : t \text{ dominates } \epsilon \text{ according to } \mathcal{F}\}, \text{ and} \\ V_{-\epsilon}(S_n) &:= \{t \in V(S_n) : t \text{ is cut from } \epsilon \text{ according to } \mathcal{F}\}. \end{aligned}$$

- Also, under these circumstances, let us write $S_n^{-\epsilon}$ for the component of the forest $S - \partial_\epsilon(S_n) - \{e \in E_S(S_n, S \setminus S_n) : e^- \in V_\epsilon(S_n)\}$ containing the root of S . Note that $S_n \subseteq S_n^{-\epsilon}$.

The following lemma contains a large part of the work needed for our inductive construction.

Lemma A.6.3 (*T -tribe refinement lemma*). *Suppose we have a thick T -tribe \mathcal{F}_n concentrated at ϵ which agrees about $\partial(S_n)$ and $V(S_n)$ for some $n \in \mathbb{N}$. Let f denote the unique edge from S_n to $S_{n+1} \setminus S_n$. Then there is a thick T -tribe \mathcal{F}_{n+1} concentrated at ϵ with the following properties:*

(i) \mathcal{F}_{n+1} agrees about $\partial(S_{n+1})$ and $V(S_{n+1})$.

(ii) $\mathcal{F}_{n+1} \cup \mathcal{F}_n$ agree about $\partial(S_n) \setminus \{f\}$ and $V(S_n)$.

(iii) $S_{n+1}^{-\epsilon} \supseteq S_n^{-\epsilon}$.

(iv) For all $H \in \mathcal{F}_{n+1}$ there is a finite $X \subseteq \Gamma$ such that

$$H(S_{n+1}^{-\epsilon}) \cap (X \cup C_\Gamma(X, \epsilon)) = H(V_\epsilon(S_{n+1})).$$

Moreover, if $f \in \partial_\epsilon(S_n)$, and $R_f = v_0 v_1 v_2 \dots \subseteq S^*$ (with $v_0 = f^+$) denotes the ray displaying self-similarity of T at f , then we may additionally assume:

(v) For every $H \in \mathcal{F}_{n+1}$ and every $k \in \mathbb{N}$, there is $H' \in \mathcal{F}_{n+1}$ with

- $H' \subseteq_r H$,
- $H'(S_n) = H(S_n)$,
- $H'(T_{v_0}) \subseteq_r H(T_{v_k})$, and
- $H'(R_f) \subseteq H(R_f)$.

Proof. Concerning (v), if $f \in \partial_\epsilon(S_n)$ recall that according to Definition A.6.2, the ray R_f satisfies that for all $k \in \mathbb{N}$ we have $T_{v_0} \leq_r T_{v_k}$ such that R_f gets embedded into itself. In particular, there is a subtree \hat{T}_1 of T_{v_1} which is a rooted subdivision of T_{v_0} with $\hat{T}_1(R_f) \subseteq R_f$, considering \hat{T}_1 as a rooted tree given by the tree order in T_{v_1} . Let us recursively define \hat{T}_k for each $k \in \mathbb{N}$ as the corresponding subdivision of \hat{T}_1 (viewed as a subgraph of T_{v_0}) in \hat{T}_{k-1} (viewed as a subdivision of T_{v_0}), that is $\hat{T}_k := \hat{T}_{k-1}(\hat{T}_1)$. Then it is clear that $(\hat{T}_k : k \in \mathbb{N})$ is a family of rooted subdivisions of T_{v_0} such that for each $k \in \mathbb{N}$

- $\hat{T}_k \subseteq T_{v_k}$;
- $\hat{T}_k \supseteq \hat{T}_{k+1}$; and
- $\hat{T}_k(R_f) \subseteq R_f$.

Hence, for every subdivision H of T with $H \in \bigcup \mathcal{F}_n$ and every $k \in \mathbb{N}$, the subgraph $H(\hat{T}_k)$ is also a rooted subdivision of T_{v_0} . Let us construct a subdivision $H^{(k)}$ of T by letting $H^{(k)}$ be the minimal subtree of H containing $H(T \setminus T_{v_0}) \cup H(\hat{T}_k)$, where $H^{(k)}(T_{v_0}) = H(\hat{T}_k)$ and $H^{(k)}(T \setminus T_{v_0}) = H(T \setminus T_{v_0})$. Note that

$$H^{(k)}(T_{v_0}) = H(\hat{T}_k) \subseteq_r H^{(k-1)}(T_{v_0}) = H(\hat{T}_{k-1}) \subseteq_r \dots \subseteq_r H(T_{v_k}).$$

In particular, for every subdivision $H \in \bigcup \mathcal{F}_n$ of T and every $k \in \mathbb{N}$, there is a subdivision $H^{(k)} \subseteq H$ of T such that $H^{(k)}(S_n^{-\epsilon}) = H(S_n^{-\epsilon})$, $H^{(k)}(T_{v_0}) \subseteq_r H(T_{v_k})$ and $H^{(k)}(R_f) \subseteq H(R_f)$. By the pigeon hole principle, there is an infinite index set $K_H = \{k_i^H : i \in \mathbb{N}\} \subseteq \mathbb{N}$ such that $\{H^{(k)} : k \in K_H\}$ agrees about $\partial(S_{n+1})$. With

$$F'_i := \{H^{(k_i^H)} : H \in F\}, \quad (\dagger)$$

consider the thick subtribe $\mathcal{F}'_n := \{F'_i : F \in \mathcal{F}_n, i \in \mathbb{N}\}$ of \mathcal{F}_n . Observe that $\mathcal{F}'_n \cup \mathcal{F}_n$ still agrees about $\partial(S_n)$ and $V(S_n)$. (If $f \in \partial_{-\epsilon}(S_n)$, then skip this part and simply let $\mathcal{F}'_n := \mathcal{F}_n$.)

Concerning (iii), observe that for every $H \in \bigcup \mathcal{F}'_n$, since the rays $H(R_e)$ for $e \in \partial_{-\epsilon}(S_n)$ do not tend to ϵ , there is a finite vertex set X_H such that $H(R_e) \cap C(X_H, \epsilon) = \emptyset$ for all $e \in \partial_{-\epsilon}(S_n)$. Furthermore, since X_H is finite, for each such extension edge e there exists $x_e \in R_e$ such that

$$H(T_{x_e}) \cap (X_H \cup C(X_H, \epsilon)) = \emptyset.$$

By definition of extension edges, cf. Definition A.6.2, for each $e \in \partial_{-\epsilon}(S_n)$ there is a rooted embedding of T_{e^+} into $H(T_{x_e})$. Hence, there is a subdivision \tilde{H} of T with $\tilde{H} \leq H$ and $\tilde{H}(S_n) = H(S_n)$ such that $\tilde{H}(T_{e^+}) \subseteq H(T_{x_e})$ for each $e \in \partial_{-\epsilon}(S_n)$.

Note that if $e \in \partial_{-\epsilon}(S_n)$ and g is an extension edge with $e \leq g \in \partial(S_{n+1}) \setminus \partial(S_n)$, then $\tilde{H}(R_g) \subseteq \tilde{H}(S_{e^+}) \subseteq H(S_{x_e})$, and so

$$\tilde{H}(R_g) \text{ does not tend to } \epsilon. \quad (\ddagger)$$

Define $\tilde{\mathcal{F}}_n$ to be the thick T -subtribe of \mathcal{F}'_n consisting of the \tilde{H} for every $H \in \bigcup \mathcal{F}'_n$.

Now use Lemma A.5.3 to choose a maximal thick flat subtribe \mathcal{F}_n^* of $\tilde{\mathcal{F}}_n$ which agrees about $\partial(S_{n+1})$ and $V(S_{n+1})$, so it satisfies (i) and (ii). By (\ddagger), the tribe \mathcal{F}_n^* satisfies (iii), and by maximality and (\dagger), it satisfies (v).

In our last step, we now arrange for (iv) while preserving all other properties. For each $H \in \bigcup \mathcal{F}_n^*$, since $H(S_{n+1})$ is finite, we may find a finite separator Y_H such that

$$H(S_{n+1}) \cap (Y_H \cup C(Y_H, \epsilon)) = H(V_\epsilon(S_{n+1})).$$

Since Y_H is finite, for every vertex $t \in V_{-\epsilon}(S_{n+1})$, say with $N^+(t) = (t_i)_{i \in \mathbb{N}}$, there exists $n_t \in \mathbb{N}$ such that $C(Y_H, \epsilon) \cap H(T_{t_j}) = \emptyset$ for all $j \geq n_t$. Using Corollary A.3.8, for every such t there is a rooted embedding

$$\{t\} \cup \bigcup_{j > n_t} T_{t_j} \leq_r \{t\} \cup \bigcup_{j > n_t} T_{t_j}.$$

fixing the root t . Hence there is a subdivision H' of T with $H' \leq H$ such that $H'(T \setminus S) = H(T \setminus S)$ and for every $t \in V_{-\epsilon}(S_{n+1})$

$$H' \left[\{t\} \cup \bigcup_{j > N_t} T_{t_j} \right] \cap C(Y_H, \epsilon) = \emptyset.$$

Moreover, note that by construction of \tilde{F}_n , every such H' automatically satisfies that

$$H(S_{e^+}) \cap C(X_H \cup Y_H, \epsilon) = \emptyset$$

for all $e \in \partial_{-\epsilon}(S_{n+1})$. Let \mathcal{F}_{n+1} consist of the set of H' as defined above for all $H \in \mathcal{F}_n^*$. Then $X_H \cup Y_H$ is a finite separator witnessing that \mathcal{F}_{n+1} satisfies (iv). \square

A.6.2. Only finitely many vertices dominate ϵ

We first note as in Lemma A.5.4, that for every finite vertex set $X \subseteq V(\Gamma)$ only a thin subtribe of \mathcal{F} can meet X , so a thick subtribe is contained in the graph $\Gamma - X$. By removing the set of vertices dominating ϵ , we may therefore assume without loss of generality that no vertex of Γ dominates ϵ .

Definition A.6.4 (Bounder, extender). Suppose that some thick T -tribe \mathcal{F} which is concentrated at ϵ agrees about $\partial(S_r)$ for some given $r \in \mathbb{N}$, and that Q_1, Q_2, \dots, Q_s are disjoint subdivisions of $S_r^{-\epsilon}$ for some given $s \in \mathbb{N}$, (note, $S_r^{-\epsilon}$ depends on \mathcal{F}).

- A *bounder* for the $(Q_i : i \in [s])$ is a finite set X of vertices in Γ separating all the Q_i from ϵ , i.e. such that

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

- An *extender* for the $(Q_i : i \in [s])$ is a family $\mathcal{E} = (E_{e,i} : e \in \partial_\epsilon(S_r), i \in [s])$ of rays in Γ tending to ϵ which are disjoint from each other and also from each Q_i except at their initial vertices, and where the start vertex of $E_{e,i}$ is $Q_i(e^-)$.

To prove Theorem A.1.2, we now assume inductively that for some $n \in \mathbb{N}$, with $r := \lfloor n/2 \rfloor$ and $s := \lceil n/2 \rceil$ we have:

- (1) A thick T -tribe \mathcal{F}_r in Γ concentrated at ϵ which agrees about $\partial(S_r)$, with a boundary $\partial_\epsilon(S_r)$ such that $S_{r-1}^{-\epsilon} \subseteq S_r^{-\epsilon}$.[§]
- (2) a family $(Q_i^n : i \in [s])$ of s pairwise disjoint T -suitable subdivisions of $S_r^{-\epsilon}$ in Γ with $Q_i^n(S_{r-1}^{-\epsilon}) = Q_i^{n-1}$ for all $i \leq s-1$,
- (3) a bounder X_n for the $(Q_i^n : i \in [s])$, and
- (4) an extender $\mathcal{E}_n = (E_{e,i}^n : e \in \partial_\epsilon(S_r^{-\epsilon}), i \in [s])$ for the $(Q_i^n : i \in [s])$.

The base case $n = 0$ is easy, as we simply may choose $\mathcal{F}_0 \leq_r \mathcal{F}$ to be any thick T -subtribe in Γ which agrees about $\partial(S_0)$, and let all other objects be empty.

So, let us assume that our construction has proceeded to step $n \geq 0$. Our next task splits into two parts: First, if $n = 2k - 1$ is odd, we extend the already existing k subdivisions $(Q_i^n : i \in [k])$ of $S_{k-1}^{-\epsilon}$ to subdivisions $(Q_i^{n+1} : i \in [k])$ of $S_k^{-\epsilon}$. And secondly, if $n = 2k$ is even, we construct a further disjoint copy Q_{k+1}^{n+1} of $S_k^{-\epsilon}$.

Construction part 1: $n = 2k - 1$ is odd. By assumption, \mathcal{F}_{k-1} agrees about $\partial(S_{k-1})$. Let f denote the unique edge from S_{k-1} to $S_k \setminus S_{k-1}$. We first apply Lemma A.6.3 to \mathcal{F}_{k-1} in order to find a thick T -tribe \mathcal{F}_k concentrated at ϵ satisfying properties (i)–(v). In particular, \mathcal{F}_k agrees about $\partial(S_k)$ and $S_{k-1}^{-\epsilon} \subseteq S_k^{-\epsilon}$.

We first note that if $f \notin \partial_\epsilon(S_{k-1})$, then $S_{k-1}^{-\epsilon} = S_k^{-\epsilon}$, and we can simply take $Q_i^{n+1} := Q_i^n$ for all $i \in [k]$, $\mathcal{E}_{n+1} := \mathcal{E}_n$ and $X_{n+1} := X_n$.

Otherwise, we have $f \in \partial_\epsilon(S_{k-1})$. By Lemma A.5.5(2) \mathcal{F}_k is concentrated at ϵ , and so we may pick a collection $\{H_1, \dots, H_N\}$ of disjoint subdivisions of T from some $F \in \mathcal{F}_k$, all of which are contained in $C(X_n, \epsilon)$, where $N = |\mathcal{E}_n|$. By Lemma A.4.3 there is some linkage $\mathcal{P} \subseteq C(X_n, \epsilon)$ from \mathcal{E}_n to $(H_j(R_f) : j \in [N])$, which is after X_n . Let us suppose that the linkage \mathcal{P} joins a vertex $x_{e,i} \in E_{e,i}^n$ to $y_{\sigma(e,i)} \in H_{\sigma(e,i)}(R_f)$ via a path $P_{e,i} \in \mathcal{P}$. Let $z_{\sigma(e,i)}$ be a vertex in R_f such that $y_{\sigma(e,i)} \leq H_{\sigma(e,i)}(z_{\sigma(e,i)})$ in the tree order on $H_{\sigma(e,i)}(T)$.

By property (v) of \mathcal{F}_k in Lemma A.6.3, we may assume without loss of generality that for each H_j there is another member $H'_j \subseteq H_j$ of \mathcal{F}_k such that $H'_j(T_{f^+}) \subseteq_r H_j(T_{z_j})$. Let $\hat{P}_j \subseteq H'_j$ denote the path from $H_j(y_j)$ to $H'_j(f^+)$.

Now for each $i \in [k]$, define

$$Q_i^{n+1} = Q_i^n \cup E_{f,i}^n x_{f,i} P_{f,i} y_{\sigma(f,i)} \hat{P}_{\sigma(f,i)} \cup H'_{\sigma(f,i)}(S_k^{-\epsilon} \setminus S_{k-1}^{-\epsilon}).$$

[§]Note that since ϵ is undominated, every thick T -tribe agrees about the fact that $V_\epsilon(S_i) = \emptyset$ for all $i \in \mathbb{N}$.

By construction, each Q_i^{n+1} is a T -suitable subdivision of $S_k^{-\epsilon}$.

By Lemma A.6.3(iv) we may find a finite set $X_{n+1} \subseteq \Gamma$ with $X_n \subseteq X_{n+1}$ such that

$$C(X_{n+1}, \epsilon) \cap \left(\bigcup_{i \in [k]} Q_i^{n+1} \right) = \emptyset.$$

This set X_{n+1} will be our bounder.

Define an extender $\mathcal{E}_{n+1} = (E_{e,i}^{n+1} : e \in \partial_\epsilon(S_k), i \in [k])$ for the Q_i^{n+1} as follows:

- For $e \in \partial_\epsilon(S_{k-1}) \setminus \{f\}$, let $E_{e,i}^{n+1} := E_{e,i}^n x_{e,i} P_{e,i} y_{\sigma(e,i)} H_{\sigma(e,i)}(R_f)$.
- For $e \in \partial_\epsilon(S_k) \setminus \partial(S_{k-1})$, let $E_{e,i}^{n+1} := H'_{\sigma(e,i)}(R_e)$.

Since each $H_{\sigma(e,i)}, H'_{\sigma(e,i)} \in \bigcup \mathcal{F}_k$, and \mathcal{F}_k determines that R_f converges to ϵ , these are indeed ϵ rays. Furthermore, since $H'_{\sigma(e,i)} \subseteq H_{\sigma(e,i)}$ and $\{H_1, \dots, H_N\}$ are disjoint, it follows that the rays are disjoint.

Construction part 2: $n = 2k$ is even. If $\partial_\epsilon(S_k) = \emptyset$, then $S_k^{-\epsilon} = S$, and so picking any element Q_{k+1}^{n+1} from \mathcal{F}_k with $Q_{k+1}^{n+1} \subseteq C(X_n, \epsilon)$ gives us a further copy of S disjoint from all the previous ones. Using Lemma A.6.3(iv), there is a suitable bounder $X_{n+1} \supseteq X_n$ for Q_{k+1}^{n+1} , and we are done. Otherwise, pick $e_0 \in \partial_\epsilon(S_k)$ arbitrary.

Since \mathcal{F}_k is concentrated at ϵ , we may pick a collection $\{H_1, \dots, H_N\}$ of disjoint subdivisions of T from \mathcal{F}_k all contained in $C(X_n, \epsilon)$, where N is large enough so that we may apply Lemma A.4.4 to find a linkage $\mathcal{P} \subseteq C(X_n, \epsilon)$ from \mathcal{E}_n to $(H_i(R_{e_0}) : i \in [N])$ after X_n , avoiding say H_1 . Let us suppose the linkage \mathcal{P} joins a vertex $x_{e,i} \in E_{e,i}^n$ to $y_{\sigma(e,i)} \in H_{\sigma(e,i)}(R_{e_0})$ via a path $P_{e,i} \in \mathcal{P}$. Define

$$Q_{k+1}^{n+1} = H_1(S_k^{-\epsilon}).$$

Note that Q_{k+1}^{n+1} is a T -suitable subdivision of $S_k^{-\epsilon}$.

By Lemma A.6.3(iv) there is a finite set $X_{n+1} \subseteq \Gamma$ with $X_n \subseteq X_{n+1}$ such that

$$C(X_{n+1}, \epsilon) \cap Q_{k+1}^{n+1} = \emptyset.$$

This set X_{n+1} will be our new bounder.

Define the extender $\mathcal{E}_{n+1} = (E_{e,i}^{n+1} : e \in \partial_\epsilon(S_{k+1}), i \in [k+1])$ of ϵ -rays as follows:

- For $i \in [k]$, let $E_{e,i}^{n+1} := E_{e,i}^n x_{e,i} P_{e,i} y_{\sigma(e,i)} H_{\sigma(e,i)}(R_{e_0})$.
- For $i = k+1$, let $E_{e,k+1}^{n+1} := H_1(R_e)$ for all $e \in \partial_\epsilon(S_{k+1})$.

Once the construction is complete, let us define $H_i := \bigcup_{n \geq 2i-1} Q_i^n$.

Since $\bigcup_{n \in \mathbb{N}} S_n^{-\epsilon} = S$, and due to the extension property (2), the collection $(H_i : i \in \mathbb{N})$ is an S -horde. \square

We remark that our construction so far suffices to give a complete proof that countable trees are \leq -ubiquitous. Indeed, it is well-known that an end of Γ is dominated by infinitely many distinct vertices if and only if Γ contains a subdivision of K_{\aleph_0} [24, Exercise 19, Chapter 8], in which case proving ubiquity becomes trivial:

Lemma A.6.5. *For any countable graph G , we have $\aleph_0 \cdot G \subseteq K_{\aleph_0}$.*

Proof. By partitioning the vertex set of K_{\aleph_0} into countably many infinite parts, we see that $\aleph_0 \cdot K_{\aleph_0} \subseteq K_{\aleph_0}$. Also, clearly $G \subseteq K_{\aleph_0}$ holds. Hence, we can conclude that $\aleph_0 \cdot G \subseteq \aleph_0 \cdot K_{\aleph_0} \subseteq K_{\aleph_0}$. \square

A.6.3. Infinitely many vertices dominate ϵ

The argument in this case is very similar to that in the previous subsection. We define bounders and extenders just as before. We once more assume inductively that for some $n \in \mathbb{N}$, with $r := \lfloor n/2 \rfloor$, we have objects given by (1)–(4) as in the last section, and which in addition satisfy

(5) \mathcal{F}_r agrees about $V(S_r)$.

(6) For any $t \in V_\epsilon(S_r)$ the vertex $Q_i^n(t)$ dominates ϵ .

The base case is again trivial, so suppose that our construction has proceeded to step $n \geq 0$. The construction is split into two parts just as before, where the case $n = 2k$, in which we need to refine our T -tribe and find a new copy Q_{k+1}^{n+1} of $S_k^{-\epsilon}$, proceeds just as in the last section.

If $n = 2k - 1$ is odd, and if $f \in \partial_{-\epsilon}(S_{k-1})$ or $\partial_\epsilon(S_{k-1})$, then we proceed as in the last subsection. But these are no longer the only possibilities. It follows from the definition of $S_k^{-\epsilon}$ that there is one more option, namely that $f^- \in V_\epsilon(S_k)$. In this case we modify the steps of the construction as follows:

We first apply Lemma A.6.3 to \mathcal{F}_{k-1} in order to find a thick T -tribe \mathcal{F}_{k-1} which agrees about $\partial(S_k)$ and $V(S_k)$.

Then, by applying Lemma A.4.5 to tails of the rays $E_{e,i}^n$ in $C_\Gamma(X_n, \epsilon)$, we obtain a family \mathcal{P}_{n+1} of paths $P_{e,i}^{n+1}$ which are disjoint from each other and from the Q_i^n

except at their initial vertices, where the initial vertex of $P_{e,i}^{n+1}$ is $Q_i^n(e^-)$ and the final vertex $y_{e,i}^{n+1}$ of $P_{e,i}^{n+1}$ dominates ϵ .

Since \mathcal{F}_k is concentrated at ϵ , we may pick a collection $\{H_1, \dots, H_k\}$ of disjoint subdivisions of T from \mathcal{F}_k all contained in $C(X_n \cup \bigcup \mathcal{P}_{n+1}, \epsilon)$.

Now for each $i \in [k]$, define

$$\hat{Q}_i^{n+1} = Q_i^n \cup H_i(f^-) \cup H_i(S_k^{-\epsilon} \setminus S_{k-1}^{-\epsilon}).$$

These are almost T -suitable subdivisions of $S_k^{-\epsilon}$, except we need to add a path between $Q_i^n(f^-)$ and $H_i(f^-)$.

By applying Lemma A.4.5 to tails of the rays $H_i(R_e)$ inside $C(X_n \cup \bigcup \mathcal{P}_{n+1}, \epsilon)$ with $e \in \partial_\epsilon(S_{k+1}) \setminus \partial(S_k)$ we construct a family $\mathcal{P}'_{n+1} := \{P_{e,i}^{n+1} : e \in \partial_\epsilon(S_{k+1}) \setminus \partial_\epsilon(S_k), i \leq k\}$ of paths which are disjoint from each other and from the \hat{Q}_i^{n+1} except at their initial vertices, where the initial vertex of $P_{e,i}^{n+1}$ is $H_i(e^-)$ and the final vertex $y_{e,i}^{n+1}$ of $P_{e,i}^{n+1}$ dominates ϵ . Therefore the family

$$\mathcal{P}_{n+1} \cup \mathcal{P}'_{n+1} = (P_{e,i}^{n+1} : e \in \partial_\epsilon(S_{k+1}), i \in [k])$$

is a family of disjoint paths, which are also disjoint from the \hat{Q}_i^{n+1} except at their initial vertices, where the initial vertex of $P_{e,i}^{n+1}$ is $H_i(e^-)$ or $Q_i^n(e^-)$ and the final vertex $y_{e,i}^{n+1}$ of $P_{e,i}^{n+1}$ dominates ϵ .

Since $Q_i^n(f^-)$ and $H_i(f^-)$ both dominate ϵ for all i , we may recursively build a sequence $\hat{\mathcal{P}}_{n+1} = \{\hat{P}_i : 1 \leq i \leq k\}$ of disjoint paths \hat{P}_i from $Q_i^n(f^-)$ to $H_i(f^-)$ with all internal vertices in $C(X_{n+1} \cup (\bigcup \mathcal{P}'_{n+1} \cup \bigcup \mathcal{P}_{n+1}), \epsilon)$. Letting $Q_i^{n+1} = \hat{Q}_i^{n+1} \cup \hat{P}_i$, we see that each Q_i^{n+1} is a T -suitable subdivision of $S_k^{-\epsilon}$ in Γ .

Our new bounder will be $X_{n+1} := X_n \cup \bigcup \hat{\mathcal{P}}_{n+1} \cup \bigcup \mathcal{P}'_{n+1} \cup \bigcup \mathcal{P}_{n+1}$.

Finally, we apply Lemma A.4.6 to $Y := \{y_{e,i}^{n+1} : e \in \partial_\epsilon(S_{n+1}), i \leq k\}$ in $\Gamma[Y \cup C(X_{n+1}, \epsilon)]$. This gives us a family of disjoint rays

$$\hat{\mathcal{E}}_{n+1} = (\hat{E}_{e,i}^{n+1} : e \in \partial_\epsilon(S_{k+1}), i \in [k])$$

such that $\hat{E}_{e,i}^{n+1}$ has initial vertex $y_{e,i}^{n+1}$. Let us define our new extender \mathcal{E}_{n+1} given by

- $E_{e,i}^{n+1} = Q_i^n(e^-)P_{e,i}^{n+1}y_{e,i}^{n+1}\hat{E}_{e,i}^{n+1}$ if $e \in \partial_\epsilon(S_k), i \in [k]$;
- $E_{e,i}^{n+1} = H_i(e^-)P_{e,i}^{n+1}y_{e,i}^{n+1}\hat{E}_{e,i}^{n+1}$ if $e \in \partial_\epsilon(S_{k+1}) \setminus \partial(S_k), i \in [k]$.

This concludes the proof of Theorem A.1.2. □

A.7. The induction argument

We consider T as a rooted tree with root r . In Section A.6 we constructed an S -horde for any countable subtree S of T . In this section we will extend an S -horde for some specific countable subtree S to a T -horde, completing the proof of Theorem A.1.1.

Recall that for a vertex t of T and an infinite cardinal κ we say that a child t' of t is κ -embeddable if there are at least κ children t'' of t such that $T_{t'}$ is a (rooted) topological minor of $T_{t''}$ (see Definition A.3.6). By Corollary A.3.7, the number of children of t which are not κ -embeddable is less than κ .

Definition A.7.1 (κ -closure). Let T be an infinite tree with root r .

- If S is a subtree of T and S' is a subtree of S , then we say that S' is κ -closed in S if for any vertex t of S' all children of t in S are either in S' or else are κ -embeddable.
- The κ -closure of S' in S is the smallest κ -closed subtree of S including S' .

Lemma A.7.2. *Let S' be a subtree of S . If κ is an uncountable regular cardinal and S' has size less than κ , then the κ -closure of S' in S also has size less than κ .*

Proof. Let $S'(0) := S'$ and define $S'(n+1)$ inductively to consist of $S'(n)$ together with all non- κ -embeddable children contained in S for all vertices of $S'(n)$. Clearly $\bigcup_{n \in \mathbb{N}} S'(n)$ is the κ -closure of S' . If κ_n denotes the size of $S'(n)$, then $\kappa_n < \kappa$ by induction with Corollary A.3.7. Therefore, the size of the κ -closure is bounded by $\sum_{n \in \mathbb{N}} \kappa_n < \kappa$, since κ has uncountable cofinality. \square

We will construct the desired T -horde via transfinite induction on the cardinals $\mu \leq |T|$. Our first lemma illustrates the induction step for regular cardinals.

Lemma A.7.3. *Let κ be an uncountable regular cardinal. Let S be a rooted subtree of T of size at most κ and let S' be a κ -closed rooted subtree of S of size less than κ . Then any S' -horde $(H_i : i \in \mathbb{N})$ can be extended to an S -horde.*

Proof. Let $(s_\alpha : \alpha < \kappa)$ be an enumeration of the vertices of S such that the parent of any vertex appears before that vertex in the enumeration, and for any α let S_α be the subtree of T with vertex set $V(S') \cup \{s_\beta : \beta < \alpha\}$. Let \bar{S}_α denote the κ -closure of S_α in S , and observe that $|\bar{S}_\alpha| < \kappa$ by Lemma A.7.2.

We will recursively construct for each α an \bar{S}_α -horde $(H_i^\alpha : i \in \mathbb{N})$ in Γ , where each of these hordes extends all the previous ones. For $\alpha = 0$ we let $H_i^0 = H_i$ for each $i \in \mathbb{N}$. For any limit ordinal λ we have $\bar{S}_\lambda = \bigcup_{\beta < \lambda} \bar{S}_\beta$, and so we can take $H_i^\lambda = \bigcup_{\beta < \lambda} H_i^\beta$ for each $i \in \mathbb{N}$.

For any successor ordinal $\alpha = \beta + 1$, if $s_\beta \in \bar{S}_\beta$, then $\bar{S}_\alpha = \bar{S}_\beta$, and so we can take $H_i^\alpha = H_i^\beta$ for each $i \in \mathbb{N}$. Otherwise, \bar{S}_α is the κ -closure of $\bar{S}_\beta + s_\beta$, and so $\bar{S}_\alpha - \bar{S}_\beta$ is a subtree of T_{s_β} . Furthermore, since s_β is not contained in \bar{S}_β , it must be κ -embeddable.

Let s be the parent of s_β . By suitability of the H_i^β , we can find for each $i \in \mathbb{N}$ some subdivision \hat{H}_i of T_s with $\hat{H}_i(s) = H_i^\beta(s)$. We now build the H_i^α recursively in i as follows:

Let t_i be a child of s such that T_{t_i} has a rooted subdivision K of T_{s_β} , and such that $\hat{H}_i(T_{t_i} + s) - \hat{H}_i(s)$ is disjoint from all H_j^α with $j < i$ and from all H_j^β . Since there are κ disjoint possibilities for K , and all H_j^α with $j < i$ and all H_j^β cover less than κ vertices in Γ , such a choice of K is always possible. Then let H_i^α be the union of H_i^β with $\hat{H}_i(K(\bar{S}_\alpha - \bar{S}_\beta) + st_i)$.

This completes the construction of the $(H_i^\alpha : i \in \mathbb{N})$. Obviously, each H_i^α for $i \in \mathbb{N}$ is a subdivision of \bar{S}_α with $H_i^\alpha(\bar{S}_\gamma) = H_i^\gamma$ for all $\gamma < \alpha$, and all of them are pairwise disjoint for $i \neq j \in \mathbb{N}$. Moreover, H_i^α is T -suitable since for all vertices $H_i^\alpha(t)$ whose t -suitability is not witnessed in previous construction steps, their suitability is witnessed now by the corresponding subtree of \hat{H}_i . Hence $(\bigcup_{\alpha < \kappa} H_i^\alpha : i \in \mathbb{N})$ is the desired S -horde extending $(H_i : i \in \mathbb{N})$. \square

Our final lemma will deal with the induction step for singular cardinals. The crucial ingredient will be to represent a tree S of singular cardinality μ as a continuous increasing union of $< \mu$ -sized subtrees $(S_\varrho : \varrho < \text{cf}(\mu))$ where each S_ϱ is $|S_\varrho|^+$ -closed in S . This type of argument is based on Shelah's singular compactness theorem, see e.g. [97], but can be read without knowledge of the paper.

Definition A.7.4 (S -representation). For a tree S with $|S| = \mu$, we call a sequence $\mathcal{S} = (S_\varrho : \varrho < \text{cf}(\mu))$ of subtrees of S with $|S_\varrho| = \mu_\varrho$ an S -representation if

- $(\mu_\varrho : \varrho < \text{cf}(\mu))$ is a strictly increasing continuous sequence of cardinals less than μ which is cofinal for μ ,
- $S_\varrho \subseteq S_{\varrho'}$ for all $\varrho < \varrho'$, i.e. \mathcal{S} is increasing,

- for every limit $\lambda < \text{cf}(\mu)$ we have $\bigcup_{\varrho < \lambda} S_\varrho = S_\lambda$, i.e. \mathcal{S} is *continuous*,
- $\bigcup_{\varrho < \text{cf}(\mu)} S_\varrho = S$, i.e. \mathcal{S} is *exhausting*,
- S_ϱ is μ_ϱ^+ -closed in S for all $\varrho < \text{cf}(\mu)$, where μ_ϱ^+ is the successor cardinal of μ_ϱ .

Moreover, for a tree $S' \subseteq S$ we say that \mathcal{S} is an *S-representation extending S'* if additionally

- $S' \subseteq S_\varrho$ for all $\varrho < \text{cf}(\mu)$.

Lemma A.7.5. *For every tree S of singular cardinality and every subtree S' of S with $|S'| < |S|$ there is an S -representation extending S' .*

Proof. Let $|S| = \mu$ be singular, and let $|S'| = \kappa$. Let $(s_\alpha : \alpha < \mu)$ be an enumeration of the vertices of S . Let γ be the cofinality of μ and let $(\mu_\varrho : \varrho < \gamma)$ be a strictly increasing continuous cofinal sequence of cardinals less than μ with $\mu_0 > \gamma$ and $\mu_0 > \kappa$. By recursion on i we choose for each $i \in \mathbb{N}$ a sequence $(S_\varrho^i : \varrho < \gamma)$ of subtrees of S of cardinality μ_ϱ , where the vertices of each S_ϱ^i are enumerated as $(s_{\varrho,\alpha}^i : \alpha < \mu_\varrho)$, such that:

- (1) S_ϱ^i is μ_ϱ^+ -closed.
- (2) S' is a subtree of S_ϱ^i .
- (3) $S_{\varrho'}^i$ is a subtree of S_ϱ^i for $\varrho' < \varrho$.
- (4) $s_\alpha \in S_\varrho^i$ for $\alpha < \mu_\varrho$.
- (5) $s_{\varrho',\alpha}^j \in S_\varrho^i$ for any $j < i$, $\varrho \leq \varrho' < \gamma$ and $\alpha < \mu_\varrho$

This is achieved by recursion on ϱ as follows: For any given $\varrho < \gamma$, let X_ϱ^i be the set of all vertices which are forced to lie in S_ϱ^i by conditions (2)–(5), that is, all vertices of S' or of $S_{\varrho'}^i$ with $\varrho' < \varrho$, all s_β with $\beta < \mu_\varrho$ and all $s_{\varrho',\alpha}^j$ with $j < i$, $\varrho \leq \varrho' < \gamma$ and $\alpha < \mu_\varrho$. Then X_ϱ^i has cardinality μ_ϱ and so it is included in a subtree of S of cardinality μ_ϱ . We take S_ϱ^i to be the μ_ϱ^+ -closure of this subtree in S . Note that, since μ_ϱ^+ is regular, it follows from Lemma A.7.2 that S_ϱ^i has cardinality μ_ϱ .

For each $\varrho < \gamma$, let $S_\varrho := \bigcup_{i \in \mathbb{N}} S_\varrho^i$. Then each S_ϱ is a union of μ_ϱ^+ -closed trees and so is μ_ϱ^+ -closed itself. Furthermore, each S_ϱ clearly has cardinality μ_ϱ .

It follows from (4) that $S = \bigcup_{\varrho < \gamma} S_\varrho$. Thus, it remains to argue that our sequence is indeed continuous, i.e. that for any limit ordinal $\lambda < \gamma$ we have $S_\lambda = \bigcup_{\varrho < \lambda} S_\varrho$. The inclusion $\bigcup_{\varrho < \lambda} S_\varrho \subseteq S_\lambda$ is clear from (3). For the other inclusion, let s be any element of S_λ . Then there is some $i \in \mathbb{N}$ with $s \in S_\lambda^i$ and so there is some $\alpha < \mu_\alpha$ with $s = s_{\lambda, \alpha}^i$. Then by continuity there is some $\sigma < \lambda$ with $\alpha < \mu_\sigma$ and so $s \in S_\sigma^{i+1} \subseteq S_\sigma \subseteq \bigcup_{\varrho < \lambda} S_\varrho$. \square

Lemma A.7.6. *Let μ be a cardinal. Then for any rooted subtree S of T of size μ and any uncountable regular cardinal $\kappa \leq \mu$, any S' -horde $(H_i : i \in \mathbb{N})$ of a κ -closed rooted subtree S' of S of size less than κ can be extended to an S -horde.*

Proof. The proof is by transfinite induction on μ . If μ is regular, we let S'' be the μ -closure of S' in S . Thus S'' has size less than μ . So by the induction hypothesis $(H_i : i \in \mathbb{N})$ can be extended to an S'' -horde, which by Lemma A.7.3 can be further extended to an S -horde.

So let us assume that μ is singular, and write $\gamma = \text{cf}(\mu)$. By Lemma A.7.5, fix an S -representation $\mathcal{S} = (S_\varrho : \varrho < \text{cf}(\mu))$ extending S' with $|S'| < |S_0|$.

We now recursively construct for each $\varrho < \gamma$ an S_ϱ -horde $(H_i^\varrho : i \in \mathbb{N})$, where each of these hordes extends all the previous ones and $(H_i : i \in \mathbb{N})$. Using that each S_ϱ is μ_ϱ^+ -closed in S , we can find $(H_i^0 : i \in \mathbb{N})$ by the induction hypothesis, and if ϱ is a successor ordinal we can find $(H_i^\varrho : i \in \mathbb{N})$ by again using the induction hypothesis. For any limit ordinal λ we set $H_i^\lambda = \bigcup_{\varrho < \lambda} H_i^\varrho$ for each $i \in \mathbb{N}$, which yields an S_λ -horde by the continuity of \mathcal{S} .

This completes the construction of the H_i^ϱ . Then $(\bigcup_{\varrho < \gamma} H_i^\varrho : i \in \mathbb{N})$ is an S -horde extending $(H_i : i \in \mathbb{N})$. \square

Finally, with the right induction start we obtain the following theorem and hence a proof of Theorem A.1.1.

Theorem A.7.7. *Let T be a tree and Γ a graph such that $nT \leq \Gamma$ for every $n \in \mathbb{N}$. Then there is a T -horde, and hence $\aleph_0 T \leq \Gamma$.*

Proof. By Theorem A.1.2, we may assume that T is uncountable. Let S' be the \aleph_1 -closure of the root $\{r\}$ in T . Then S' is countable by Lemma A.7.2 and so

there is an S' -horde in Γ by Theorem A.1.2. This can be extended to a T -horde in Γ by Lemma A.7.6 with $\mu = |T|$. \square

B. Ubiquity of graphs with nowhere-linear end structure

B.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 denotes the disjoint union of α many copies of G . For example, a classic result of Halin [49] says that the ray is \subseteq -ubiquitous, where \subseteq is the subgraph relation.

Examples of graphs which are not ubiquitous with respect to the subgraph or topological minor relation are known (see [2] for some particularly simple examples). In [3] 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 relied on the existence of a counterexample to the well-quasi-ordering of infinite graphs under the minor relation, for which only examples of size at least the continuum are known [103]. In particular, the question of whether there exists a countable graph which is not \preceq -ubiquitous remains open. Most importantly, however, Andreae [3] 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 [2] Andreae proved that his conjecture holds for a large class of locally finite graphs. The exact definition of this class is technical, but in particular his result implies the following.

Theorem B.1.1 (Andreae, [2, Corollary 2]). *Let G be a connected, locally finite graph of finite tree-width such that every block of G is finite. Then G is \preceq -ubiquitous.*

*A precise definitions of rays, the ends of a graph, their degree, and what it means for a ray to converge to an end can be found in Section B.2.

Note that every end in such a graph must have degree* one.

Andreae’s proof employs deep results about well-quasi-orderings of labelled (infinite) trees [71]. Interestingly, the way these tools are used does not require the extra condition in Theorem B.1.1 that every block of G is finite and so it is natural to ask if his proof can be adapted to remove this condition. And indeed, it is the purpose of the present and subsequent paper [15], to show that this is possible, i.e. that all connected, locally finite graphs of finite tree-width are \preceq -ubiquitous.




Figure B.1.: A linkage between \mathcal{R} and \mathcal{S} .

The present paper lays the groundwork for this extension of Andreae’s result. The fundamental obstacle one encounters when trying to extend Andreae’s methods is the following: In the proof we often have two families of disjoint rays $\mathcal{R} = (R_i : i \in I)$ and $\mathcal{S} = (S_j : j \in J)$ in Γ , which we may assume all converge* to a common end of Γ , and we wish to find a *linkage* between \mathcal{R} and \mathcal{S} , that is, an injective function $\sigma : I \rightarrow J$ and a set \mathcal{P} of disjoint finite paths P_i from $x_i \in R_i$ to $y_{\sigma(i)} \in S_{\sigma(i)}$ such that the walks

$$\mathcal{T} = (R_i x_i P_i y_{\sigma(i)} S_{\sigma(i)} : i \in I)$$

formed by following each R_i along to x_i , then following the path P_i to $y_{\sigma(i)}$, then following the tail of $S_{\sigma(i)}$, form a family of disjoint rays (see Figure B.1). Broadly, we can think of this as ‘re-routing’ the rays \mathcal{R} to some subset of the rays in \mathcal{S} . Since all the rays in \mathcal{R} and \mathcal{S} converge to the same end of Γ , it is relatively simple

to show that, as long as $|I| \leq |J|$, there is enough connectivity between the rays in Γ to ensure that such a linkage always exists.

However, in practice it is not enough for us to be guaranteed the existence of some injection σ giving rise to a linkage, but instead we want to choose σ in advance, and be able to find a corresponding linkage afterwards.

In general, however, it is possible that for certain choices of σ no suitable linkage exists. Consider, for example, the case where Γ is the *half-grid* (which we denote by $\mathbb{Z} \square \mathbb{N}$), which is the graph whose vertex set is $\mathbb{Z} \times \mathbb{N}$ and where two vertices are adjacent if they differ in precisely one co-ordinate and the difference in that co-ordinate is one. If we consider two sufficiently large families of disjoint rays \mathcal{R} and \mathcal{S} in Γ , then it is not hard to see that both \mathcal{R} and \mathcal{S} inherit a linear ordering from the planar structure of Γ , which must be preserved by any linkage between them.

By analysing the possible kind of linkages which can arise between two families of rays converging to a given end, we will give a classification of ends of infinite degree, which we call *thick*, into three different types depending on the possible linkages they support. Roughly all such ends will either be *pebbly*, meaning that we can always find suitable linkages for all σ as above, *half-grid-like*, and exhibit behaviour similar to that of the half-grid $\mathbb{Z} \square \mathbb{N}$, or *grid-like*, and exhibit behaviour similar to that of the full-grid $\mathbb{Z} \square \mathbb{Z}$ (which is analogously defined as the half-grid but with $\mathbb{Z} \times \mathbb{Z}$ as vertex set). We will give precise definitions of these terms in Sections [B.5](#) and [B.7](#).

Theorem B.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.*

If appropriate ends of Γ are pebbly, then this freedom in choosing our linkages would allow us to follow Andreae's proof strategy in order to prove the ubiquity of G . However, in fact the property of an end being pebbly is so strong that we do not need to follow Andreae's strategy for such graphs. More precisely, in a pebbly end we can use the existence of such linkages to directly build a K_{\aleph_0} -minor of Γ (See Lemma [B.5.2](#)), from which it follows that $\aleph_0 G \preceq \Gamma$ for any countable graph G . In this way, Theorem [B.1.2](#) can be thought of as a local structure theorem for the ends of a graph which don't contain a K_{\aleph_0} -minor.

In this way, Theorem [B.1.2](#) allows us to make some structural assumptions

on the ‘host’ graph Γ when considering the question of \preceq -ubiquity. However, more importantly, it also allows us to make some structural assumptions about G . Roughly, if the ends of G do not have a particularly simple structure then the fact that $nG \preceq \Gamma$ for each $n \in \mathbb{N}$ will imply that Γ must have a pebbly end.

Analysing this situation gives rise to the following definition: We say that an end ϵ of a graph G is *linear* if for every finite set \mathcal{R} of at least three disjoint rays in G which converge to ϵ we can order the elements of \mathcal{R} as $\mathcal{R} = \{R_1, R_2, \dots, R_n\}$ such that for each $1 \leq k < i < \ell \leq n$, the rays R_k and R_ℓ belong to different ends of $G - V(R_i)$.

For example, the half-grid has a unique end and it is linear. On the other end of the spectrum, let us say that a graph G has *nowhere-linear end structure* if no end of G is linear.

Our main theorem in this paper is the following.

Theorem B.1.3. *Every locally finite connected graph with nowhere-linear end structure is \preceq -ubiquitous.*

More generally, these ideas will allow us to assume, when following the proof strategy of Andreae, that certain ends of Γ are grid-like or half-grid-like, and that certain ends of G are linear. The fact that G is linear will mean that the only functions σ that we have to consider are ones which preserve the linear ordering on the rays, and the fact that Γ is grid- or half-grid-like will allow us to deduce that appropriate linkages exist for such functions. This will be a key part of our extension of Theorem B.1.1 in [15].

However, independently of these potential later developments, our methods already allow us to establish new ubiquity results for many natural graphs and graph classes.

As a first concrete example, consider the full-grid $G = \mathbb{Z} \square \mathbb{Z}$. G is one-ended, and for any ray R in G , the graph $G - V(R)$ still has at most one end. Hence the unique end of G is non-linear, and so Theorem B.1.3 has the following corollary:

Corollary B.1.4. *The full-grid is \preceq -ubiquitous.*

Using an argument similar in spirit to that of Halin [46], we also establish the following theorem in this paper:

Theorem B.1.5. *Any connected minor of the half-grid $\mathbb{N} \square \mathbb{Z}$ is \preceq -ubiquitous.*

Since every countable tree is a minor of the half-grid, Theorem B.1.5 implies that all countable trees are \preceq -ubiquitous, see Corollary B.9.4. We remark that while it has been shown that all trees are ubiquitous with respect to the topological minor relation, [13], the question of whether all uncountable trees are \preceq -ubiquitous has remains open, and we hope to resolve this in a paper in preparation.

In a different direction, if G is any locally finite connected graph, then it is possible to show that $G \square \mathbb{Z}$ or $G \square \mathbb{N}$ either have nowhere-linear end structure, or are either the full-grid or a subgraph of the half-grid. Hence, Theorems B.1.3 and B.1.5 and Corollary B.1.4 have the following corollary.

Theorem B.1.6. *For every locally finite connected graph G , both $G \square \mathbb{Z}$ and $G \square \mathbb{N}$ are \preceq -ubiquitous.*

Finally, we will also show the following result about non-locally finite graphs. For $k \in \mathbb{N}$, we let the k -fold dominated ray be the graph DR_k formed by taking a ray together with k additional vertices, each of which we make adjacent to every vertex in the ray. For $k \leq 2$, DR_k is a minor of the half-grid, and so ubiquitous by Theorem B.1.5. In our last theorem, we show that DR_k is ubiquitous for all $k \in \mathbb{N}$.

Theorem B.1.7. *The k -fold dominated ray DR_k is \preceq -ubiquitous for every $k \in \mathbb{N}$.*

The paper is structured as follows: In Section B.2 we introduce some basic terminology for talking about minors. In Section B.3 we introduce the concept of a *ray graph* and *linkages* between families of rays, which will help us to describe the structure of an end. In Sections B.4 and B.5 we introduce a pebble-pushing game which encodes possible linkages between families of rays and use this to give a sufficient condition for an end to contain a countable clique minor. In Sections B.6 and B.7 we prove Theorem B.1.2, classifying the thick ends which are non-pebbly. In Section B.8 we re-introduce some concepts from [13] and show that we may assume that the G -minors in Γ are *concentrated* towards some end ϵ of Γ . In Section B.9 we use the results of the previous section to prove Theorem B.1.5 and finally in Section B.10 we prove Theorem B.1.3 and its corollaries.

B.2. Preliminaries

In our graph theoretic notation we generally follow the textbook of Diestel [24]. Given two graphs G and H the *cartesian product* $G \square H$ is a graph with vertex

set $V(G) \times V(H)$ with an edge between (a, b) and (c, d) if and only if $a = c$ and $bd \in E(H)$ or $ac \in E(G)$ and $b = d$.

Definition B.2.1. A one-way infinite path is called a *ray* and a two-way infinite path is called a *double ray*.

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 . 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 .

Given two paths or rays P and Q which are disjoint but for one of their endvertices, 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 B.2.2 (Ends of a graph, cf. [24, Chapter 8]). An *end* of an infinite graph Γ is an equivalence class of rays, where two rays R and S 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 $C(X, \epsilon)$.

For an end $\epsilon \in \Omega(\Gamma)$ we define the *degree* of ϵ in Γ 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 [49, Satz 1]. If an end has finite degree, we call it *thin*. Otherwise, we call it *thick*.

A vertex $v \in V(\Gamma)$ *dominates* an end $\epsilon \in \Omega(\Gamma)$ if there is a ray $R \in \epsilon$ such that there are infinitely many $v-R$ -paths in Γ that are vertex disjoint apart from v .

We will use the following two basic facts about infinite graphs.

Proposition B.2.3. [24, Proposition 8.2.1] *An infinite connected graph contains either a ray or a vertex of infinite degree.*

Proposition B.2.4. [24] *Exercise 8.19 A graph G contains a subdivided K_{\aleph_0} as a subgraph if and only if G has an end which is dominated by infinitely many vertices.*

Definition B.2.5 (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 B.2.6 (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 .

Definition B.2.7 (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 $v \in V(G) \cap V(G')$.

Note 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 $H(w)$.

Definition B.2.8 (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 B.2.9 (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 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 $\varphi^{-1}(v)$ for each $(u, v) \in E(M)$.

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 B.2.10 (Pullback). Let G be a graph, $R \subseteq G$ a ray, and let H be a tidy IG . The *pullback of R to H* is the subgraph $H^\downarrow(R) \subseteq H$ where $H^\downarrow(R)$ is subgraph minimal such that $(H^\downarrow(R), \varphi \upharpoonright V(H^\downarrow(R)))$ is an IM .

Note that, since H is tidy, $H^\downarrow(R)$ is well defined. As we shall see, $H^\downarrow(R)$ will be a ray.

Lemma B.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.*

Proof. Let $R = x_1x_2\dots$. For each integer $i \geq 1$ there is a unique edge $v_iw_i \in E(H)$ between the branch sets $H(x_i)$ and $H(x_{i+1})$. By the tidiness assumption, $H(x_{i+1})$ induces a tree in H , and so there is a unique path $P_i \subset H(x_{i+1})$ from w_i to v_{i+1} in H .

By minimality of $H^\downarrow(R)$, it follows that $H^\downarrow(R)(x_1) = \{v_1\}$ and $H^\downarrow(R)(x_{i+1}) = V(P_i)$ for each $i \geq 1$. Hence $H^\downarrow(R)$ is a ray. \square

B.3. The Ray Graph

Definition B.3.1 (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 in Γ 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 $\text{RG}_\Gamma(\mathcal{R})$.

The following lemmas are simple exercises. For a family \mathcal{R} of disjoint rays in G tending to the same end and $H \subseteq \Gamma$ being an IG the aim is to establish the following: if \mathcal{S} is a family of disjoint rays in Γ which contains the pullback $H^\downarrow(R)$ of each $R \in \mathcal{R}$, then the subgraph of the ray graph $\text{RG}_\Gamma(\mathcal{S})$ induced on the vertices given by $\{H^\downarrow(R) : R \in \mathcal{R}\}$ is connected.

Lemma B.3.2. *Let G be a graph and let $\mathcal{R} = (R_i : i \in I)$ be a finite family of disjoint rays in G . Then $\text{RG}_G(\mathcal{R})$ is connected if and only if all rays in \mathcal{R} tend to*

a common end $\omega \in \Omega(G)$. Moreover, if R'_i is a tail of R_i for each $i \in I$, then we have that $\text{RG}(R_i: i \in I) = \text{RG}(R'_i: i \in I)$.

Lemma B.3.3. *Let G be a graph, $\mathcal{R} = (R_i: i \in I)$ be a finite family of disjoint rays in G and let H be an IG. If $\mathcal{R}' = (H^\downarrow(R_i): i \in I)$ is the set of pullbacks of the rays in \mathcal{R} in H , then $\text{RG}_G(\mathcal{R}) = \text{RG}_H(\mathcal{R}')$.*

Lemma B.3.4. *Let G be a graph, $H \subseteq G$, $\mathcal{R} = (R_i: i \in I)$ be a finite disjoint family of rays in H and let $\mathcal{S} = (S_j: j \in J)$ be a finite disjoint family of rays in $G - V(H)$, where I and J are disjoint. Then $\text{RG}_H(\mathcal{R})$ is a subgraph of $\text{RG}_G(\mathcal{R} \cup \mathcal{S})[I]$. In particular, if all rays in \mathcal{R} tend to a common end in H , then $\text{RG}_G(\mathcal{R} \cup \mathcal{S})[I]$ is connected.*

Recall that an end ω of a graph G is called *linear* if for every finite set \mathcal{R} of at least three disjoint ω -rays in G we can order the elements of \mathcal{R} as $\mathcal{R} = \{R_1, R_2, \dots, R_n\}$ such that for each $1 \leq k < i < \ell \leq n$, the rays R_k and R_ℓ belong to different ends of $G - V(R_i)$.

Lemma B.3.5. *An end ω of a graph G is linear if and only if the ray graph of every finite family of disjoint ω -rays is a path.*

Proof. For the forward direction suppose ω is linear and $\{R_1, R_2, \dots, R_n\}$ converge to ω , with the order given by the definition of linear. It follows that there is no $1 \leq k < i < \ell \leq n$ such that $k\ell$ is an edge in $\text{RG}(R_j: j \in [n])$. However, by Lemma B.3.2 $\text{RG}(R_j: j \in [n])$ is connected, and hence it must be the path $12 \dots n$.

Conversely, suppose that the ray graph of every finite family of ω -rays is a path. Then, every such family \mathcal{R} can be ordered as $\{R_1, R_2, \dots, R_n\}$ such that $\text{RG}(\mathcal{R})$ is the path $12 \dots n$. In particular, for each i , we have that $k\ell \notin E(\text{RG}(\mathcal{R}))$ whenever $1 \leq k < i < \ell \leq n - 1$.

Suppose for a contradiction that there exists $1 \leq k < i < \ell \leq n - 1$ such that R_k and R_ℓ belong to the same end of $G - V(R_i)$, and so there is an infinite family of vertex disjoint paths \mathcal{P} from R_k to R_ℓ in $G - V(R_i)$. Each of these paths must contain a subpath which goes from a ray R_r for some $1 \leq r < i$ to a ray R_s for some $i < s \leq n - 1$, and which meets no other ray in \mathcal{R} . Since there are infinitely many paths, by the pigeon hole principle there is some $1 \leq r < i < s \leq n - 1$ such that there are infinitely many vertex disjoint paths from R_r to R_s in $G \setminus V(R_i)$ which meet not other ray in \mathcal{R} , and so $rs \in E(\text{RG}(\mathcal{R}))$, a contradiction. \square

We will also use the following lemma, whose proof is an easy exercise.

Lemma B.3.6. *Let $\mathcal{R} = (R_i: i \in I)$ be a finite family of disjoint rays in G and let $\mathcal{R}' = (R_i: i \in J)$ be a subfamily of \mathcal{R} . Then $\text{RG}(\mathcal{R}')$ contains an edge between $i \in J$ and $j \in J$ if and only if i and j lie in the same component of $\text{RG}(\mathcal{R}) - (J \setminus \{i, j\})$.*

Definition B.3.7 (Tail of a ray after a set). Given a ray R in a graph G and a finite set $X \subseteq V(G)$ the *tail of R after X* , denoted by $T(R, X)$, is the unique infinite component of R in $G - X$.

Definition B.3.8 (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 of G , where the initial vertex of each R_i is denoted x_i . A family $\mathcal{P} = (P_i: i \in I)$ of paths in G 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.

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 . Given a vertex set $X \subseteq V(G)$ 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} . We say that a function $\sigma: I \rightarrow J$ is a *transition function* from \mathcal{R} to \mathcal{S} if for any finite vertex set $X \subseteq V(G)$ there is a linkage from \mathcal{R} to \mathcal{S} after X that induces σ .

We will need the following lemma from [13], which asserts the existence of linkages.

Lemma B.3.9 (Weak linking lemma [13, Lemma 4.3]). *Let G be a graph, $\omega \in \Omega(G)$ and let $n \in \mathbb{N}$. Then for any two families $\mathcal{R} = (R_i: i \in [n])$ and $\mathcal{S} = (S_j: j \in [n])$ of vertex disjoint ω -rays and any finite vertex set $X \subseteq V(G)$, there is a linkage from \mathcal{R} to \mathcal{S} after X .*

B.4. A pebble-pushing game

Suppose we have a family of disjoint rays $\mathcal{R} = (R_i: i \in I)$ in a graph G and a subset $J \subseteq I$. Often we will be interested in which functions we can obtain as

transition functions between $(R_i: i \in J)$ and $(R_i: i \in I)$. We can think of this as trying to ‘re-route’ the rays $(R_i: i \in J)$ to the tails of a different set of $|J|$ rays in $(R_i: i \in I)$.

To this end, it will be useful to understand the following pebble-pushing game on a graph.

Definition B.4.1 (Pebble-pushing game). Let $G = (V, E)$ be a finite graph. For any fixed positive integer k we call a tuple $(x_1, x_2, \dots, x_k) \in V^k$ a *game state* if $x_i \neq x_j$ for all $i, j \in [k]$ with $i \neq j$.

The *pebble-pushing game (on G)* is a game played by a single player. Given a game state $Y = (y_1, y_2, \dots, y_k)$, we imagine k labelled pebbles placed on the vertices (y_1, y_2, \dots, y_k) . We move between game states by moving a pebble from a vertex to an adjacent vertex which does not contain a pebble, or formally, a Y -move is a game state $Z = (z_1, z_2, \dots, z_k)$ such that there is an $\ell \in [k]$ such that $y_\ell z_\ell \in E$ and $y_i = z_i$ for all $i \in [k] \setminus \{\ell\}$.

Let $X = (x_1, x_2, \dots, x_k)$ be a game state. The X -pebble-pushing game (on G) is a pebble-pushing game where we start with k labelled pebbles placed on the vertices (x_1, x_2, \dots, x_k) .

We say a game state Y is *achievable* in the X -pebble-pushing game if there is a sequence $(X_i: i \in [n])$ of game states for some $n \in \mathbb{N}$ such that $X_1 = X$, $X_n = Y$ and X_{i+1} is an X_i -move for all $i \in [n-1]$, that is, if it is a sequence of moves that pushes the pebbles from X to Y .

A graph G is k -pebble-win if Y is an achievable game state in the X -pebble-pushing game on G for every two game states X and Y .

The following lemma shows that achievable game states on the ray graph $\text{RG}(\mathcal{R})$ yield transition functions from a subset of \mathcal{R} to itself. Therefore, it will be useful to understand which game states are achievable, and in particular the structure of graphs on which there are unachievable game states.

Lemma B.4.2. *Let Γ be a graph, $\omega \in \Omega(\Gamma)$, $m \geq k$ be positive integers and let $(S_j: j \in [m])$ be a family of disjoint rays in ω . For every achievable game state $Z = (z_1, z_2, \dots, z_k)$ in the $(1, 2, \dots, k)$ -pebble-pushing game on $\text{RG}(S_j: j \in [m])$, the map σ defined via $\sigma(i) := z_i$ for every $i \in [k]$ is a transition function from $(S_i: i \in [k])$ to $(S_j: j \in [m])$.*

Proof. Note first that if σ is a transition function from $(S_i: i \in [k])$ to $(S_j: j \in [m])$ and τ is a transition function from $(S_i: i \in \sigma([k]))$ to $(S_j: j \in [m])$, then clearly $\tau \circ \sigma$ is a transition function from $(S_i: i \in [k])$ to $(S_j: j \in [m])$.

Hence, it is sufficient to show the statement holds when σ is obtained from $(1, 2, \dots, k)$ by a single move, that is, there is some $t \in [k]$ and a vertex $\sigma(t) \notin [k]$ such that $\sigma(t)$ is adjacent to t in $\text{RG}(S_j: j \in [m])$ and $\sigma(i) = i$ for $i \in [k] \setminus \{t\}$.

So, let $X \subseteq V(G)$ be a finite set. We will show that there is a linkage from $(S_i: i \in [k])$ to $(S_j: j \in [m])$ after X that induces σ . By assumption, there is an edge $t\sigma(t)$ of $\text{RG}(S_j: j \in [m])$. Hence, there is a path P between $T(S_t, X)$ and $T(S_{\sigma(t)}, X)$ which avoids X and all other S_j .

Then the family $\mathcal{P} = (P_1, P_2, \dots, P_k)$ where $P_t = P$ and $P_i = \emptyset$ for each $i \neq t$ is a linkage from $(S_i: i \in [k])$ to $(S_j: j \in [m])$ after X that induces σ . \square

We note that this pebble-pushing game is sometimes known in the literature as “permutation pebble motion” [66] or “token reconfiguration” [19]. Previous results have mostly focused on computational questions about the game, rather than the structural questions we are interested in, but we note that in [66] the authors give an algorithm that decides whether or not a graph is k -pebble-win, from which it should be possible to deduce the main result in this section, Lemma B.4.9. However, since a direct derivation was shorter and self contained, we will not use their results. We present the following simple lemmas without proof.

Lemma B.4.3. *Let G be a finite graph and X a game state.*

- *If Y is an achievable game state in the X -pebble-pushing game on G , then X is an achievable game state in the Y -pebble-pushing game on G .*
- *If Y is an achievable game state in the X -pebble-pushing game on G and Z is an achievable game state in the Y -pebble-pushing game on G , then Z is an achievable game state in the X -pebble-pushing game on G .*

Definition B.4.4. Let G be a finite graph and let $X = (x_1, x_2, \dots, x_k)$ be a game state. Given a permutation σ of $[k]$ let us write $X^\sigma = (x_{\sigma(1)}, x_{\sigma(2)}, \dots, x_{\sigma(k)})$. We define the *pebble-permutation group* of (G, X) to be the set of permutations σ of $[k]$ such that X^σ is an achievable game state in the X -pebble-pushing game on G .

Note that by Lemma B.4.3, the pebble-permutation group of (G, X) is a subgroup of the symmetric group S_k .

Lemma B.4.5. *Let G be a graph and let X be a game state. If Y is an achievable game state in the X -pebble-pushing game and σ is in the pebble-permutation group of Y , then σ is in the pebble-permutation group of X .*

Lemma B.4.6. *Let G be a finite connected graph and let X be a game state. Then G is k -pebble-win if and only if the pebble-permutation group of (G, X) is S_k .*

Proof. Clearly, if the pebble-permutation group is not S_k then G is not k -pebble-win. Conversely, since G is connected, for any game states X and Y there is some τ such that Y^τ is an achievable game state in the X -pebble-pushing game, since we can move the pebbles to any set of k vertices, up to some permutation of the labels. We know by assumption that $X^{\tau^{-1}}$ is an achievable game state in the X -pebble-pushing game. Therefore, by Lemma B.4.3, Y is an achievable game state in the X -pebble-pushing game. \square

Lemma B.4.7. *Let G be a finite connected graph and let $X = (x_1, x_2, \dots, x_k)$ be a game state. If G is not k -pebble-win, then there is a two colouring $c: X \rightarrow \{r, b\}$ such that both colour classes are non trivial and for all $i, j \in [k]$ with $c(x_i) = r$ and $c(x_j) = b$ the transposition (ij) is not in the pebble-permutation group.*

Proof. Let us draw a graph H on $\{x_1, x_2, \dots, x_k\}$ by letting $x_i x_j$ be an edge if and only if (ij) is in the pebble-permutation group of (G, X) . It is a simple exercise to show that the pebble-permutation group of (G, X) is S_k if and only if H has a single component.

Since G is not k -pebble-win, we know by Lemma B.4.6 that there are at least two components in H . Let us pick one component C_1 and set $c(x) = r$ for all $x \in V(C_1)$ and $c(x) = b$ for all $x \in X \setminus V(C_1)$. \square

Definition B.4.8. Given a graph G , a path $x_1 x_2 \dots x_n$ in G is a *bare path* if $d_G(x_i) = 2$ for all $2 \leq i \leq n - 1$.

Lemma B.4.9. *Let G be a finite connected graph with vertex set $V := V(G)$ which is not k -pebble-win and with $|V| \geq k + 2$. Then there is a bare path $P = p_1 p_2 \dots p_n$ in G such that $|V \setminus V(P)| \leq k$. Furthermore, either every edge in P is a bridge in G , or G is a cycle.*

Proof. Let $X = (x_1, x_2, \dots, x_k)$ be a game state. By Lemma B.4.7, since G is not k -pebble-win, there is a two colouring $c: \{x_i: i \in [k]\} \rightarrow \{r, b\}$ such that both colour classes are non trivial and for all $i, j \in [k]$ with $c(x_i) = r$ and $c(x_j) = b$ the transposition (ij) is not in the pebble permutation group. Let us consider this as a three colouring $c: V \rightarrow \{r, b, 0\}$ where $c(v) = 0$ if $v \notin \{x_1, x_2, \dots, x_k\}$.

For every achievable game state $Z = (z_1, z_2, \dots, z_k)$ in the X -pebble-pushing game we define a three colouring c_Z given by $c_Z(z_i) = c(x_i)$ for all $i \in [k]$ and by $c_Z(v) = 0$ for all $v \notin \{z_1, z_2, \dots, z_k\}$. We note that, for any achievable game state Z there is no $z_i \in c_Z^{-1}(r)$ and $z_j \in c_Z^{-1}(b)$ such that (ij) is in the pebble permutation group of (G, Z) . Indeed, if it were, then by Lemma B.4.3 $X^{(ij)}$ is an achievable game state in the X -pebble-pushing game, contradicting the fact that $c(x_i) = r$ and $c(x_j) = b$.

Since G is connected, for every achievable game state Z there is a path $P = p_1 p_2 \dots p_m$ in G with $c_Z(p_1) = r$, $c_Z(p_m) = b$ and $c_Z(p_i) = 0$ otherwise. Let us consider an achievable game state Z for which G contains such a path P of maximal length.

We first claim that there is no $v \notin P$ with $c_Z(v) = 0$. Indeed, suppose there is such a vertex v . Since G is connected there is some v - P path Q in G and so, by pushing pebbles towards v on Q , we can achieve a game state Z' such that $c_{Z'} = c_Z$ on P and there is a vertex v' adjacent to P such that $c_{Z'}(v') = 0$. Clearly v' cannot be adjacent to p_1 or p_m , since then we can push the pebble on p_1 or p_m onto v' and achieve a game state Z'' for which G contains a longer path than P with the required colouring. However, if v' is adjacent to p_ℓ with $2 \leq \ell \leq m - 1$, then we can push the pebble on p_1 onto p_ℓ and then onto v' , then push the pebble from p_m onto p_1 and finally push the pebble on v' onto p_ℓ and then onto p_m .

If $Z' = (z'_1, z'_2, \dots, z'_k)$ with $p_1 = z'_i$ and $p_m = z'_j$, then above shows that (ij) is in the pebble-permutation group of (G, Z') . However, we have $c_{Z'}(z'_i) = c_Z(p_1) = r$ as well as $c_{Z'}(z'_j) = c_Z(p_m) = b$, contradicting our assumptions on $c_{Z'}$.

Next, we claim that each p_i with $3 \leq i \leq m - 2$ has degree 2. Indeed, suppose first that p_i with $3 \leq i \leq m - 2$ is adjacent to some other p_j with $1 \leq j \leq m$ such that p_i and p_j are not adjacent in P . Then it is easy to find a sequence of moves which exchanges the pebbles on p_1 and p_m , contradicting our assumptions on c_Z .

Suppose then that p_i is adjacent to a vertex v not in P . Then, $c_Z(v) \neq 0$, say without loss of generality $c_Z(v) = r$. However then, we can push the pebble on p_m

onto p_{i-1} , push the pebble on v onto p_i and then onto p_m and finally push the pebble on p_{i-1} onto p_i and then onto v . As before, this contradicts our assumptions on c_Z .

Hence $P' = p_2 p_3 \dots p_{m-1}$ is a bare path in G , and since every vertex in $V - V(P')$ is coloured using r or using b , there are at most k such vertices.

Finally, suppose that there is some edge in P' which is not a bridge of G , and so no edge of P' is a bridge of G . Before we show that G is a cycle, we make the following claim:

Claim B.4.10. *There is no achievable game state $W = (w_1, w_2, \dots, w_k)$ such that there is a cycle $C = c_1 c_2 \dots c_r c_1$ and a vertex $v \notin C$ such that:*

- *There exist distinct positive integers i, j, s and t such that $c_W(c_i) = r$, $c_W(c_j) = b$ and $c_W(c_s) = c_W(c_t) = 0$;*
- *v adjacent to some $c_v \in C$.*

Proof of Claim B.4.10. Suppose for a contradiction there exists such an achievable game state W . Since C is a cycle, we may assume without loss of generality that $c_i = c_1$, $c_s = c_2 = c_v$, $c_t = c_3$ and $c_j = c_4$. If $c_W(v) = b$, then we can push the pebble at v to c_2 and then to c_3 , push the pebble at c_1 to c_2 and then to v , and then push the pebble at c_3 to c_1 . This contradicts our assumptions on c_W . The case where $c_W(v) = r$ is similar. Finally, if $c_W(v) = 0$, then we can push the pebble at c_1 to c_2 and then to v , then push the pebble at c_4 to c_1 , then push the pebble at v to c_2 and then to c_4 . Again this contradicts our assumptions on c_W . \square

Since no edge of P' is a bridge, it follows that G contains a cycle C containing P' . If G is not a cycle, then there is a vertex $v \in V \setminus C$ which is adjacent to C . However by pushing the pebble on p_1 onto p_2 and the pebble on p_m onto p_{m-1} , which is possible since $|V| \geq k + 2$, we achieve a game state Z' such that C and v satisfy the assumptions of the above claim, a contradiction. \square

B.5. Pebbly and non-pebby ends

Definition B.5.1 (Pebbly). Let Γ be a graph and ω an end of Γ . We say ω is *pebby* if for every $k \in \mathbb{N}$ there is an $n \geq k$ and a family $\mathcal{R} = (R_i : i \in [n])$ of

disjoint rays in ω such that $\text{RG}(\mathcal{R})$ is k -pebble-win. If for some k there is no such family \mathcal{R} , we say ω is *non-pebbly* and in particular *not k -pebble-win*.

Clearly an end of degree k is not k -pebble-win, since no graph on at most k vertices is k -pebble-win, and so every pebbly end is thick. However, as we shall see, pebbly ends are particularly rich in structure.

Lemma B.5.2. *Let Γ be a graph and let $\omega \in \Omega(\Gamma)$ be a pebbly end. Then $K_{\aleph_0} \preceq \Gamma$.*

Proof. By assumption, there exists a sequence $\mathcal{R}_1, \mathcal{R}_2, \dots$ of families of disjoint ω -rays such that, for each $k \in \mathbb{N}$, $\text{RG}(\mathcal{R}_k)$ is k -pebble-win. Let us suppose that

$$\mathcal{R}_i = (R_1^i, R_2^i, \dots, R_{m_i}^i) \text{ for each } i \in \mathbb{N}.$$

Let us enumerate the vertices and edges of K_{\aleph_0} with a bijection $\sigma: \mathbb{N} \cup \mathbb{N}^{(2)} \rightarrow \mathbb{N}$ such that $\sigma(i, j) > \max\{\sigma(i), \sigma(j)\}$ for every $\{i, j\} \in \mathbb{N}^{(2)}$ and also $\sigma(1) < \sigma(2) < \dots$. For each $k \in \mathbb{N}$ let G_k be the graph on vertex set $V_k = \{i \in \mathbb{N} : \sigma(i) \leq k\}$ and edge set $E_k = \{\{i, j\} \in \mathbb{N}^{(2)} : \sigma(i, j) \leq k\}$.

We will inductively construct subgraphs H_k of Γ such that H_k is an IG_k extending H_{k-1} . Furthermore for each $k \in \mathbb{N}$ if $V(G_k) = [n]$ then there will be tails T_1, T_2, \dots, T_n of n distinct rays in \mathcal{R}_n such that for every $i \in [n]$ the tail T_i meets H_k in a vertex of the branch set of i , and is otherwise disjoint from H_k . We will assume without loss of generality that T_i is a tail of R_i^n .

Since $\sigma(1) = 1$ we can take H_1 to be the initial vertex of R_1^1 . Suppose then that $V(G_{n-1}) = [r]$ and we have already constructed H_{n-1} together with appropriate tails T_i of R_i^r for each $i \in [r]$. Suppose firstly that $\sigma^{-1}(n) = r + 1 \in \mathbb{N}$.

Let $X = V(H_{n-1})$. There is a linkage from $(T_i : i \in [r])$ to $(R_1^{r+1}, R_2^{r+1}, \dots, R_r^{r+1})$ after X by Lemma B.3.9, and, after relabelling, we may assume this linkage induces the identity on $[r]$. Let us suppose the linkage consists of paths P_i from $x_i \in T_i$ to $y_i \in R_i^{r+1}$.

Since $X \cup \bigcup_i P_i \cup \bigcup_i T_i x_i$ is a finite set, there is some vertex y_{r+1} on R_{r+1}^{r+1} such that the tail $y_{r+1} R_{r+1}^{r+1}$ is disjoint from $X \cup \bigcup_i P_i \cup \bigcup_i T_i x_i$.

To form H_n we add the paths $T_i x_i \cup P_i$ to the branch set of each $i \leq r$ and set y_{r+1} as the branch set for $r + 1$. Then H_n is an IG_n extending H_{n-1} and the tails $y_j R_j^{r+1}$ are as claimed.

Suppose then that $\sigma^{-1}(n) = \{u, v\} \in \mathbb{N}^{(2)}$ with $u, v \leq r$. We have tails T_i of R_i^r for each $i \in [r]$ which are disjoint from H_{n-1} apart from their initial vertices.

Let us take tails T_j of R_j^r for each $j > r$ which are also disjoint from H_{n-1} . Since $\text{RG}(\mathcal{R}_r)$ is r -pebble-win, it follows that $\text{RG}(T_i: i \in [m_r])$ is also r -pebble-win. Furthermore, since by Lemma B.3.2 $\text{RG}(T_i: i \in [m_r])$ is connected, there is some neighbour $w \in [m_r]$ of u in $\text{RG}(T_i: i \in [m_r])$.

Let us first assume that $w \notin [r]$. Since $\text{RG}(T_i: i \in [m_r])$ is r -pebble-win, the game state $(1, 2, \dots, v-1, w, v+1, \dots, r)$ is an achievable game state in the $(1, 2, \dots, r)$ -pebble-pushing game and hence by Lemma B.4.2 the function φ_1 given by $\varphi_1(i) = i$ for all $i \in [r] \setminus \{v\}$ and $\varphi_1(v) = w$ is a transition function from $(T_i: i \in [r])$ to $(T_i: i \in [m_r])$.

Let us take a linkage from $(T_i: i \in [r])$ to $(T_i: i \in [m_r])$ inducing φ_1 which is after $V(H_{n-1})$. Let us suppose the linkage consists of paths P_i from $x_i \in T_i$ to $y_i \in T_i$ for $i \neq v$ and P_v from $x_v \in T_v$ to $y_v \in T_w$. Let

$$X = V(H_{n-1}) \cup \bigcup_{i \in [r]} P_i \cup \bigcup_{i \in [r]} T_i x_i$$

Since u is adjacent to w in $\text{RG}(T_i: i \in [m_r])$ there is a path \hat{P} between $T(T_u, X)$ and $T(T_w, X)$ which is disjoint from X and from all other T_i , say \hat{P} is from $\hat{x} \in T_u$ to $\hat{y} \in T_w$.

Finally, since $\text{RG}(T_i: i \in [m_r])$ is r -pebble-win, the game state $(1, 2, \dots, r)$ is an achievable game state in the $(1, 2, \dots, v-1, w, v+1, \dots, r)$ -pebble-pushing game and hence by Lemma B.4.2 the function φ_2 given by $\varphi_2(i) = i$ for all $i \in [r] \setminus \{v\}$ and $\varphi_2(w) = v$ is a transition function from $(T_i: i \in [r] \setminus \{v\} \cup \{w\})$ to $(T_i: i \in [m_r])$.

Let us take a further linkage from $(T_i: i \in [r] \setminus \{v\} \cup \{w\})$ to $(T_i: i \in [m_r])$ inducing φ_2 which is after $X \cup \hat{P} \cup T_u \hat{x} \cup T_w \hat{y}$. Let us suppose the linkage consists of paths P'_i from $x'_i \in T_i$ to $y'_i \in T_i$ for $i \in [r] \setminus \{v\}$ and P'_v from $x'_v \in T_w$ to $y'_v \in T_v$.

In the case that $w \in [r]$, $w < v$, say, the game state

$$(1, 2, \dots, w-1, v, w+1, \dots, v-1, w, v+1, \dots, r)$$

is an achievable game state in the $(1, 2, \dots, r)$ -pebble pushing-game and we get, by a similar argument, all $P_i, x_i, y_i, P'_i, x'_i, y'_i$ and \hat{P} .

We build H_n from H_{n-1} by adjoining the following paths:

- for each $i \neq v$ we add the path $T_i x_i P_i y_i T_i x'_i P'_i y'_i$ to H_{n-1} , adding the vertices to the branch set of i ;

- we add \hat{P} to H_{n-1} , adding the vertices of $V(\hat{P}) \setminus \{\hat{y}\}$ to the branch set of u ;
- we add the path $T_v x_v P_v y_v T_w x'_v P'_v y'_v$ to H_{n-1} , adding the vertices to the branch set of v .

We note that, since $\hat{y} \in y_v T_w x'_v$ the branch sets for u and v are now adjacent. Hence H_n is an IG_n extending H_{n-1} . Finally the rays $y'_i T_i$ for $i \in [r]$ are appropriate tails of the used rays of \mathcal{R}_r . \square

As every countable graph is a subgraph of K_{\aleph_0} , a graph with a pebbly end contains every countable graph as a minor. Thus, as $\aleph_0 G$ is countable, if G is countable, we obtain the following corollary:

Corollary B.5.3. *Let Γ be a graph with a pebbly end ω and let G be a countable graph. Then $\aleph_0 G \preceq \Gamma$.*

So, at least when considering the question of \preceq -ubiquity for countable graphs, Corollary B.5.3 allows one to restrict one's attention to host graphs Γ in which each end is non-pebbly. For this reason it will be useful to understand the structure of such ends.

On immediate observation we can make is the following corollary of Lemma B.4.9.

Corollary B.5.4. *Let ω be an end of a graph Γ which is not k -pebble-win for some positive integer k and let $\mathcal{R} = (R_i : i \in [m])$ be a family of $m \geq k + 2$ disjoint rays in ω . Then there is a bare path $P = p_1 p_2 \dots p_n$ in $\text{RG}(R_i : i \in [m])$ such that $|[m] \setminus V(P)| \leq k$. Furthermore, either each edge in P is a bridge in $\text{RG}(R_i : i \in [m])$, or $\text{RG}(R_i : i \in [m])$ is a cycle.*

So, if ω is not pebbly, then the ray graph of every family of ω -rays is either close in structure to a path, or close in structure to a cycle. In fact, this dichotomy is not just true for each ray graph individually, but rather uniformly for each ray graph in the end. That is, we will show that either every ray graph of a family of ω -rays will be close in structure to a path, or every ray graph will be close in structure to a cycle. Furthermore, the structure of this end will restrict the possible transition functions between families of ω -rays.

As motivating examples consider the half-grid $\mathbb{N} \square \mathbb{Z}$ and the full-grid $\mathbb{Z} \square \mathbb{Z}$. Both graphs have a unique end ω_h / ω_f and it is easy to show that the ray graph

of every family of ω_h -rays is a path, and the ray graph of every family of ω_f -rays is a cycle (and so in particular $\mathbb{N} \square \mathbb{Z}$ is not 2-pebble-win and $\mathbb{Z} \square \mathbb{Z}$ is not 3-pebble-win).

There is a natural way to order any family of disjoint ω_h -rays, if you imagine them drawn on a page their tails will appear in some order from left to right. Then, it can be shown that any transition function between two large enough families of ω_h -rays must preserve this ordering.

Similarly, there is a natural way to cyclically order any family of disjoint ω_f -rays. As before, it can be shown that any transition function between two large enough families of ω_f -rays must preserve this ordering.

The aim of the next few sections is to demonstrate that the above dichotomy holds for all non-pebbly ends: that either every ray-graph is close in structure to a path or close in structure to a cycle, and furthermore that in each of these cases the possible transition functions between families of rays are restricted in a similar fashion as those of the half-grid or full-grid, in which case we will say the end is *half-grid-like* or *grid-like* respectively. These results, whilst not used in this paper, will be a vital part of the proof in [15].

We note that, in principle, this trichotomy that an end of a graph is either pebbly, grid-like or half-grid-like, and the information that this implies about its finite rays graphs and the transitions between them, could in principle be derived from earlier work of Diestel and Thomas [30], who gave a structural characterisation of graphs without a K_{\aleph_0} -minor. However, to introduce their result and derive what we needed from it would have been at least as hard as our work in Section B.6, if not more complicated, and so we have opted for a straightforward and self-contained presentation.

B.6. The structure of non-pebbly ends

B.6.1. Polypods

It will be useful for our analysis of the structure of non-pebbly ends to consider the possible families of disjoint rays in the end with a fixed set of start vertices, and the relative structure of these rays.

Definition B.6.1. Given an end ϵ of a graph Γ , a *polypod* (for ϵ in Γ) is a pair (X, Y) of disjoint finite sets of vertices of Γ such that there is at least one family $(R_y: y \in Y)$ of disjoint ϵ -rays, where R_y begins at y and all the R_y are disjoint from X . Such a family $(R_y: y \in Y)$ is called a *family of tendrils* for (X, Y) . The *order* of the polypod is $|Y|$. The *connection graph* $K_{X,Y}$ of a polypod (X, Y) is a graph with vertex set Y . It has an edge between vertices v and w if and only if there is a family $(R_y: y \in Y)$ of tendrils for (X, Y) such that there is an R_v - R_w -path in Γ disjoint from X and from every other R_y .

Note that the ray graph of any family of tendrils for a polypod must be a subgraph of the connection graph of that polypod.

Definition B.6.2. We say that a polypod (X, Y) for ϵ in Γ is *tight* if its connection graph is minimal amongst connection graphs of polypods for ϵ in Γ with respect to the spanning isomorphic subgraph relation, i.e. for no other polypod (X', Y') for ϵ in Γ of order $|Y'| = |Y|$ is the graph $K_{X',Y'}$ isomorphic to a proper subgraph of $K_{X,Y}$. (Let us write $H \subsetneq G$ if H is isomorphic to a subgraph of G .) We say that a polypod *attains* its connection graph if there is some family of tendrils for that polypod whose ray graph is equal to the connection graph.

Lemma B.6.3. *Let (X, Y) be a tight polypod, $(R_y: y \in Y)$ a family of tendrils and for every $y \in Y$ let v_y be a vertex on R_y . Let X' be a finite vertex set disjoint from all $v_y R_y$ and including X as well as each of the initial segments $R_y \hat{v}_y$. Let $Y' = \{v_y: y \in Y\}$. Then (X', Y') is a tight polypod with the same connection graph as (X, Y) .*

Proof. The family $(v_y R_y: y \in Y)$ witnesses that (X', Y') is a polypod. Moreover every family of tendrils for (X', Y') can be extended by the paths $R_y v_y$ to obtain a family of tendrils for (X, Y) . Hence if there is an edge $v_y v_z$ in $K_{X',Y'}$ then there must also be the edge yz in $K_{X,Y}$. Thus $K_{X',Y'} \subsetneq K_{X,Y}$. But since (X, Y) is tight we must have equality. Therefore (X', Y') is tight as well. \square

Lemma B.6.4. *Any tight polypod (X, Y) attains its connection graph.*

Proof. We must construct a family of tendrils for (X, Y) whose ray graph is $K_{X,Y}$. We will recursively build larger and larger initial segments of the rays, together with disjoint paths between them.

Precisely this means that, after partitioning \mathbb{N} into infinite sets A_e , one for each edge e of $K_{X,Y}$, we will construct, for each $n \in \mathbb{N}$, a family $(P_y^n : y \in Y)$ of disjoint paths, and also paths Q_n such that for some arbitrary fixed ray $R \in \epsilon$:

- Each P_y^n starts at y . We write y_n for the last vertex of P_y^n .
- Each P_y^n has length at least n and there are at least n disjoint paths from P_y^n to R .
- For $m \leq n$, the path P_y^n extends P_y^m .
- If $n \in A_{vw}$ for $vw \in E(K_{X,Y})$, then Q_n is a path from P_v^n to P_w^n .
- If $n \in A_{vw}$ for $vw \in E(K_{X,Y})$, then Q_n meets no P_y^m with $y \in Y \setminus \{v, w\}$ for any $m \in \mathbb{N}$.
- All the Q_n are pairwise disjoint.
- All the P_y^n and all the Q_n are disjoint from X .
- For any $n \in \mathbb{N}$ there is a family $(R_y^n : y \in Y)$ of tendrils for (X, Y) such that each P_y^n is an initial segment of the corresponding R_y^n , and the R_y^n meet the Q_m with $m \leq n$ in $P_y^{\circ} y_n$.

Once the construction is complete, we obtain a family of tendrils by letting each R_y be the union of all the P_y^n – indeed, R_y clearly is an ϵ -ray since there are arbitrarily many disjoint paths from R_y to R . Furthermore, for any edge e of $K_{X,Y}$ the family $(Q_n : n \in A_e)$ will witness that e is in the ray graph of this family. So that ray graph will be all of $K_{X,Y}$, as required.

So it remains to show how to carry out this recursive construction. Let vw be the edge of $K_{X,Y}$ with $1 \in A_{vw}$. By the definition of the connection graph there is a family $(R_y^1 : y \in Y)$ of tendrils for (X, Y) such that there is a path Q_1 from R_v^1 to R_w^1 , disjoint from all other R_y^1 and from X .

For each $y \in Y$ let P_y^1 be an initial segment of R_y^1 with end vertex y_1 of length at least 1 such that $Q_1 \cap R_y^1 \subseteq P_y^1 y_1$, and such that there is a path from P_z^1 to R – which is possible, since both R and R_y^1 are ϵ -rays.

This choice of the P_y^1 and of Q_1 clearly satisfies the conditions above.

Suppose that we have constructed suitable P_y^m and Q_m for all $m \leq n$. For each $y \in Y$, let y_n be the endvertex of P_y^n . Let Y_n be $\{y_n : y \in Y\}$ and

$$Z_n = X \cup \bigcup_{m \leq n} \bigcup_{y \in Y} (V(P_y^m) \cup V(Q_m)).$$

Let X_n be $Z_n \setminus Y_n$, and note that every $V(Q_m) \subseteq X_n$ for every $m \leq n$. Then by Lemma B.6.3 (X_n, Y_n) is a tight polypod with the same connection graph as (X, Y) .

In particular, letting vw be the edge of $K_{X,Y}$ with $n+1 \in A_{vw}$, we have that $v_n w_n$ is an edge of K_{X_n, Y_n} . So there is a family $(S_{y_n}^{n+1} : y_n \in Y_n)$ of tendrils for (X_n, Y_n) together with a path Q_{n+1} from $S_{v_n}^{n+1}$ to $S_{w_n}^{n+1}$ disjoint from all other $S_{y_n}^{n+1}$ and from X_n . Now for any $y \in Y$ we let R_y^{n+1} be the ray $y P_y^n y_n S_{y_n}^{n+1}$. Let $P_y^{n+1} = R_y^{n+1} y_{n+1}$ be an initial segment of R_y^{n+1} of length at least $n+1$ and long enough to include P_y^n , and such that $Q_{n+1} \cap R_y^{n+1} \subseteq P_y^{n+1} \dot{y}_{n+1}$, and such that there are at least $n+1$ -disjoint paths between P_y^{n+1} and R - which is possible since both R and R_y^{n+1} are ϵ -rays. This completes the recursion step, and so the construction is complete. \square

Lemma B.6.5. *Let (X, Y) be a polypod of order n for ϵ in Γ with connection graph $K_{X,Y}$, $(S_y : y \in Y)$ be a family of tendrils for (X, Y) , and $(R_i : i \in I)$ be a set of disjoint ϵ -rays. Then for any transition function σ from \mathcal{S} to \mathcal{R} and every pair $y, y' \in Y$ such that there is a path from $\sigma(y)$ to $\sigma(y')$ otherwise avoiding $\sigma(Y)$ in $E(\text{RG}(R_i : i \in I))$, the edge yy' is in $E(K_{X,Y})$.*

Proof. Since σ is a transition function there exists a linkage from \mathcal{S} to \mathcal{R} after X which induces σ . This linkage gives us a family of tendrils $(S'_y : y \in Y)$ for (X, Y) such that S'_y is a tail of $R_{\sigma(y)}$ for each $y \in Y$. Then, by Lemmas B.3.2 and B.3.6, if $y, y' \in Y$ are such that there is a path from $\sigma(y)$ to $\sigma(y')$ otherwise avoiding $\sigma(Y)$ in $E(\text{RG}(R_i : i \in I))$, then S'_y and $S'_{y'}$ are adjacent in $\text{RG}(S'_y : y \in Y)$, and so y and y' are adjacent in $K_{X,Y}$. \square

Corollary and Definition B.6.6. Any two polypods for ϵ in Γ of the same order which attain their connection graphs have isomorphic connection graphs.

We will refer to the graph arising in this way for polypods of order n for ϵ in Γ as the n^{th} shape graph of the end ϵ . \square

B.6.2. Frames

Given a family of tendrils $(R_y : y \in Y)$ for a polypod (X, Y) there may be different families of tendrils $(R'_y : y \in Y)$ for (X, Y) such that each R_y shares a tail with some $R'_{\pi(y)}$. In order to understand the possible transition functions between different families of rays in ϵ it will be useful to understand the possible functions π that arise in this fashion.

To do so will we consider *frames*, finite subgraphs L which contain a path family between two sets of vertices $\alpha(Y)$ and $\beta(Y)$. For appropriate choices of $\alpha(Y)$ and $\beta(Y)$ these will be the subgraphs arising from a linkage from the family of tendrils $(R_y : y \in Y)$ to itself after X , each of which gives rise to a family $(R'_y : y \in Y)$ as above.

Some frames will contain multiple such path families, linking $\alpha(Y)$ to $\beta(Y)$ in different ways. For appropriately chosen frames the possible ways we can link $\alpha(Y)$ to $\beta(Y)$ will be restricted by the structure of $K_{X,Y}$, which will allow us relate this to the possible transition functions from $(R_y : y \in Y)$ to itself, and from there to the possible transition functions between different families of rays.

Definition B.6.7. Let Y be a finite set. A Y -*frame* (L, α, β) consists of a finite graph L together with two injections α and β from Y to $V(L)$. The set $A = \alpha(Y)$ is called the *source set* and the set $B = \beta(Y)$ is called the *target set*. A *weave* of the Y -frame is a family $\mathcal{Q} = (Q_y : y \in Y)$ of disjoint paths in L from A to B , where the initial vertex of Q_y is $\alpha(y)$ for each $y \in Y$. The *weave pattern* $\pi_{\mathcal{Q}}$ of \mathcal{Q} is the bijection from Y to itself sending y to the inverse image under β of the endvertex of Q_y . In other words, $\pi_{\mathcal{Q}}$ is the function so that every Q_y is an $\alpha(y)$ - $\beta(\pi_{\mathcal{Q}}(y))$ path. The *weave graph* $K_{\mathcal{Q}}$ of \mathcal{Q} has vertex set Y and an edge joining distinct vertices u and v of Y precisely when there is a path from Q_u to Q_v in L disjoint from all other Q_y . For a graph K with vertex set Y , we say that the Y -frame is *K -spartan* if all its weave graphs are subgraphs of K and all its weave patterns are automorphisms of K .

Connection graphs of polypods and weave graphs of frames are closely connected.

Lemma B.6.8. *Let (X, Y) be a polypod for ϵ in Γ attaining its connection graph $K_{X,Y}$ and let $\mathcal{R} = (R_y : y \in Y)$ be a family of tendrils for (X, Y) . Let L be any finite subgraph of Γ disjoint from X but meeting all the R_y . For each $y \in Y$*

let $\alpha(y)$ be the first vertex of R_y in L and $\beta(y)$ the last vertex of R_y in L . Then the Y -frame (L, α, β) is $K_{X,Y}$ -spartan.

Proof. Since there is some family of tendrils $(S_y: y \in Y)$ attaining $K_{X,Y}$ and there is, by Lemma B.3.9, a linkage from $(R_y: y \in Y)$ to $(S_y: y \in Y)$ after X and $V(L)$, we may assume without loss of generality that $\text{RG}(R_y: y \in Y)$ is isomorphic to $K_{X,Y}$.

For a given weave $\mathcal{Q} = (Q_y: y \in Y)$, applying the definition of the connection graph to the rays $R'_y = R_y\alpha(y)Q_y\beta(\pi_{\mathcal{Q}}(y))R_{\pi_{\mathcal{Q}}(y)}$ shows that $K_{\mathcal{Q}}$ is a subgraph of $K_{X,Y}$. Furthermore, since $\text{RG}(R_y: y \in Y)$ is isomorphic to $K_{X,Y}$, for any $uv \in E(K_{X,Y})$ there is a path from R_u to R_v which is disjoint from $R_u\alpha(u) \cup R_v\alpha(v) \cup L \cup X$ and which doesn't meet any other R_y , and so joins $R'_{\pi^{-1}(u)}$ to $R'_{\pi^{-1}(v)}$. So, the family of tendrils $(R'_y: y \in Y)$ witness that $\pi^{-1}(u)\pi^{-1}(v) \in E(K_{X,Y})$, and so $\pi_{\mathcal{Q}}$ is an automorphism of $K_{X,Y}$. \square

Corollary B.6.9. *Let (X, Y) be a polypod for ϵ in Γ attaining its connection graph $K_{X,Y}$ and let $\mathcal{R} = (R_y: y \in Y)$ be a family of tendrils for (X, Y) . Then for any transition function σ from \mathcal{R} to itself there is a $K_{X,Y}$ -spartan Y -frame for which both σ and the identity are weave patterns.*

Proof. Let $(P_y: y \in Y)$ be a linkage from \mathcal{R} to itself after X inducing σ , and let L be a finite subgraph of Γ containing $\bigcup_{y \in Y} P_y$ as well as a finite segment of each R_y , such that each P_y is a path between two such segments. Then the Y -frame on L which exists by Lemma B.6.8 has the desired properties. \square

Lemma B.6.10. *Let (X, Y) be a polypod for ϵ in Γ attaining its connection graph $K_{X,Y}$ and let $\mathcal{R} = (R_y: y \in Y)$ be a family of tendrils for (X, Y) . Then there is a $K_{X,Y}$ -spartan Y -frame for which both $K_{X,Y}$ and $\text{RG}(R_y: y \in Y)$ are weave graphs.*

Proof. By adding finitely many vertices to X if necessary, we may obtain a superset X' of X such that for any two of the R_y , if there is any path between them disjoint from all the other rays and X' , then there are infinitely many disjoint such paths. Let $(S_y: y \in Y)$ be any family of tendrils for (X, Y) with connection graph $K_{X,Y}$.

For each edge $e = uv$ of $\text{RG}(\mathcal{R})$ let P_e be a path from R_u to R_v disjoint from all the other R_y and from X' . Similarly for each edge $f = uv$ of $K_{X,Y}$ let Q_f be a

path from S_u to S_v disjoint from all the other S_y and from X' . Let $(P'_y : y \in Y)$ be a linkage from the S_y to the R_y after

$$X' \cup \bigcup_{e \in E(\text{RG}(\mathcal{R}))} P_e \cup \bigcup_{f \in E(K_{X,Y})} Q_f.$$

Let the initial vertex of P'_y be $\gamma(y)$ and the end vertex be $\beta(y)$. Let π be the permutation of Y by setting $\pi(y)$ to be the element of Y with $\beta(y)$ on $R_{\pi(y)}$. Let L be the graph given by the union of all paths of the form $S_y\gamma(y)$ and $R_{\pi(y)}\beta(y)$ together with P'_y , P_e and Q_e .

Letting α be the identity function on Y , it follows from Lemma B.6.8 that (L, α, β) is a $K_{X,Y}$ -spartan Y -frame. The paths Q_f witness that the weave graph for the paths $S_y\gamma(y)P'_y$ includes $K_{X,Y}$ and so, by $K_{X,Y}$ -spartanness, must be equal to $K_{X,Y}$. The paths P_e witness that the weave graph for the paths $R_y\beta(y)$ includes the ray graph $\text{RG}(\mathcal{R})$. However conversely, since $V(L)$ is disjoint from X' , if two of the R_y are joined in L by a path disjoint from the other rays in \mathcal{R} then they are joined by infinitely many, and hence adjacent in $\text{RG}(\mathcal{R})$. It follows that the weave graph is equal to $\text{RG}(\mathcal{R})$. \square

Hence to understand ray graphs and the transition functions between them it is useful to understand the possible weave graphs and weave patterns of spartan frames. Their structure can be captured in terms of automorphisms and cycles.

Definition B.6.11. Let K be a finite graph. An automorphism σ of K is called *local* if it is a cycle $(z_1 \dots z_t)$ where, for any $i \leq t$, there is an edge from z_i to $\sigma(z_i)$ in K . If $t \geq 3$ this means that $z_1 \dots z_t z_1$ is a cycle of K , and we call such cycles *turnable*. If $t = 2$ then we call the edge $z_1 z_2$ of K *flippable*. We say that an automorphism of K is *locally generated* if it is a product of local automorphisms.

Remark B.6.12. A cycle C in K is turnable if and only if all its vertices have the same neighbourhood in $K - C$, and whenever a chord of length $\ell \in \mathbb{N}$, i.e. a chord whose endvertices have distance ℓ on C , is present in $K[C]$, then all chords of length ℓ are present. Similarly an edge e of K is flippable if and only if its two endvertices have the same neighbourhood in $K - e$. Thus, if K is connected and contains at least three vertices, no vertex of degree one or cutvertex of K can lie on a turnable cycle or a flippable edge. So vertices of degree one and cutvertices in such graphs are preserved by locally generated automorphisms.

Lemma B.6.13. *Let $\mathcal{L} = (L, \alpha, \beta)$ be a K -spartan Y -frame which is K -spartan. Then for any two of its weave patterns π and π' the automorphism $\pi^{-1} \cdot \pi'$ of K is locally generated. Furthermore, if K is a weave graph for \mathcal{L} then each weave graph for \mathcal{L} contains a turnable cycle or a flippable edge of K .*

Proof. Let us suppose, for a contradiction, that the conclusion does not hold and let $\mathcal{L} = (L, \alpha, \beta)$ be a counterexample in which $|E(L)|$ is minimal. Let $\mathcal{P} = (P_y : y \in Y)$ and $\mathcal{Q} = (Q_y : y \in Y)$ be weaves for \mathcal{L} such that either $\pi_{\mathcal{P}} \neq \pi_{\mathcal{Q}}$ and $\pi_{\mathcal{P}}^{-1} \cdot \pi_{\mathcal{Q}}$ is not locally generated, or $K_{\mathcal{P}} = K$ and $K_{\mathcal{Q}}$ does not contain a turnable cycle or a flippable edge of K .

Each edge of L is in one path of \mathcal{P} or \mathcal{Q} since otherwise we could simply delete it. Similarly no edge appears in both \mathcal{P} and \mathcal{Q} since otherwise we could simply contract it. No vertex appears on just one of P_y or Q_y since otherwise we could contract one of the two incident edges. Vertices of L appearing in neither $\bigcup \mathcal{P}$ nor $\bigcup \mathcal{Q}$ are isolated and so may be ignored. Thus we may suppose that each edge of L appears in precisely one of \mathcal{P} or \mathcal{Q} , and that each vertex of L appears in both.

Let Z be the set of those $y \in Y$ for which $\alpha(y) \neq \beta(y)$. For any $z \in Z$ let $\gamma(z)$ be the second vertex of P_z , i.e. the neighbour of $\alpha(z)$ on P_z , and let $f(z) \in Y$ be chosen such that $\gamma(z)$ lies on $Q_{f(z)}$. Then since $\gamma(z) \neq \alpha(f(z))$ we have $f(z) \in Z$ for all $z \in Z$. Furthermore, Z is nonempty as \mathcal{P} and \mathcal{Q} are distinct. Let z be any element of Z . Then since Z is finite there must be $i < j$ with $f^i(z) = f^j(z)$, which means that $f^i(z) = f^{j-i}(f^i(z))$. Let $t > 0$ be minimal such that there is some $z_1 \in Z$ with $z_1 = f^t(z_1)$.

If $t = 1$ then we may delete the edge $\alpha(z_1)\gamma(z_1)$ and replace the path P_{z_1} with $\alpha(z_1)Q_{z_1}\gamma(z_1)P_{z_1}$. This preserves all of $\pi_{\mathcal{P}}$, $\pi_{\mathcal{Q}}$ and $K_{\mathcal{Q}}$, and can only make $K_{\mathcal{P}}$ bigger, contradicting the minimality of our counterexample. So we must have $t \geq 2$.

For each $i \leq t$ let z_i be $f^{i-1}(z_1)$ and let σ be the bijection $(z_1 z_2 \dots z_t)$ on Y . Let L' be the graph obtained from L by deleting all vertices of the form $\alpha(z_i)$. Let α' be the injection from Y to $V(L')$ sending z_i to $\gamma(z_i)$ for $i \leq t$ and sending any other $y \in Y$ to $\alpha(y)$. Then (L', α', β) is a Y -frame. For any weave $(\hat{P}_y : y \in Y)$ in this Y -frame, $(P''_y : y \in Y)$ where $P''_{z_i} = \alpha(z_i)\gamma(z_i)\hat{P}_{z_i}$ for every $i \leq t$ and $P''_y = \alpha(y)$ for every $y \in Y \setminus \{z_1, \dots, z_t\}$ is a weave in (L, α, β) with the same weave pattern and whose weave graph includes that of $(\hat{P}_y : y \in Y)$. Thus (L', α', β) is K -spartan.

Let P'_y be $\alpha'(y)P_y$ and Q'_y be $\alpha'(y)Q_{\sigma(y)}$ for each $y \in Y$. Now set $\mathcal{P}' = (P'_y : y \in Y)$ and $\mathcal{Q}' = (Q'_y : y \in Y)$. Then we have $\pi_{\mathcal{Q}'} = \pi_{\mathcal{Q}} \cdot \sigma$ and so $\sigma = \pi_{\mathcal{Q}'}^{-1} \cdot \pi_{\mathcal{Q}}$ is an au-

tomorphism of K since π_Q is an automorphism of K by the K -spartanness. For any $i \leq t$ the edge $\alpha(z_i)\gamma(z_i)$ witnesses that $z_i\sigma(z_i)$ is an edge of K_Q , and hence, since \mathcal{L} is K -spartan, also an edge of K , and so σ is a local automorphism of K . It follows that K_Q includes a turnable cycle or a flippable edge. Finally, by the minimality of $|E(L)|$ we know that $\pi_{\mathcal{P}'}^{-1} \cdot \pi_Q$ is locally generated and hence so is $\pi_{\mathcal{P}'}^{-1} \cdot \pi_Q = \pi_{\mathcal{P}'}^{-1} \cdot \pi_Q \cdot \sigma^{-1}$. This is the desired contradiction. \square

Finally, the following two lemmas are the main conclusions of this section:

Lemma B.6.14. *Let (X, Y) be a polypod attaining its connection graph $K_{X,Y}$ such that $K_{X,Y}$ is a cycle of length at least 4. Then for any family of tendrils \mathcal{R} for this polypod the ray graph is $K_{X,Y}$. Furthermore, any transition function from \mathcal{R} to itself preserves each of the cyclic orientations of $K_{X,Y}$.*

Proof. By Lemma B.6.10 there is some $K_{X,Y}$ -spartan Y -frame for which both $K_{X,Y}$ and the ray graph $\text{RG}(\mathcal{R})$ are weave graphs. Since $K_{X,Y}$ is a cycle of length at least 4 and hence has no flippable edges, the ray graph must include a cycle by Lemma B.6.13 and so since it is a subgraph of $K_{X,Y}$ it must be the whole of $K_{X,Y}$. Similarly Lemma B.6.13 together with Corollary B.6.9 shows that all transition functions must be locally generated and so must preserve the orientation. \square

Lemma B.6.15. *Let (X, Y) be a polypod attaining its connection graph $K_{X,Y}$ such that $K_{X,Y}$ includes a bare path P whose edges are bridges. Let \mathcal{R} be a family of tendrils for (X, Y) whose ray graph is $K_{X,Y}$. Then for any transition function σ from \mathcal{R} to itself, the restriction of σ to P is the identity.*

Proof. By Lemmas B.6.9 and B.6.13 any transition function must be a locally generated automorphism of $K_{X,Y}$, and so by Remark B.6.12 it cannot move the vertices of the bare path, which are vertices of degree one or cutvertices. \square

B.7. Grid-like and half-grid-like ends

We are now in a position to analyse the different kinds of thick ends which can arise in a graph in terms of the possible ray graphs and the transition functions between them. The first kind of ends are the pebbly ends, in which, by Corollary B.5.3, for any n we can find a family of n disjoint rays whose ray graph is K_n and for which every function $\sigma: [n] \rightarrow [n]$ is a transition function.

So, in the following let us fix a graph Γ with a thick non-pebbly end ϵ and a number $N \in \mathbb{N}$, where $N \geq 3$, such that ϵ is not N -pebble win. Under these circumstances we get nontrivial restrictions on the ray graphs and the transition functions between them. There are two essentially different cases, corresponding to the two cases in Corollary B.5.4: The grid-like and the half-grid-like case.

B.7.1. Grid-like ends

The first case focuses on ends which behave like that of the infinite grid. In this case, all large enough ray graphs are cycles and all transition functions between them preserve the cyclic order.

Formally, we say that the end ϵ is *grid-like* if the $(N + 2)^{\text{nd}}$ shape graph for ϵ is a cycle. For the rest of this subsection we will assume that ϵ is grid-like. Let us fix some polypod (X, Y) of order $N + 2$ attaining its connection graph. Let $(S_y: y \in Y)$ be a family of tendrils for (X, Y) whose ray graph is the cycle $C_{N+2} = K_{X,Y}$.

Lemma B.7.1. *The ray graph K for any family $(R_i: i \in I)$ of ϵ -rays in Γ with $|I| \geq N + 2$ is a cycle.*

Proof. By Corollary B.5.4, K is either a cycle or contains a bridge. However, given any edge $ij \in E(K)$, let $J \subseteq I$ be such that $i, j \in J$ and $|J| = N + 2$. Let $(T_y: y \in Y)$ be a family of tendrils for (X, Y) obtained by transitioning from $(S_y: y \in Y)$ to $(R_j: j \in J)$ after X along some linkage. By Lemma B.3.2, the ray graph K_J of $(R_i: i \in J)$ is isomorphic to the ray graph of $(T_y: y \in Y)$, which is a cycle by Lemma B.6.14.

Hence, ij is not a bridge of K_J , and it is easy to see that this implies that ij is not a bridge of K . Hence, K is a cycle. \square

Given a cycle C a *cyclic orientation* of C is an orientation of the graph C which does not have any sink. Note that any cycle has precisely two cyclic orientations. Given a cyclic orientation and three distinct vertices x, y, z we say that they *appear consecutively in the order* (x, y, z) if y lies on the unique directed path from x to z . Given two cycles C, C' , each with a cyclic orientation, we say that an injection $f: V(C) \rightarrow V(C')$ *preserves the cyclic orientation* if whenever three

distinct vertices x, y and z appear on C in the order (x, y, z) then their images appear on C' in the order $(f(x), f(y), f(z))$.

We will now choose cyclic orientations of every large enough ray-graph such that the transition functions preserve the cyclic orders corresponding to those orientations. To that end, we fix a cyclic orientation of $K_{X,Y}$. We say that a cyclic orientation of the ray graph for a family $(R_i: i \in I)$ of at least $N + 3$ disjoint ϵ -rays is *correct* if there is a transition function σ from the S_y to the R_i which preserves the cyclic orientation of $K_{X,Y}$.

Lemma B.7.2. *For any family $(R_i: i \in I)$ of at least $N + 3$ disjoint ϵ -rays there is precisely one correct cyclic orientation of its ray graph.*

Proof. We first claim that there is at least one correct cyclic orientation. By Lemma B.3.9, there is a transition function σ from the S_y to some subset J of I , and we claim that there is some cyclic orientation of the ray graph K of $(R_i: i \in I)$ such that σ preserves the cyclic orientation of $K_{X,Y}$.

We first note that the ray graph K_J of $(R_i: i \in J)$ is a cycle by Lemma B.7.1, and it is obtained from K by subdividing edges, which doesn't affect the cyclic order. Hence it is sufficient to show that there is some cyclic orientation of K_J such that σ preserves the cyclic orientation of $K_{X,Y}$.

Since each linkage inducing σ gives rise to a family of tendrils $(S'_y: y \in Y)$ where S'_y shares a tail with $R_{\sigma(y)}$, it follows that if $\sigma(y)$ and $\sigma(y)'$ are adjacent in K_J then y and y' are adjacent in $K_{X,Y}$. Since both K_J and $K_{X,Y}$ are cycles, it follows that there is some cyclic orientation of K_J such that σ preserves the cyclic orientation of $K_{X,Y}$.

Suppose for a contradiction that there are two, and let σ and σ' be transition functions witnessing that both orientations of the ray graph are correct. By Lemma B.4.2 we may assume without loss of generality that the images of σ and σ' are the same. Call this common image I' . Since the ray graphs of $(R_i: i \in I)$ and $(R_i: i \in I')$ are both cycles, the former is obtained from the latter by subdivision of edges. Since this does not affect the cyclic order, we may assume without loss of generality that $I' = I$. By Lemma B.3.9 again, there is some transition function τ from the R_i to the S_y . By Lemma B.6.14, both $\tau \cdot \sigma$ and $\tau \cdot \sigma'$ must preserve the cyclic order, which is the desired contradiction. \square

It therefore makes sense to refer to *the* correct orientation of a ray graph.

Corollary B.7.3. *Any transition function between two families of at least $N + 3$ ϵ -rays preserves the correct orientations of their ray graphs.*

Proof. Suppose that $\mathcal{R} = (R_i : i \in I)$ and $\mathcal{T} = (T_j : j \in J)$ are families of at least $N + 3$ rays and σ is a transition function from \mathcal{R} to \mathcal{T} .

Let us fix some transition function τ from $(S_y : y \in Y)$ to \mathcal{R} and let \mathcal{P} be a linkage from $(S_y : y \in Y)$ to \mathcal{R} which induces τ . For any finite $X \subseteq V(\Gamma)$ there is a linkage \mathcal{P}' from \mathcal{R} to \mathcal{T} which is after $\bigcup \mathcal{P} \cup X$ and which induces σ . Then, $((S_y : y \in Y) \circ_{\mathcal{P}} \mathcal{R}) \circ_{\mathcal{P}'} \mathcal{T}$ is a linkage from $(S_y : y \in Y)$ to \mathcal{T} which is after X and induces $\sigma \cdot \tau$. It follows that $\sigma \cdot \tau$ is a transition function from $(S_y : y \in Y)$ to \mathcal{T} .

However, by the definition of correct orientation and Lemma B.7.2, τ and $\sigma \cdot \tau$ both preserve the cyclic orientation of $K_{X,Y}$, and hence σ must preserve the correct orientation of the ray graphs of \mathcal{R} and \mathcal{T} . \square

B.7.2. Half-grid-like ends

In this subsection we suppose that ϵ is thick but neither pebbly nor grid-like. We shall call such ends *half-grid-like*, since as we shall shortly see in this case the ray graphs and the transition functions between them behave similarly to those for the unique end of the half-grid. Note that, by definition Theorem B.1.2 holds.

We will need to carefully consider how the ray graphs are divided up by their cutvertices. In particular, for a graph K and vertices x and y of K we will denote by $C^{xy}(K)$ the union of all components of $K - x$ which do not contain y , and we will denote by K^{xy} the graph $K - C^{xy}(K) - C^{yx}(K)$. We will refer to K^{xy} as the part of K *between* x and y .

As in the last subsection, let (X, Y) be a polypod of order $N + 2$ attaining its connection graph and let $(S_y : y \in Y)$ be a family of tendrils for (X, Y) with ray graph $K_{X,Y}$, which by assumption is not a cycle. By Corollary B.5.4 there is a bare path of length at least 1 in $K_{X,Y}$ all of whose edges are bridges. Let $y_1 y_2$ be any edge of that path. Without loss of generality we have $C^{y_1 y_2}(K_{X,Y}) \neq \emptyset$.

Throughout the remainder of this section we will always consider arbitrary families $\mathcal{R} = (R_i : i \in I)$ of disjoint ϵ -rays with $|I| \geq N + 3$. We will write K to denote the ray graph of \mathcal{R} .

Remark B.7.4. For any transition function σ from the S_y to the R_i we have the inclusions $\sigma[V(C^{y_1 y_2}(K_{X,Y}))] \subseteq V(C^{\sigma(y_1)\sigma(y_2)}(K))$ and $\sigma[V(C^{y_2 y_1}(K_{X,Y}))] \subseteq$

$V(C^{\sigma(y_2)\sigma(y_1)}(K))$ by Lemma B.6.5. Thus $\sigma[Y]$ and $V(K^{\sigma(y_1)\sigma(y_2)})$ meet precisely in $\sigma(y_1)$ and $\sigma(y_2)$.

Lemma B.7.5. *For any transition function σ from the S_y to the R_i the graph $K^{\sigma(y_1)\sigma(y_2)}$ is a path from $\sigma(y_1)$ to $\sigma(y_2)$. This path is a bare path in K and all of its edges are bridges.*

Proof. Since K is connected, $K^{\sigma(y_1)\sigma(y_2)}$ must include a path P from $\sigma(y_1)$ to $\sigma(y_2)$. If it is not equal to that path then it follows from Lemma B.4.2 that the function σ' , which we define to be just like σ except for $\sigma'(y_1) = \sigma(y_2)$ and $\sigma'(y_2) = \sigma(y_1)$, is a transition function from the S_y to the R_i . But then by Remark B.7.4 we have

$$\begin{aligned} \sigma[V(C^{y_1y_2}(K_{X,Y}))] &\subseteq V(C^{\sigma(y_1)\sigma(y_2)}(K)) \cap V(C^{\sigma'(y_1)\sigma'(y_2)}(K)) \\ &= V(C^{\sigma(y_1)\sigma(y_2)}(K)) \cap V(C^{\sigma(y_2)\sigma(y_1)}(K)) = \emptyset, \end{aligned}$$

a contradiction. The last sentence of the lemma follows from the definition of $K^{\sigma(y_1)\sigma(y_2)}$. \square

Given a path P with endvertices s and t we say *the orientation of P from s to t* to mean the total order \leq on the vertices of P where $a \leq b$ if and only if a lies on sPb , in this case we say that a lies before b . Note that every path with at least one edge has precisely two orientations.

Now, we fix a transition function σ_{\max} from the S_y to the R_i so that the path $P := K^{\sigma_{\max}(y_1)\sigma_{\max}(y_2)}$ is as long as possible. We call P the *central* path of K and the orientation of P from $\sigma_{\max}(y_1)$ to $\sigma_{\max}(y_2)$ the *correct* orientation.

We first note that, for large enough families of rays almost all of the ray graph lies on the central path.

Lemma B.7.6. *At most N vertices of K are not on the central path.*

Proof. By Remark B.7.4 we have $\sigma_{\max}[V(C^{y_1y_2}(K_{X,Y}))] \subseteq V(C^{\sigma_{\max}(y_1)\sigma_{\max}(y_2)}(K))$. If it were a proper subset, then we would be able to use Lemma B.4.2 to produce a transition function in which this path is longer. So we must have $\sigma_{\max}[V(C^{y_1y_2}(K_{X,Y}))] = V(C^{\sigma_{\max}(y_1)\sigma_{\max}(y_2)}(K))$. Similarly we also get that $\sigma_{\max}[V(C^{y_2y_1}(K_{X,Y}))] = V(C^{\sigma_{\max}(y_2)\sigma_{\max}(y_1)}(K))$. However, since y_1y_2 is a bridge, $|V(C^{y_1y_2}(K_{X,Y})) \cup V(C^{y_2y_1}(K_{X,Y}))| = N$ and so at most N vertices of K are not on the central path. \square

We call P the *central path* of K and the orientation of P from $\sigma_{\max}(y_1)$ to $\sigma_{\max}(y_2)$ the *correct orientation*. We note the following simple corollary, which will be useful in later work.

Corollary B.7.7. *For any $i \in I$ if $\text{RG}(\mathcal{R}) - i$ has precisely two components, each of size at least $N + 1$, then i is an inner vertex of the central path of $\text{RG}(\mathcal{R})$.*

Proof. By Lemma B.7.6 both components of $\text{RG}(\mathcal{R}) - i$ contain a vertex of the central path. However, since all the edges of the central path are bridges, it follows that i lies between these two vertices on the central path. \square

We can in fact determine the central path and its correct orientation by considering the possible transition functions from the S_y to the R_i .

Lemma B.7.8. *For any two vertices v_1 and v_2 of K , there exists a transition function $\sigma: V(K_{X,Y}) \rightarrow V(K)$ with $\sigma(y_1) = v_1$ and $\sigma(y_2) = v_2$ if and only if v_1 and v_2 both lie on P , with v_1 before v_2 .*

Proof. The ‘if’ direction is clear by applying Lemma B.4.2 to σ_{\max} . For the ‘only if’ direction, we begin by setting $c_1 = |V(C^{y_1 y_2}(K_{X,Y}))|$ and $c_2 = |V(C^{y_2 y_1}(K_{X,Y}))|$. We enumerate the set $V(C^{y_1 y_2}(K_{X,Y}))$ as $y_3 \dots y_{c_1+2}$ and $V(C^{y_2 y_1}(K_{X,Y}))$ as $y_{c_1+3} \dots y_{c_1+c_2+2}$. Then for any $(N + 2)$ -tuple (x_1, \dots, x_{N+2}) of distinct vertices which is achievable in the $(\sigma_{\max}(y_1), \dots, \sigma_{\max}(y_{N+2}))$ -pebble-pushing game on K we must have the following three properties, since they are preserved by any single move:

- x_1 and x_2 lie on P , with x_1 before x_2 .
- $\{x_3, \dots, x_{c_1+2}\} \subseteq V(C^{x_1 x_2}(K))$.
- $\{x_{c_1+3}, \dots, x_{c_1+c_2+2}\} \subseteq V(C^{x_2 x_1}(K))$.

Now let σ be any transition function from the S_y to the R_i . Let (x_1, \dots, x_{N+2}) be an $(N + 2)$ -tuple achievable in the $(\sigma_{\max}(y_1), \dots, \sigma_{\max}(y_{N+2}))$ -pebble-pushing game such that $\{x_1, \dots, x_{N+2}\} = \sigma[Y]$. By Lemma B.4.2, the function σ' sending y_i to x_i for each $i \leq N + 2$ is also a transition function and $\sigma'[Y] = \sigma[Y]$. Let τ be a transition function from $(R_i: i \in \sigma[Y])$ to the S_y . Then, by Lemma B.6.15, both $\tau \cdot \sigma$ and $\tau \cdot \sigma'$ keep both y_1 and y_2 fixed. Thus $\sigma(y_1) = \sigma'(y_1) = x_1$ and furthermore $\sigma(y_2) = \sigma'(y_2) = x_2$. As noted above, this means that $\sigma(y_1)$ and $\sigma(y_2)$ both lie on P with $\sigma(y_1)$ before $\sigma(y_2)$, as desired. \square

Thus the central path and the correct orientation depend only on our choice of y_1 and y_2 . Hence, we get the following corollary.

Corollary B.7.9. *Each ray graph on at least $N + 3$ vertices contains a unique central path with a correct orientation and every transition function between two families of at least $N + 3$ ϵ -rays sends vertices of the central path to vertices of the central path and preserves the correct orientation.*

Proof. Consider the family $\mathcal{R} = (R_i: i \in I)$ with its ray graph K and another family $\mathcal{T} = (T_j: j \in J)$ of at least $N + 3$ rays, with ray graph $K_{\mathcal{T}}$, and let τ be a transition function from \mathcal{R} to \mathcal{T} .

Let v_1, v_2 be two vertices in the central path P of K with v_1 before v_2 . By Lemma B.7.8 there is transition function σ from $(S_y: y \in Y)$ to \mathcal{R} with $\sigma(y_1) = v_1$ and $\sigma(y_2) = v_2$.

Then, as in Lemma B.7.3, it is clear that $\tau \cdot \sigma$ is a transition function from $(S_y: y \in Y)$ to \mathcal{T} . However since $\tau \cdot \sigma(y_1) = \tau(v_1)$ and $\tau \cdot \sigma(y_2) = \tau(v_2)$, it follows from Lemma B.7.8 that $\tau(v_1)$ and $\tau(v_2)$ both lie on the central path $P_{\mathcal{T}}$ of $K_{\mathcal{T}}$ with $\tau(v_1)$ before $\tau(v_2)$, and hence τ sends vertices of P to vertices of $P_{\mathcal{T}}$ and preserves the correct orientation. \square

Lemma B.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 $\text{RG}(\mathcal{T})$.*

Proof. That $\sigma(x)$ is an inner vertex of the central path of $\text{RG}(\mathcal{T})$ follows from Corollary B.7.9. We note, by Lemma B.6.5, given any family of rays \mathcal{K} and a transition function γ from \mathcal{S} to \mathcal{K} , if y separates x from z in $K_{X,Y}$ then $\gamma(y)$ separates $\gamma(x)$ from $\gamma(z)$ in $\text{RG}(\mathcal{K})$.

Let $\tau: V(K_{X,Y}) \rightarrow V(\text{RG}(\mathcal{R}))$ be a transition function with $\tau(y_1) = x$ which exists by Lemma B.7.8. Since x is an inner vertex of the central path of $\text{RG}(\mathcal{R})$, there are exactly two components of $\text{RG}(\mathcal{R}) - x$, one containing v_1 and one containing v_2 . Furthermore, by Lemma B.6.5, it follows that $\tau(C^{y_1 y_2}(K_{X,Y}))$ and $\tau(C^{y_2 y_1}(K_{X,Y} \cup \{y_2\}))$ are contained in different components of $\text{RG}(\mathcal{R}) - x$.

Hence, by Lemma B.4.2 we may assume without loss of generality that $v_1, v_2 \in \tau(V(K_{X,Y}))$, where y_1 separates $w_1 := \tau^{-1}(v_1)$ and $w_2 := \tau^{-1}(v_2)$ in $K_{X,Y}$.

However, by the remark above applied to the transition function $\sigma \cdot \tau$ we conclude that $\sigma(x) = \sigma \cdot \tau(y_1)$ separates $\sigma(v_1) = \sigma \cdot \tau(w_1)$ from $\sigma(v_2) = \sigma \cdot \tau(w_2)$. \square

B.8. G -tribes and concentration of G -tribes towards an end

To show that a given graph G is \preceq -ubiquitous, we shall assume that $nG \preceq \Gamma$ holds for every $n \in \mathbb{N}$ and show that this implies $\aleph_0 G \preceq \Gamma$. To this end we use the following notation for such collections of nG in Γ , most of which we established in [13].

Definition B.8.1 (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}' in Γ is a G -subtribe* 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 $\varphi_{F'}: F' \rightarrow \Psi(F')$ such that $V(H') \subseteq V(\varphi_{F'}(H'))$ for each $H' \in F'$. 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 set X of vertices 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} . It is *strongly concentrated at* ϵ if additionally, for every finite vertex set X of Γ , every member H of \mathcal{F} intersects $C(X, \epsilon)$.

We note that 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 [13].

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

Lemma B.8.2 (Removing a thin subtribe, [13, 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 B.8.3 (Pigeon hole principle for thick G -tribes, [13, Lemma 5.3]). *Suppose for some $k \in \mathbb{N}$, we have a k -colouring $c: \bigcup \mathcal{F} \rightarrow [k]$ 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} .*

Note that, in the following lemmas, it is necessary that G is connected, so that every member of the G -tribe is a connected graph.

Lemma B.8.4 ([13, Lemma 5.4]). *Let G be a connected graph and Γ a graph containing a thick 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 B.8.5 ([13, Lemma 5.5]). *Let G be a connected graph and Γ a graph containing a thick 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 , too.*

Lemma B.8.6. *Let G be a connected graph and Γ a graph containing a thick G -tribe \mathcal{F} concentrated at an end $\epsilon \in \Omega(\Gamma)$. Then either $\aleph_0 G \preceq \Gamma$, or there is a thick flat subtribe of \mathcal{F} which is strongly concentrated at ϵ .*

Proof. Suppose that no thick flat subtribe of \mathcal{F} is strongly concentrated at ϵ . We construct an $\aleph_0 G \preceq \Gamma$ by recursively choosing disjoint IGs H_1, H_2, \dots in Γ as follows: Having chosen H_1, H_2, \dots, H_n such that for some finite set X_n we have

$$H_i \cap C(X_n, \epsilon) = \emptyset$$

for all $i \in [n]$, then by Lemma B.8.5(1), there is still a thick flat subtribe \mathcal{F}'_n of \mathcal{F} contained in $C(X_n, \epsilon)$. Since by assumption, \mathcal{F}'_n is not strongly concentrated at ϵ ,

we may pick $H_{n+1} \in \mathcal{F}'_n$ and a finite set $X_{n+1} \supseteq X_n$ with $H_{n+1} \cap C(X_{n+1}, \epsilon) = \emptyset$. Then the union of all the H_i is an $\aleph_0 G \preceq \Gamma$. \square

The following lemma will show that we can restrict ourselves to thick G -tribes which are concentrated at thick ends.

Lemma B.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 $\aleph_0 G \preceq \Gamma$.*

Proof. Since ϵ is thin, we may assume by Proposition B.2.4 that only finitely many vertices dominate ϵ . Deleting these yields a subgraph of Γ in which there is still a thick G -tribe concentrated at ϵ . Hence we may assume without loss of generality that ϵ is not dominated by any vertex in Γ .

Let $k \in \mathbb{N}$ be the degree of ϵ . By [42, Corollary 5.5] there is a sequence of vertex sets $(S_n : n \in \mathbb{N})$ such that:

- $|S_n| = k$,
- $C(S_{n+1}, \epsilon) \subseteq C(S_n, \epsilon)$, and
- $\bigcap_{n \in \mathbb{N}} C(S_n, \epsilon) = \emptyset$.

Suppose there is a thick subtribe \mathcal{F}' of \mathcal{F} which is strongly concentrated at ϵ . For any $F \in \mathcal{F}'$ there is an $N_F \in \mathbb{N}$ such that $H \setminus C(S_{N_F}, \epsilon) \neq \emptyset$ for all $H \in F$ by the properties of the sequence. Furthermore, since \mathcal{F}' is strongly concentrated, $H \cap C(S_{N_F}, \epsilon) \neq \emptyset$ as well for each $H \in F$.

Let $F \in \mathcal{F}'$ be such that $|F| > k$. Since G is connected, so is H , and so from the above it follows that $H \cap S_{N_F} \neq \emptyset$ for each $H \in F$, contradicting the fact that $|S_{N_F}| = k < |F|$. Thus $\aleph_0 G \preceq \Gamma$ by Lemma B.8.6. \square

Note that, whilst concentration is hereditary for subtribes, strong concentration is not. However if we restrict to *flat* subtribes, then strong concentration is a hereditary property.

Let us show see how ends of the members of a strongly concentrated tribe relate to ends of the host graph Γ . Let G be a connected graph and $H \subseteq \Gamma$ an IG . By Lemmas B.3.2 and B.3.4, if $\omega \in \Omega(G)$ and R_1 and $R_2 \in \omega$ then the pullbacks $H^\downarrow(R_1)$ and $H^\downarrow(R_2)$ belong to the same end $\omega' \in \Omega(\Gamma)$. Hence, H determines for every end $\omega \in G$ a *pullback end* $H(\omega) \in \Omega(\Gamma)$. The next lemma is where we need to use the assumption that G is locally finite.

Lemma B.8.8. *Let G be a locally finite connected graph and Γ a graph containing a thick G -tribe \mathcal{F} strongly concentrated at an end $\epsilon \in \Omega(\Gamma)$, where every member is a tidy IG . Then either $\aleph_0 G \preceq \Gamma$, or there is a flat subtribe \mathcal{F}' of \mathcal{F} such that for every $H \in \bigcup \mathcal{F}'$ there is an end $\omega_H \in \Omega(G)$ such that $H(\omega_H) = \epsilon$.*

Proof. Since G is locally finite and every $H \in \bigcup \mathcal{F}$ is tidy, the branch sets $H(v)$ are finite for each $v \in V(G)$. If ϵ is dominated by infinitely many vertices, then $\aleph_0 G \preceq \Gamma$ by Proposition B.2.4, since every locally finite connected graph is countable. If this is not the case, then there is some $k \in \mathbb{N}$ such that ϵ is dominated by k vertices and so for every $F \in \mathcal{F}$ at most k of the $H \in F$ contain vertices which dominate ϵ in Γ . Therefore, there is a thick flat subtribe \mathcal{F}' of \mathcal{F} such that no $H \in \bigcup \mathcal{F}'$ contains a vertex dominating ϵ in Γ . Note that \mathcal{F}' is still strongly concentrated at ϵ , and every branch set of every $H \in \bigcup \mathcal{F}'$ is finite.

Since \mathcal{F}' is strongly concentrated at ϵ , for every finite vertex set X of Γ , every $H \in \bigcup \mathcal{F}'$ intersects $C(X, \epsilon)$. By a standard argument, since H as a connected infinite graph does not contain a vertex dominating ϵ in Γ , instead H contains a ray $R_H \in \epsilon$.

Since each branch set $H(v)$ is finite, R_H meets infinitely many branch sets. Let us consider the subgraph $K \subseteq G$ consisting of all the edges (v, w) such that R_H uses an edge between $H(v)$ and $H(w)$. Note that, since there is a edge in H between $H(v)$ and $H(w)$ if and only if $(v, w) \in E(G)$, K is well-defined and connected.

K is then an infinite connected subgraph of a locally finite graph, and as such contains a ray S_H in G . Since the edges between $H(v)$ and $H(w)$, if they exist, were unique, it follows that the pullback $H^\downarrow(S_H)$ of S_H has infinitely many edges in common with R_H , and so tends to ϵ in Γ . Therefore, if S_H tends to ω_H in $\Omega(G)$, then $H(\omega_H) = \epsilon$. \square

B.9. Ubiquity of minors of the half-grid

Here, and in the following, we denote by \mathbb{H} the infinite, one-ended, cubic hexagonal half-grid (see Figure B.2). The following theorem of Halin is one of the cornerstones of infinite graph theory.

Theorem B.9.1 (Halin, see [24, Theorem 8.2.6]). *Whenever a graph Γ contains a thick end, then $\mathbb{H} \leq \Gamma$.* \square




Figure B.2.: The hexagonal half-grid \mathbb{H} .

In [46], Halin used this result to show that every topological minor of \mathbb{H} is ubiquitous with respect to the topological minor relation \leq . In particular, trees of maximum degree 3 are ubiquitous with respect to \leq .

However, the following argument, which is a slight adaptation of Halin's, shows that every connected minor of \mathbb{H} is ubiquitous with respect to the minor relation. In particular, the dominated ray, the dominated double ray, and all countable trees are ubiquitous with respect to the minor relation.

The main difference to Halin's original proof is that, since he was only considering locally finite graphs, he was able to assume that the host graph Γ was also locally finite.

We will need the following result of Halin.

Lemma B.9.2 ([46, (4) in Section 3]). $\aleph_0\mathbb{H}$ is a topological minor of \mathbb{H} .

Theorem B.1.5. Any connected minor of the half-grid $\mathbb{N}\square\mathbb{Z}$ is \preceq -ubiquitous.

Proof. Suppose $G \preceq \mathbb{N}\square\mathbb{Z}$ is a minor of the half-grid, and Γ is a graph such that $nG \preceq \Gamma$ for each $n \in \mathbb{N}$. By Lemma B.8.4 we may assume there is an end ϵ of Γ and a thick G -tribe \mathcal{F} which is concentrated at ϵ . By Lemma B.8.7 we may assume that ϵ is thick. Hence $\mathbb{H} \leq \Gamma$ by Theorem B.9.1, and with Lemma B.9.2 we obtain

$$\aleph_0 G \preceq \aleph_0(\mathbb{N}\square\mathbb{Z}) \preceq \aleph_0\mathbb{H} \leq \mathbb{H} \leq \Gamma. \quad \square$$

Lemma B.9.3. \mathbb{H} contains every countable tree as a minor.

Proof. It is easy to see that the infinite binary tree T_2 embeds into \mathbb{H} as a topological minor. It is also easy to see that countably regular tree T_∞ where every vertex has

infinite degree embeds into T_2 as a minor. And obviously, every countable tree T is a subgraph of T_∞ . Hence we have

$$T \subseteq T_\infty \preceq T_2 \preceq \mathbb{H}$$

from which the result follows. \square

Corollary B.9.4. *All countable trees are ubiquitous with respect to the minor relation.*

Proof. This is an immediate consequence of Lemma B.9.3 and Theorem B.1.5. \square

B.10. Proof of main results

The following technical result contains most of the work for the proof of Theorem B.1.3, but is stated so as to be applicable in a later paper [15].

Lemma B.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 $H \in \bigcup \mathcal{F}'$.*

Proof. Let \mathcal{F}'' be the flat subtribe of \mathcal{F} given by $\mathcal{F}'' = \{F'' : F \in \mathcal{F}\}$ with

$$F'' = \{H : H \in F \text{ and } \omega_H \text{ is not linear}\}.$$

Suppose for a contradiction that \mathcal{F}'' is thick. Then, there is some $F \in \mathcal{F}$ which contains $k + 2$ disjoint IG s, H_1, H_2, \dots, H_{k+2} , where k is such that ϵ is not k -pebble-win. By assumption ω_{H_i} is not linear for each i , and so for each i there is a family of disjoint rays $\{R_1^i, R_2^i, \dots, R_{m_i}^i\}$ in G tending to ω_{H_i} whose ray graph in G is not a path. Let

$$\mathcal{S} = (H_i^\downarrow(R_j^i) : i \in [k + 2], j \in [m_i]).$$

By construction, \mathcal{S} is a disjoint family of ϵ -rays in Γ , and by Lemmas B.3.3 and B.3.4, $\text{RG}_\Gamma(\mathcal{S})$ contains disjoint subgraphs K_1, K_2, \dots, K_{k+2} that satisfy $K_i \cong \text{RG}_G(R_j^i : j \in [m_i])$. However, by Corollary B.5.4, there is a set X of vertices of size at most k such that $\text{RG}_\Gamma(\mathcal{S}) - X$ is a bare path P . However, then some $K_i \subseteq P$ is a path, a contradiction.

Remark B.10.2. In fact, every ray graph in the full-grid is a cycle (see Figure B.3).

Theorem B.1.6. *For every locally finite connected graph G , both $G \square \mathbb{Z}$ and $G \square \mathbb{N}$ are \preceq -ubiquitous.*

Proof. If G is a path or a ray, then $G \square \mathbb{Z}$ is a subgraph of the half-grid $\mathbb{N} \square \mathbb{Z}$ and thus \preceq -ubiquitous by Theorem B.1.5. If G is a double ray, then $G \square \mathbb{Z}$ is the full-grid and thus \preceq -ubiquitous by Corollary B.1.4.

Otherwise, let G' be a finite connected subgraph of G which is not a path and let H be \mathbb{Z} or \mathbb{N} . We note first that $G \square H$ has a unique end. Furthermore, for any ray R of H it is clear that G' is a subgraph of $\text{RG}_{G \square H}(\{\{v\} \square R\}_{v \in V(G')})$, and so this ray graph is not a path. Hence by Lemma B.3.5, $G \square H$ has nowhere-linear end structure and is therefore \preceq -ubiquitous by Theorem B.1.3. \square

Finally let us prove Theorem B.1.7. Recall that for $k \in \mathbb{N}$ we let DR_k denote the graph formed by taking a ray R together with k vertices v_1, v_2, \dots, v_k adjacent to every vertex in R . We shall need the following strengthening of Proposition B.2.3.

A *comb* is a union of a ray R with infinitely many disjoint finite paths, all having precisely their first vertex on R . The last vertices of these paths are the *teeth* of the comb.

Proposition B.10.3. [24, Proposition 8.2.2] *Let U be an infinite set of vertices in a connected graph G . Then G either contains a comb with all teeth in U or a subdivision of an infinite star with all leaves in U .*

Theorem B.1.7. *The k -fold dominated ray DR_k is \preceq -ubiquitous for every $k \in \mathbb{N}$.*

Proof. Let $R = x_1x_2x_3\dots$ be the ray as stated in the definition of DR_k and let v_1, v_2, \dots, v_k denote the vertices adjacent to each vertex of R . Note that if $k \leq 2$ then DR_k is a minor of the half-grid, and hence \preceq -ubiquity follows from Theorem B.1.5.

Suppose then that $k \geq 3$ and Γ is a graph which contains a thick DR_k -tribe \mathcal{F} each of whose members is tidy. We may further assume, without loss of generality, that for each $H \in \bigcup \mathcal{F}$, each $i \in [k]$, and each vertex x of $H(v_i)$, every component of $H(v_i) - x$ contains a vertex y such that there is some vertex $r \in R$ and vertex $z \in H(r)$ with yz the unique edge between $H(v_i)$ and $H(r)$

By Lemma B.8.6 we may assume that there is an end ϵ of Γ such that \mathcal{F} is strongly concentrated at ϵ . If there are infinitely many vertices dominating ϵ ,

then $\aleph_0 DR_k \preceq K_{\aleph_0} \leq \Gamma$ holds by Proposition B.2.4. So, we may assume that only finitely many vertices dominate ϵ . By taking a thick subtribe if necessary, we may assume that no member of \mathcal{F} contains such a vertex.

As before, if we can show that ϵ is pebbly, then we will be done by Corollary B.5.3. So suppose for a contradiction that ϵ is not r -pebble-win for some $r \in \mathbb{N}$.

We first claim that for each $H \in \mathcal{F}$ the pullback $R_H = H^\downarrow(R)$ of R in H is an ϵ -ray. Indeed, since \mathcal{F} is strongly concentrated at ϵ for every finite vertex set X of Γ , H intersects $C(X, \epsilon)$. As in Lemma B.8.8, since H is a connected graph and does not contain a vertex dominating ϵ in Γ , H must contain a ray $S \in \epsilon$. If S meets infinitely many branch sets then it must meet infinitely many branch sets of the form $H(x_i)$ for some x and hence, since R_H meets every $H(x_i)$, which are all connected subgraphs, we have that $R_H \sim S$ and so $R_H \in \epsilon$. Conversely, if S meets only finitely many branch sets then there must be some v_i such that $H(v_i)$ contains a tail of S . By our assumption on $H(v_i)$, for any tail of S the component of $H(v_i)$ containing that tail meets some edge between $H(v_i)$ and some $H(x_j)$. In this case it is also easy to see that $S \sim R_H$, and so $R_H \in \epsilon$.

For each $H \in \bigcup \mathcal{F}$ and each $i \in [k]$ we have that $H(v_i)$ is a connected subgraph of Γ . Let U be the set of all vertices in $H(v_i)$ which are the endpoint of some edge in H between $H(v_i)$ and $H(w)$ with $w \in R$. Since v_i dominates R , U is infinite, and so by Proposition B.10.3, $H(v_i)$ either contains a comb with all teeth in U or a subdivision of an infinite star with all leaves in U . However in the latter case the centre of the star would dominate ϵ , and so each $H(v_i)$ contains such a comb, whose spine we denote by $R_{H,i}$. Now we set $\mathcal{R}_H = (R_{H,1}, R_{H,2}, \dots, R_{H,k}, R_H)$.

Since $R_{H,i}$ is the spine of a comb, all of whose leaves are in U , it follows that in the graph $\text{RG}_H(\mathcal{R}_H)$ each $R_{H,i}$ is adjacent to R_H . Hence $\text{RG}_H(\mathcal{R}_H)$ contains a vertex of degree $k \geq 3$.

There is some layer $F \in \mathcal{F}$ of size $\ell \geq r + 1$, say $F = (H_i : i \in [\ell])$. For every $i \in [r + 1]$ we set $\mathcal{R}_{H_i} = (R_{H_i,1}, R_{H_i,2}, \dots, R_{H_i,k}, R_{H_i})$. Let us now consider the family of disjoint rays

$$\mathcal{R} = \bigcup_{i=1}^{r+1} \mathcal{R}_{H_i}.$$

By construction \mathcal{R} is a family of disjoint rays which tend to ϵ in Γ and by Lemmas B.3.3 and B.3.4, $\text{RG}_\Gamma(\mathcal{R})$ contains $r + 1$ vertices whose degree is at least $k \geq 3$. However, by Corollary B.5.4, there is a vertex set X of size at most r

such that $\text{RG}_\Gamma(\mathcal{R}) - X$ is a bare path P . But then some vertex whose degree is at least 3 is contained in the bare path, a contradiction. \square

C. Ubiquity of locally finite graphs with extensive tree-decompositions

C.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 [49, 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 [48]. However, not all graphs are \subseteq -ubiquitous, and in fact even trees can fail to be \subseteq -ubiquitous (see for example [112]).

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 [6, 69]), where \leq is the topological minor relation, it was shown by Andreae that all rayless countable graphs [4] and all locally finite trees [5] are \leq -ubiquitous. The latter result was recently extended to the class of all trees by the present authors [13].

In [3] 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 [65, 85, 103]. 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 [2] 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 C.1.1 (Andreae, [2, 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 C.1.2 (Andreae, [2, 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.* The main result of this paper is a far-reaching extension of Andreae's results, removing the assumption of finite blocks.

Theorem C.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 C.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 C.1.5 (Extensive tree-decompositions and ubiquity). *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 C.4 up to Theorem C.4.6. Roughly, however, it implies that we can find many self-minors of the graph at spots whose precise

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

positions are governed by the decomposition tree. We hope that the proof sketch in Section C.2 is a good source for additional intuition before the reader delves into the technical details.

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

This analysis occupies Section C.4 and C.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 [102] 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 C.1.3 and C.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 C.9.

But first, in Section C.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.

C.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 C.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 C.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 [13] 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 $\aleph_0 G \not\preceq \Gamma$.

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 $(R_{m,i}^n : i \leq k)$ 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 C.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 $(R_{m,i}^n : i \leq k)$ of rays starting at the various H_m^n as the *bundles* of rays from those H_m^n .



Figure C.1.: Stage n of the construction with disjoint $G[\bigcup_{m \leq n} V_m]$ -minors H_i^n with their bundles of disjoint rays.

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:

Definition C.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.[†] 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 .

[†]Where we use the notation as in [24], see also Definition C.3.3.

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 [13]:

Lemma C.2.2 (Weak linking lemma [13, Lemma 4.3]). *Let Γ be a graph and let $\omega \in \Omega(\Gamma)$. 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 [13, Lemma 4.4] or Lemma C.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:

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 C.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 C.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 [14] we analysed the possibilities for how the ray graphs and transition functions associated to a given thick[‡] 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 [14], 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

[‡]An end is *thick* if it contains infinitely many disjoint rays.

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 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 C.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.

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

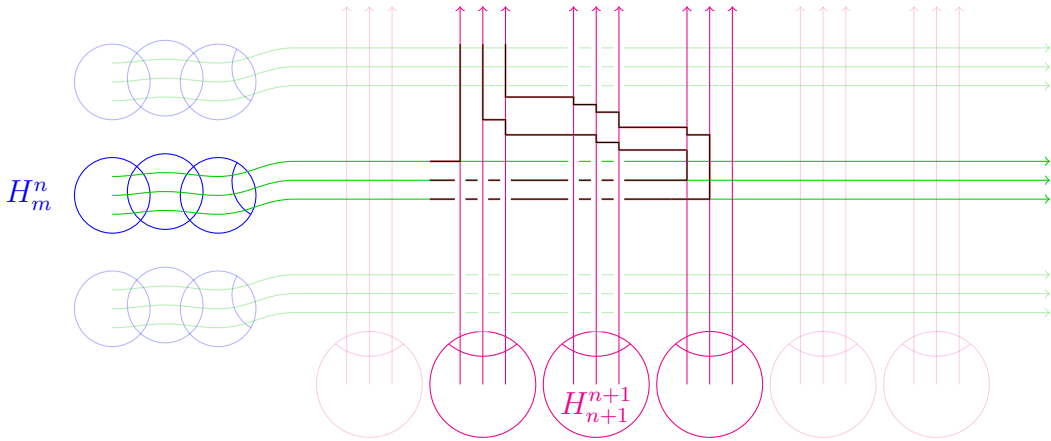


Figure C.2.: Extending the H_m^n by routing onto a set of disjoint G -minors might cause problems with introducing a new H_{n+1}^{n+1} disjoint to the rest.

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 C.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 C.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

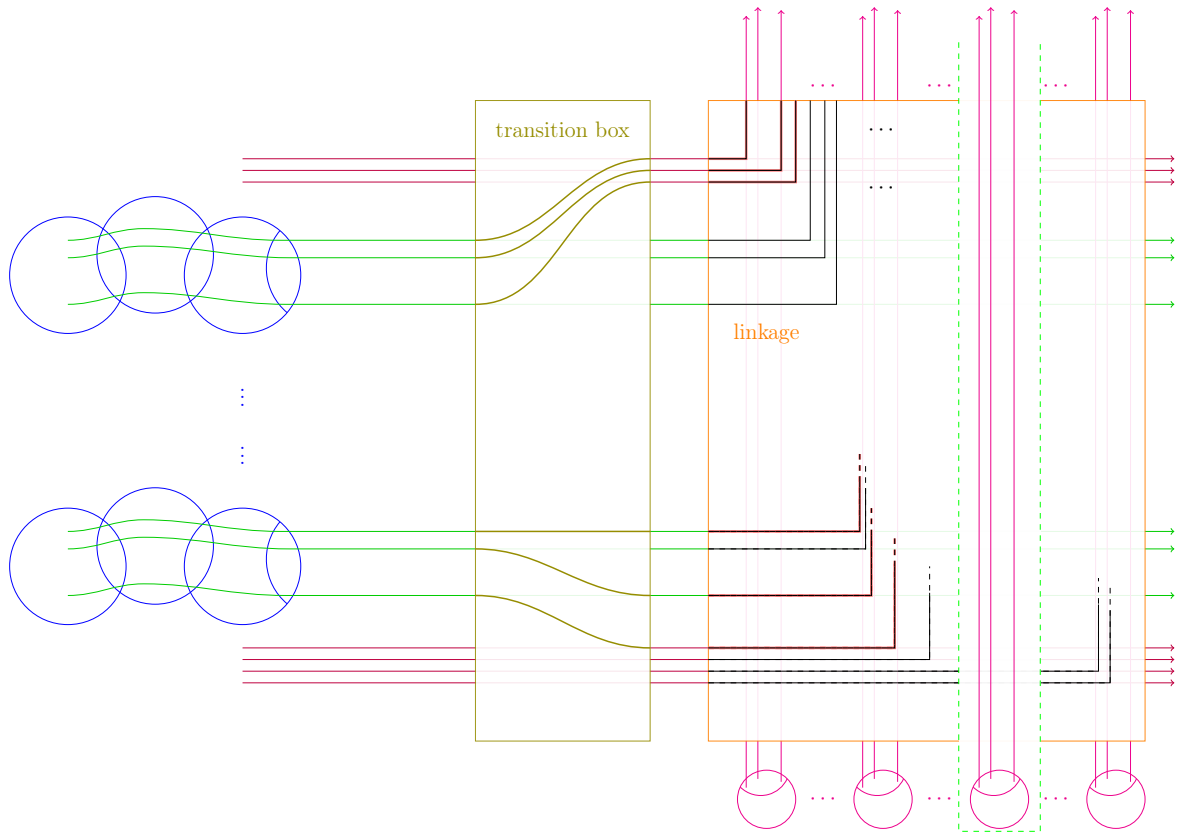


Figure C.3.: The transitioning strategy between the old and new bundles.

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 C.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 [13], where they appear in slightly simpler circumstances.

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

C.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 [24]. 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 $0 < i < n$.

C.3.1. Rays and ends

Definition C.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 C.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 C.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 C.3.4 (Ends of a graph, cf. [24, 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 $|\mathcal{R}| = \deg(\omega)$, as proved by Halin [49, Satz 1]. An end with finite degree is called *thin*, otherwise the end is called *thick*.

C.3.2. Inflated copies of graphs

Definition C.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 C.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 C.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 $v \in V(G)$. Note 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 C.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 C.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 C.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 C.3.11 ([14, 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 C.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 C.3.11) which are equivalent in H , and hence also equivalent in Γ .

Definition C.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 $R \in \omega$.

C.3.3. Transitional linkages and the strong linking lemma

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

Definition C.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 C.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 C.2.2, we will also need the following stronger linking lemma, which is a slight modification of [13, Lemma 4.4]:

Lemma C.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 C.3.15. We apply the strong linking lemma established in [13, Lemma 4.4] to the set $X \cup Y$ to obtain this version of the strong linking lemma. \square

Lemma and Definition C.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 C.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 a transitional linkage from \mathcal{R}_2 to \mathcal{R}_3 after $V(\bigcup \mathcal{P}_2)$. 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 ;[§]
- (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 C.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 C.3.20 ([14, 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*

C.3.4. Separations and tree-decompositions of graphs

Definition C.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:

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

Definition C.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 $T - T_x$.

[§]Formally, it is only the subset of \mathcal{P}_2 starting at the endpoints of \mathcal{P}_1 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.

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

A reader unfamiliar with tree-decompositions may also consult [24, Chapter 12.3].

Definition C.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 C.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 .

C.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 [5, Lemma 2].

C.4.1. Extensive tree-decompositions

Definition C.4.1 (Separations induced by tree-decompositions). Given 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. [24, 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 C.3.21.)

Definition C.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 C.4.3 (Self-similar bough). Let (T, \mathcal{V}) be a tree-decomposition of a graph G . Given $e \in E(T)$, 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 C.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 C.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 C.4.6. *Every locally finite connected graph admitting an extensive tree-decomposition is \preceq -ubiquitous.*

C.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 C.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 $\overline{G[B]}$, 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 practise (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 C.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 a subgraph $W \subseteq G[B(e')]$, which is an inflated copy 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 C.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 C.4.3 such that $W \subseteq G[B(e')]$. By Definition C.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 C.4.3.

Since $G[A(e)]$ is an $IG[A(e)]$ and W is an inflated copy of $\overline{G[B(e)]}$, by Definitions C.4.7 and C.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 C.8.

C.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 C.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 [89], proved over a series of 20 papers, shows that finite graphs are well-quasi-ordered under the minor relation. Thomas [102] 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 C.5.3 and Lemma C.5.11, to show that our two classes of graphs admit extensive tree-decompositions.

In Section C.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.

C.5.1. Finitely many thin ends

We will consider the following strengthening of the minor relation.

Definition C.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 C.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 [88, 1.7]. \square

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

Lemma C.5.4 ([49, 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 C.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 C.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 C.5.4 where $R' = t'_0 t'_1 \dots$

By Remark C.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 $G_i := G[W'_i]$ and $\pi_i : [\ell] \rightarrow V(G_i)$ where

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

By Lemma C.5.3 and Remark C.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 now 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 C.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 C.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 graphs H_{m_j} are pointed minors of $G_{n_0+m_j}$, we can deduce 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 C.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 C.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

C.5.2. Finite tree-width

Definition C.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 C.5.9. Kříž and Thomas [67] 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 C.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 C.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 [102]. \square

Lemma C.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 C.5.9, G has a lean tree-decomposition of width at most m and so, by Lemma C.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 many 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 C.5.11 and Remark C.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 $(G[B(t_{i_{k-1}}^\epsilon t_{i_k}^\epsilon)], \pi_{i_k}^\epsilon)_{i \in I_\epsilon}$, 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 C.5.6. \square

Remark C.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.

C.5.3. Sporadic examples

We note that, whilst Lemmas C.5.7 and C.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 C.5.14. The grid $\mathbb{Z} \times \mathbb{Z}$ has an extensive tree-decomposition, which can be seen in Figure C.4. 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 [14].




Figure C.4.: In the grid the boughs are self-similar.

In fact, we do not know of any construction of a locally finite connected graph which does not admit an extensive tree-decomposition.

Question C.5.15. Do all locally finite connected graphs admit an extensive tree-decomposition?

C.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 [14], 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 C.6.1 ([14, 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.*

The precise technical definition of such ends is not relevant, in what follows we will simply need to use the following results from [14].

Corollary C.6.2 ([14, Corollary 5.3]). *Let Γ be a graph with a pebbly end ϵ and let G be a countable graph. Then $\aleph_0 G \preccurlyeq \Gamma$.*

Lemma C.6.3 ([14, 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 C.6.4 ([14, 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 C.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 C.6.3 and C.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 C.8.2, we will shortly sketch the differences for the grid-like case.

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

By Lemma C.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 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 C.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 C.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 C.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 C.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 C.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 C.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 C.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 C.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 C.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 C.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 C.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 C.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 C.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 C.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}'' .

Indeed, let us take a finite separator X which separates ϵ' and ϵ'' in $\Gamma - V(R_c)$. By Lemma C.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 \mathcal{S}' in $\Gamma - V(R_c)$, and similarly there is a linkage after $X \cup Y$ from \mathcal{R}''_1 to \mathcal{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 \mathcal{S}' , and \mathcal{R}''_2 to \mathcal{S}'' . By Lemma C.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. Therefore, 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 C.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 C.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 C.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 C.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 C.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 C.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 C.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, we obtain that $T \in \top(\epsilon, S)$ as well and hence $S \leq_\epsilon T$. The other case is analogous. \square

Since, by C.6.14 (3), the order \leq_ϵ will agree with correct order on the central path, which is preserved by transition functions by Lemma C.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 C.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 C.6.14(2), 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 C.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 C.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 C.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 C.6.9, there are no other components of $\text{RG}(\overline{\mathcal{R}}) - i$. Finally, by the comment after Definition C.6.5, it follows that R_i is an inner vertex of the central path of this ray graph. \square

Lemma C.6.17 ([14, 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 C.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 C.6.19. *Let $\mathcal{R} = (R_i : i \in I)$ be a finite family of disjoint core ϵ -rays and let $\mathcal{S} = (S_j : j \in J)$ be a finite family of disjoint ϵ -rays. Let $\overline{\mathcal{R}} = (R_i : i \in \overline{I})$ be as in Lemma C.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 C.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 C.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 C.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 C.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 C.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 C.6.21. A sub-linkage of a transitional linkage is transitional.

The following remarks are an immediate consequence of the definitions and Lemma C.6.4.

Remark C.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 C.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 $\mathcal{R} \circ_{\mathcal{P}_1} \mathcal{S}$, then the concatenation $\mathcal{P}_1 + \mathcal{P}_2$ is preserving on \mathcal{R} .

Lemma C.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 C.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 C.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 C.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 C.6.17 and Lemma C.6.19 there must be two components of $\text{RG}(\mathcal{T}) \setminus \sigma(i)$ of size at least $N + 1$, a contradiction. \square

C.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 [13] and [14].

Definition C.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 -subtribe[¶] 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 [13].

Lemma C.7.2 (Removing a thin subtribe, [13, 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 C.7.3 (Pigeon hole principle for thick G -tribes, [13, 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 C.7.4 ([13, 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 C.7.5 ([13, 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} .*

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

(2) Every thick subtribe \mathcal{F}' of \mathcal{F} is concentrated at ϵ .

The following lemma from [14] shows that we can restrict ourself to thick G -tribes which are concentrated at thick ends.

Lemma C.7.6 ([14, 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 C.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 C.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 C.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 [14, Lemma 10.1], 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 \preceq_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 C.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 an 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 C.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* if for each $e \in E$, the set \mathcal{H} determines whether ω_e (cf. Definition C.4.4) converges to ϵ . 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 $\partial(S)$* 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 C.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 $H \in \mathcal{H}$.

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 w.r.t. ϵ .
- 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 $\leq_{\mathcal{F},e}$ for $\leq_{\bigcup \mathcal{F},e}$.

Remark C.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 C.7.9 ([14, 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 C.7.10 ([14, 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 C.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. For any $e \in \partial_\epsilon(S)$ apply Lemma C.7.10 to \mathcal{F} with $\omega_H = \omega_e$ for each $H \in \bigcup \mathcal{F}$. □

Lemma C.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 C.6.4. Then, by Remark C.6.14(4), 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 finite[‡]. 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 C.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 C.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 C.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 $C_\Gamma(X', H(\omega_e)) \neq C_\Gamma(X', \epsilon)$ 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

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

[‡]Note 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 C.7.11.

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 C.4.8 and Lemma C.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 C.4.8, $V(W_e) \subseteq V(H(G[B(e')])) \subseteq C_\Gamma(X', H(\omega_e))$. \square

Lemma C.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 C.7.12, we may even assume that \mathcal{F} strongly agrees about $\partial(\{r\})$. Using Lemma C.7.5(2), this \mathcal{F} would then satisfy (1) and (2). So, it remains to arrange for (3):

For every member H of \mathcal{F} and for every $e \in \partial_{-\epsilon}(\{r\})$, there exists a finite set $X_e \subseteq V(\Gamma)$ by Lemma C.7.13 such that for every finite vertex set $X' \supseteq X_e$ there is a push-out $H_e = H(G[A(e)]) \oplus_{H^\downarrow(\mathcal{R}_e)} W_e$ 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 C.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 C.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 C.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 $\{r\}$. However, $H'(\{r\}^{-\epsilon}) \setminus \bigcup_{e \in \partial_{-\epsilon}(\{r\})} W_e$ 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 C.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 C.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 $H(S^{-\epsilon}) \cap C_\Gamma(X', \epsilon) = \emptyset$. In case $\partial_{-\epsilon}(\tilde{S}) \setminus \partial(S) = \emptyset$, 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 C.7.13 to every $e \in \partial_{-\epsilon}(\tilde{S}) \setminus \partial(S)$, we obtain a finite set $X \supseteq X'$ and a family $(H_e : e \in \partial_{-\epsilon}(\tilde{S}) \setminus \partial(S))$ satisfying that 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 C.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 C.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 $\partial(\tilde{S})$, that is 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 C.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 C.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 C.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 $d(H) = d_H(H^{(1)})$ ($= d_H(H^{(i)})$ 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 C.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 C.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

C.8. The inductive argument

In this section we prove Theorem C.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 C.7.4, we may assume that \mathcal{F} is concentrated at some end ϵ of Γ , which, by Lemma C.7.6, we may assume is a thick end, and, by Lemma C.6.2, we may assume is not pebbly. Hence, by Theorem C.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 C.4.6 in

the case where ϵ is half-grid-like, and then in Section C.8.2 we will shortly 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.

C.8.1. The half-grid-like case

As explained in Section C.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 $H^\downarrow(R_{e,s})$ is an ϵ -ray.

Furthermore, since $\partial(T_n)$ is finite, we may assume by Lemma C.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 C.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 C.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 C.6.13. Furthermore, if \mathcal{F} strongly agrees about $\partial(S)$ then, as in Definition C.7.7, for each $e \in \partial_\epsilon(S)$, there is a linear order $\leq_{\mathcal{F},e}$ on $S(e)$.

Definition C.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 $\varrho := \lfloor n/2 \rfloor$ 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 $\partial(T_\varrho)$;
- is well-separated from ϵ at T_ϱ ;
- whenever $k < l \leq \varrho$, 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 $G(T_\varrho^{-\epsilon})$ (where $T_\varrho^{-\epsilon}$ 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 $\partial_\epsilon(T_\varrho) \neq \emptyset$, we require that there is a set \mathcal{J}_ϱ of disjoint core ϵ -rays disjoint to \mathcal{E}_n , with $|\mathcal{J}_\varrho| \geq (|\partial_\epsilon(T_\varrho)| + 1) \cdot |\mathcal{E}_n|$.

Suppose we have inductively constructed \mathcal{Q}_n for all $n \in \mathbb{N}$. Let us define $H_i := \bigcup_{n \geq 2i-1} Q_i^n$. 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 C.7.14 there is a thick subtribe \mathcal{F}_0 of \mathcal{F} which satisfies (I1). 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 the graph 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 C.5)

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 C.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 C.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 C.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 C.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

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

Note that Q_{k+1}^{n+1} is an inflated copy of $G(T_k^{-\epsilon})$. Moreover, let $Q_i^{n+1} := Q_i^n$ for all $i \in [k]$ and $\mathcal{Q}_{n+1} := (Q_i^{n+1} : i \in [k+1])$, 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 $C(X_{n+1}, \epsilon) \cap Q_{k+1}^{n+1} = \emptyset$. This set X_{n+1} is a bounder for \mathcal{Q}_{n+1} .

Since \mathcal{P} is transitional, Lemma C.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 C.3.17, C.6.23 and C.6.19, and Remark C.6.22(1), 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 C.6.14(1). Since \mathcal{P}' is

preserving on \mathcal{A}' and \mathcal{P}'' is preserving on \mathcal{E}_n^- , Remark C.6.22(2) 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 (I1) is still satisfied by \mathcal{F}_k and T_k , and (I4) is vacuously satisfied.




Figure C.5.: Adding a new copy when $n = 2k$ is even.

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 (I1), (I2), (I3) and (I4).

Case 2: $f \in \partial_\epsilon(T_{k-1})$. (See Figure C.6)

By (I1) we can apply Lemma C.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 (I1) for the next step.

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

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

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 C.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 C.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 we get that $H(B(e)) \cap (X_n \cup Z \cup V(\bigcup \mathcal{P})) = \emptyset$ for every $H \in F$. Let $F' \in \mathcal{F}_k$ be as in Lemma C.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 C.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

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

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 $\bigcup Q_{n+1}$. 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 C.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 C.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 $H_i^{\downarrow} \in F' \in \mathcal{F}_k$ 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 C.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^{\downarrow} to which the tails of the rest of \mathcal{E}_{n+1}^- are disjoint. Furthermore, by choice of Z and Lemma C.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. \square




Figure C.6.: Extending the copies when $n = 2k - 1$ is odd.

C.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 C.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 C.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 $i \in [n]$ **.

Every large enough ray graph is a cycle, which has a correct orientation by Lemma C.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 C.6.1, one can show the following:

Lemma C.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 C.7.9).*

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

Definition C.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 C.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.

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

Again, by a similar argument as in Section C.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 C.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 C.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 $H^\perp(\mathcal{R}_e)(\leq_{\mathcal{F},e})$ agrees with C_ϵ on $H^\perp(\mathcal{R}_e)$ for all $H \in \bigcup \mathcal{F}$.

Using this definition, the G -tribe refinement lemma (Lemma C.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 C.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 $\overline{\mathcal{C}}$ 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 $\overline{\mathcal{E}_n^-}$. 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

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

for each $i \in [N + 3]$, 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.

C.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 C.5.7 and C.5.12. Lemma C.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 C.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 [83] 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 C.1.1 and C.1.2:

Conjecture C.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 [102] 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.

Chapter II.

Hamiltonicity of locally finite graphs

D. Forcing Hamiltonicity in locally finite graphs via forbidden induced subgraphs

I: nets and bulls

D.1. Introduction

The question whether certain graphs have a Hamilton cycle, i.e. a cycle through all vertices of the graph, and the search for necessary as well as sufficient conditions forcing Hamiltonicity is a prominent subject within graph theory. Most results in this area focus on finite graphs, also for the reason that it is not clear what a Hamilton cycle in an infinite graphs should be. Although two-way infinite paths might be the canonical choice for such objects, considering only them limits the class of potential Hamiltonian graphs immensely: only graphs with at most two *ends* have a chance of being Hamiltonian. The ends of a graph are the equivalence classes of *rays*, i.e. one-way infinite paths, under the relation of being inseparable by finitely many vertices.

Nevertheless, the study of Hamiltonicity has quite successfully been transferred to infinite graphs, especially for *locally finite* ones, i.e. graphs where each vertex has finite degree. For a locally finite connected graph G the Freudenthal compactification $|G|$ is considered. This is a topological space arising from G seen as a 1-complex by adding additional points ‘at infinity’. These additional points are precisely the ends of G and the corresponding topology is defined in such a way that each ray converges to the end it is contained in. For more on the space $|G|$ see [24, 25, 29]. Following the topological approach by Diestel and Kühn [27, 28], the notion of cycles of a graph G is extended to *circles* in $|G|$, which are homeomorphic images of the unit circle $S^1 \subseteq \mathbb{R}^2$ in $|G|$. This definition now allows a rather big variety of infinite cycles. We call a circle a *Hamilton circle* of G if it contains all vertices of G . Since Hamilton circles are closed subspaces of $|G|$, they also contain

all ends of G .

Several Hamiltonicity results have been extended to locally finite infinite graphs so far, although not always completely, but with additional requirements [18, 20, 38, 50, 53–56, 75–77, 86]. For finite graphs many sufficient conditions guaranteeing Hamiltonicity exist which make use of global assumptions such as for example degree conditions involving the total number of vertices. To locally finite infinite graphs, however, such conditions do not seem to be easily transferable. For this reason we focus on sufficient conditions forcing Hamiltonicity with a local character in this series of papers, namely ones in terms of forbidden induced subgraphs. The specific subgraphs we are focusing on in this first paper out of the series are the *claw*, the *net* and the *bull*, which are depicted in Figure D.1. Specifically for the bull, we shall refer to its vertices b_1, b_2 of degree 1 as the *horns* of the bull. In general, given two graphs G and H we shall call G a *H-free* graph if G does not contain an induced subgraph isomorphic to H . So far sufficient conditions for Hamiltonicity in terms of forbidden induced subgraph conditions have not been analysed very much in the context of infinite graphs, although some results on claw-free graphs with additional constraints exist [50, 55, 56].




Figure D.1.: The induced subgraphs considered in this paper.

Our main results in this paper are centered around the following theorem by Shepherd. In order to state it, we have to give one additional definition. A finite graph G is called *k-leaf-connected* if $|V(G)| > k \in \mathbb{N}$ and given any vertex set $S \subseteq V(G)$ with $|S| = k$, then G has a spanning tree whose set of leaves is precisely S .

Theorem D.1.1. [98, Thm. 2.9] *Let G be a finite graph. If G is claw-free and net-free, then*

- (1) G is connected implies G has a Hamilton path.
- (2) G is 2-connected implies G is Hamiltonian.
- (3) For $k \geq 2$, G is $(k + 1)$ -connected or $G = K_{k+1}$ if and only if G is k -leaf-connected.

Note that statement (1) and (2) of Theorem D.1.1 were already proven in [31]. Regarding statement (3) note that Shepherd did not include the case $G = K_{k+1}$ in the equivalence. This is actually a tiny mistake in the original proof by Shepherd, because K_{k+1} is k -leaf-connected but not $(k + 1)$ -connected as the usual definition, which was also used in Shepherd's paper, requires such a graph to have more than $k + 1$ vertices. However, apart from this exception the proof by Shepherd is correct.

We shall extend all three statements of this theorem to infinite locally finite graphs. Before that, we analyse the structure of infinite locally finite claw-free and net-free graphs, and give examples of such graphs in Section D.3. Especially, we shall prove that such graphs have at most two ends.

In contrast to this, we consider locally finite graphs with potentially up to 2^{\aleph_0} many ends in the second paper of this series [58], where we focus on the *paw*, i.e. the graph obtained by attaching an edge to a triangle, and a slightly relaxed forbidden induced subgraph condition.

Regarding the first two statements of Theorem D.1.1 we shall prove the following theorems.

Theorem D.1.2. *For an infinite locally finite connected graph G that is claw-free and net-free, precisely one of the following statements holds:*

- (1) G has only one end and admits a spanning ray.
- (2) G has only two ends and admits a spanning double ray.

Theorem D.1.3. *Every locally finite, 2-connected claw-free and net-free graph is Hamiltonian.*

For statement (3) of Theorem D.1.1 it might not entirely be clear at first sight how to phrase an extension of the theorem. Note that k -leaf-connectivity has to be replaced in the statement. To see this, observe that for finite graphs 2-leaf-connectivity coincides with *Hamilton connectivity*, i.e. the existence of a

Hamilton path with any two previously chosen vertices as endpoints. While an infinite path within an infinite graph G can never meet this condition, an arc within the Freudenthal compactification $|G|$ of G may. So we define a topological analogue, called *topological k -leaf-connectedness*, whose definition can be found in Section D.2.2. We obtain the following extension of statement (3) of Theorem D.1.1.

Theorem D.1.4. *Let G be a locally finite, connected, claw-free and net-free graph, and let $k \in \mathbb{N}$ satisfy $k \geq 2$. Then G is $(k + 1)$ -connected or $G = K_{k+1}$ if and only if G is topologically k -leaf-connected.*

The key in [98] to prove Theorem D.1.1 is the following structural characterisation of the involved graphs. In order to state this characterisation we have to give another definition first. A graph G with a vertex $v \in V(G)$ is called *distance-2-complete centered at v* if $G - v$ has exactly two components and in each component C and for each $i \in \mathbb{N}$, the vertices at distance i from v in $G[V(C) \cup \{v\}]$ induce a complete graph.

Theorem D.1.5. [98, Thm 2.1] *A finite connected graph G is claw-free and net-free if and only if for every minimal separator $S \subseteq V(G)$ and every $v \in S$, the graph $G - (S \setminus \{v\})$ is distance-2-complete centered at v .*

We extend Theorem D.1.5 to locally finite graphs via the following result.

Theorem D.1.6. *A locally finite connected graph G is claw-free and net-free if and only if for every minimal finite separator $S \subseteq V(G)$ and every v in S , the graph $G - (S \setminus \{v\})$ is distance-2-complete centered at v .*

Beside Theorem D.1.1 we shall also extend the following theorem by Ryjáček, which is about claw-free graphs where induced bulls may exist, but only under an additional assumption.

Theorem D.1.7. [91, main theorem] *Let G be a finite, 2-connected, claw-free graph. If for every induced bull B in G its horns have a common neighbour in G , then G is Hamiltonian.*

Our key to extend Theorem D.1.7 to locally finite graphs is the following structural result about the involved graphs.

Theorem D.1.8. *Let G be a locally finite connected claw-free graph such that the horns of every induced bull have a common neighbour. If G contains an induced bull, then G is a finite graph.*

While the class of claw-free and bull-free graphs is a proper subclass of the class of claw-free graphs where the horns of every induced bull have a common neighbour for finite graphs, they now coincide for infinite locally finite graphs. To see that these classes differ in finite graphs, just consider a bull itself with an additional vertex only adjacent to the two horns of the bull.

Now since bull-free graphs are especially net-free, we obtain the following result for locally finite graphs as an immediate corollary of Theorem D.1.8 together with Theorem D.1.7 and of Theorem D.1.3.

Corollary D.1.9. *Let G be a locally finite, 2-connected, claw-free graph such that the horns of every induced bull have a common neighbour. If G contains an induced bull, then G is finite and Hamiltonian. Otherwise, G is net-free, (potentially infinite) and Hamiltonian.*

The structure of this paper is as follows. In Section D.2 we will first introduce the needed definitions and notation. Furthermore, we shall state all tools which we shall use to prove our main results. In Section D.3 we shall analyse the structure of the graphs we consider in this paper and give examples of them. We also prove Theorem D.1.8 and Theorem D.1.6 in that section. Finally, we prove our main results regarding Hamiltonicity, i.e. Theorem D.1.2, Theorem D.1.3, Theorem D.1.4 in Section D.4.

D.2. Preliminaries

We will follow the graph theoretical notation and use basic facts without quoting them from [24], which includes especially the topological approach to locally finite graphs in [24, Ch. 8.5]. For a wider survey of topological infinite graph theory, see [25].

D.2.1. Basic notions

All graphs which are considered in this paper are undirected and simple. In general, we do not assume a graph to be finite. A graph is called *locally finite* if every

vertex has finite degree.

For the rest of this section let G denote some graph. Later in this section, however, we shall make further assumptions on G .

We shall use the symbol \mathbb{N} to denote the set of natural numbers, which is meant to include 0 within this paper.

Let X be a subset of $V(G)$. We denote by $G[X]$ the induced subgraph of G with vertex set X . For small vertex sets, we sometimes omit the set brackets, i.e. we write $G[a, b, c]$ as a short form for $G[\{a, b, c\}]$. We write $G - X$ for the graph $G[V \setminus X]$. If H is a subgraph of G we shall write $G - H$ instead of $G - V(H)$. Again we omit set brackets around small vertex sets, especially for singleton sets. We briefly denote the cut $E(X, V \setminus X)$ by $\delta(X)$. For any $i \in \mathbb{N}$ we denote by $N_i(X)$ and $N_i(v)$ the set of vertices of distance at most i in G from the vertex set X or from a vertex $v \in V(G)$.

Let C be a cycle of G and u be a vertex of C . We implicitly fix an orientation of the cycle and we write u^+ and u^- for the neighbour of u in C in positive and negative, respectively, direction of C using a fixed orientation of C . Later on we will not always mention that we fix an orientation for the considered cycle using this notation.

If v and w are vertices of a tree T , then we denote by vTw the unique v - w path in T .

For some $k \in \mathbb{N}$, we say that a finite graph G' is *k-leaf-connected* if $|V(G')| > k$ and given any vertex set $S \subseteq V(G')$ with $|S| = k$, then G' has a spanning tree T whose set of leaves is precisely S . Note that for any graph G being 2-leaf-connected is equivalent to being *Hamilton connected*, namely that any two distinct vertices v, w of G are connected via a Hamilton path with v and w as its endpoints.

For any $v \in V(G)$ we call G *distance-2-complete centered at v* if $G - v$ has exactly two components and in each component K and for each $i \in \mathbb{N}$, the vertices at distance i from v in $G[V(K) \cup \{v\}]$ induce a complete graph.

A one-way infinite path R in G is called a *ray* of G . A subgraph of a ray R is called a *tail* of R if it is itself a ray. The unique vertex of degree 1 of R is called the *start vertex* of R . For a vertex r on a ray R , we denote the tail of R with start vertex r by rR . A two-way infinite path in G is called a *double ray*.

An equivalence relation can be defined on the set of all rays of G by saying that two rays in G are *equivalent* if they cannot be separated by finitely many

vertices. It is easy to check that this defines in fact an equivalence relation. The corresponding equivalence classes of rays under this relation are called the *ends* of G . We denote the sets of ends of a graph G with $\Omega(G)$. If $R \in \omega$ for some end $\omega \in \Omega(G)$, then we briefly call R an ω -ray.

Note that for any end ω of G and any finite vertex set $S \subseteq V(G)$ there exists a unique component $C(S, \omega)$ that contains tails of all ω -rays. We say that a finite vertex set $S \subseteq V(G)$ *separates* two ends ω_1 and ω_2 of G if $C(S, \omega_1) \neq C(S, \omega_2)$. Note that any two different ends can be separated by a finite vertex set.

We say that a (double) ray of G is *geodetic* if and only if for any two vertices on the (double) ray there is no shorter path between these two vertices in G than the one on the (double) ray.

Let R be a ray in G and $X \subseteq V(G)$ be finite. We call R *distance increasing with respect to X* if $|V(R) \cap N_i(X)| = 1$ for every $i \in \mathbb{N}$, where $N_0(X) := X$. Note that a distance increasing ray with respect to X has its start vertex in X .

D.2.2. Topological notions

For the rest of this section, we assume G to be locally finite and connected. A topology can be defined on G together with its ends to obtain a topological space which we call $|G|$. Note that inside $|G|$, every ray of G converges to the end of G it is contained in. For a precise definition of $|G|$, see [24, Ch. 8.5]. Apart from the definition of $|G|$ as in [24, Ch. 8.5], there is an equivalent way of defining the topological space $|G|$, namely, by endowing G with the topology of a 1-complex and considering the Freudenthal compactification of G . This connection was examined in [29]. For the original paper of Freudenthal about the Freudenthal compactification, see [36].

For a point set X in $|G|$, we denote its closure in $|G|$ by \overline{X} and its interior by $\overset{\circ}{X}$. A subspace Z of $|G|$ is called *standard subspace* of $|G|$ if $Z = \overline{H}$ where H is a subgraph of G .

A *circle* of G is the image of a homeomorphism which maps from the unit circle $S^1 \subseteq \mathbb{R}^2$ to $|G|$. The graph G is called *Hamiltonian* if there exists a circle in $|G|$ which contains all vertices of G , and hence, by the closedness of circles, also all ends of G . This circle is called a *Hamilton circle* of G . We note that, for finite graphs, this coincides with the usual notion of Hamiltonicity.

The image of a homeomorphism which maps from the closed real unit interval $[0, 1]$ to $|G|$ is called an *arc* in $|G|$. For an arc α in $|G|$, we call the images of 0 and 1 of the homeomorphism defining the arc, the *endpoints* of the arc. A subspace Z of $|G|$ is called *arc-connected* if for every two points of Z there is an arc in Z which has these two points as its endpoints. Finally, an arc in $|G|$ is called a *Hamilton arc* of G if it contains all vertices of G .

Let ω be an end of G and Z be a standard subspace of $|G|$ containing ω . Then we define the *degree* of ω in Z as a value in $\mathbb{N} \cup \{\infty\}$, namely the supremum of the number of edge-disjoint arcs in Z that have ω as one of their endpoint.

We make a further definition with respect to end degrees which allows us to distinguish the parity of degrees of ends when they are infinite. This definition has been introduced by Bruhn and Stein [17]. We call the degree of an end ω of G in a standard subspace X of $|G|$ *even* if there is a finite set $S \subseteq V(G)$ such that for every finite set $S' \subseteq V(G)$ with $S \subseteq S'$ the maximum number of edge-disjoint arcs in X with ω and some $s \in S'$ as endpoints is even. Otherwise, we call the degree of ω in X *odd*.

A *topological tree* of G is a connected standard subspace of $|G|$ which contains no circle of G . A topological tree of G is called *spanning* if it contains all vertices of G . We denote such a tree also as a *topological spanning tree* of G . Let T be a subgraph of G such that \overline{T} is a topological tree of G . We call a point $x \in \overline{T}$ a *leaf* of \overline{T} if either $x \in V(G)$ and has degree 1 in T or $x \in \Omega(G)$ and has degree 1 in \overline{T} .

We extend the notion of k -leaf-connectedness to locally finite connected graphs as follows. We call G *topologically k -leaf-connected* if $|V(G)| > k$ and given any set $S \subseteq V(G) \cup \Omega(G)$ with $|S| = k$, then G has a topological spanning tree \overline{T} whose set of leaves is precisely S . Similarly as for finite graphs, being topologically 2-leaf-connected coincides with the notion for locally finite connected graphs G of being *Hamilton connected*, i.e. for any two distinct $x, y \in V(G) \cup \Omega(G)$ there exists a Hamilton arc of G that has x and y as its endpoints.

D.2.3. Tools

In this section we introduce some basic lemmas we will use to prove our results. We begin with stating a lemma which allows us to make slightly limited, but still very helpful compactness arguments.

Lemma D.2.1. [24, Lemma 8.1.2 (König’s Infinity Lemma)] *Let $(V_i)_{i \in \mathbb{N}}$ be a sequence of disjoint non-empty finite sets, and let G be a graph on their union. Assume that for every $n > 0$ each vertex in V_n has a neighbour in V_{n-1} . Then G contains a ray $v_0v_1 \dots$ with $v_n \in V_n$ for all $n \in \mathbb{N}$.*

An immediate consequence of Lemma D.2.1 is the following proposition.

Proposition D.2.2. [24, Prop. 8.2.1] *Every infinite connected graph has a vertex of infinite degree or contains a ray.*

Since we shall only consider locally finite connected graphs, we know by Proposition D.2.2 that such graphs contain a ray as soon as they are infinite.

The next lemma is also an immediate consequence of Lemma D.2.1 and ensures the existence of distance increasing rays with respect to finite vertex sets in locally finite graphs. For the sake of completeness we give a proof here.

Lemma D.2.3. *Let G be an infinite locally finite connected graph and $X \subseteq V(G)$ be finite. Then there exists a distance increasing ray with respect to X .*

Proof. Since G is locally finite and connected, each $N_i(X)$ is non-empty, but finite. Also each vertex in $N_{i+1}(X)$ has a neighbour in $N_i(X)$ by definition for every $i \in \mathbb{N}$. By Lemma D.2.1 we obtain the desired ray. \square

We state a similar lemma about the existence of geodesic double rays.

Lemma D.2.4. [108, Thm. 2.2] *Let G be a locally finite connected graph and ω_1 and ω_2 two distinct ends of G . Then there is a geodesic double ray that is the union of an ω_1 -ray and an ω_2 -ray.*

The next lemma tells us that arcs within $|G|$ have to cross a finite cut as soon as they meet both sides of the cut.

Lemma D.2.5. [24, Lemma 8.5.3 (Jumping Arc Lemma)] *Let G be a locally finite connected graph and $F \subseteq E(G)$ be a cut with the sides V_1 and V_2 .*

(1) *If F is finite, then $\overline{V_1} \cap \overline{V_2} = \emptyset$, and there is no arc in $|G| \setminus \overset{\circ}{F}$ with one endpoint in V_1 and the other in V_2 .*

(2) *If F is infinite, then $\overline{V_1} \cap \overline{V_2} \neq \emptyset$, and there may be such an arc.*

The following lemma gives us a combinatorial criterion when standard subspaces of $|G|$ are topologically connected.

Lemma D.2.6. [24, Lemma 8.5.5] *If a standard subspace X of $|G|$ contains an edge from every finite cut of G whose both sides meet X , then X is topologically connected.*

Although topological connectedness and arc-connectedness differ for general topological spaces, they do not for closed subspaces of $|G|$ as shown by the following lemma.

Lemma D.2.7. [26, Thm. 2.6] *If G is a locally finite connected graph, then every closed topologically connected subset of $|G|$ is arc-connected.*

We shall make use of Lemma D.2.6 and Lemma D.2.7 to verify that ends in standard subspaces we construct have degree at least 1 by showing that those spaces intersect every finite cut. Similarly, the following theorem gives us a way to verify even degrees at ends.

Theorem D.2.8. [25, Thm. 2.5] *Let G be a locally finite connected graph. Then the following are equivalent for $D \subseteq E(G)$:*

- (1) *D meets every finite cut in an even number of edges.*
- (2) *Every vertex and every end of G has even degree in \overline{D} .*

The following lemma, combined with Lemma D.2.6 gives us a nearly purely combinatorial characterisation of those standard subspaces of $|G|$ which form a circle.

Lemma D.2.9. [17, Prop. 3] *Let C be a subgraph of a locally finite connected graph G . Then \overline{C} is a circle if and only if \overline{C} is topologically connected and every $v \in V(C)$ has degree 2 in C as well as every $\omega \in \Omega(G) \cap \overline{C}$ has degree 2 in \overline{C} .*

Our general strategy to verify that the closure \overline{H} within $|G|$ of a subgraph H of G , which we usually construct in countably many steps, is in fact a Hamilton circle of G works as follows. First we check that H contains every vertex of G . Then we prove that each vertex has degree 2 in H , which is usually an easy task. Now we prove that H intersects every finite cut of G , but in even number of edges. This

already proves that \overline{H} is topologically connected by Lemma D.2.6 and that every end of G has even degree, but least 2 in \overline{H} by Lemma D.2.7 and Theorem D.2.8. By Lemma D.2.9 it only remains to bound the degrees of the ends, for which we use Lemma D.2.5 adjusted to the way we construct H .

D.3. About the structure of the graphs considered in this paper

In this section we shall analyse the structure of the graphs that occur in the main results of this paper. Furthermore, we shall give examples of the considered graphs at the end of this section.

Let us now start with a very easy observation about claw-free graphs. The result is probably folklore and we do not give a proof here. However, in case a proof is desired, one is given in [56, Prop. 3.7.].

Proposition D.3.1. *Let G be a connected claw-free graph and S be a minimal vertex separator in G . Then $G - S$ has exactly two components.*

Next we generalise Theorem D.1.5 to locally finite graphs and obtain a structural characterisation of locally finite claw-free and net-free graphs. Note that Theorem D.1.5 was essential for the proof of Theorem D.1.1 and its generalisation will also be crucial for us to extend Theorem D.1.1 to locally finite graphs. We recall Theorem D.1.5 below. Recall that we call a graph G with a vertex $v \in V(G)$ *distance-2-complete centered at v* if $G - v$ has exactly two components and in each component C and for each $i \in \mathbb{N}$, the vertices at distance i from v in $G[V(C) \cup \{v\}]$ induce a complete graph.

Theorem D.1.5. [98, Thm 2.1] *A finite connected graph G is claw-free and net-free if and only if for every minimal separator $S \subseteq V(G)$ and every $v \in S$, the graph $G - (S \setminus \{v\})$ is distance-2-complete centered at v .*

Now we extend Theorem D.1.5 to locally finite graphs by proving Theorem D.1.6, which is recalled below. We would like to point out that for one implication we can actually use the same proof which was given in [98] for the corresponding implication for finite graphs of Theorem D.1.5. For the sake of completeness, we include this argument here. Let us now recall the theorem we are proving.

Theorem D.1.6. *A locally finite connected graph G is claw-free and net-free if and only if for every minimal finite separator $S \subseteq V(G)$ and every v in S , the graph $G - (S \setminus \{v\})$ is distance-2-complete centered at v .*

Proof. Suppose first for a contradiction that there exists a locally finite connected graph G for which every minimal finite separator $S \subseteq V(G)$ and every $v \in S$, the graph $G - (S \setminus \{v\})$ is distance-2-complete centered at v , but G contains a claw or a net as an induced subgraph H . Let h_1, h_2, h_3 denote the three vertices of degree 1 in H and let us call any other vertex of H a *central vertex* of H . By the local finiteness of G we know that $S_1 := N(h_1)$ is a finite vertex set containing a central vertex c_1 of $H^1 := H$ such that $L_1 := \{h_1, h_2, h_3\}$ intersects more than one component of $G - S_1$.

Next we shall recursively define sets S_i, H^i, L_i and vertices c_i until S_i is a \subseteq -minimal vertex separator of G such that the following holds for every $i \in \mathbb{N}$ with $i \geq 1$:

- (1) H^i is either an induced claw or an induced net of G and $H^i \subseteq G - (S_i \setminus \{c_i\})$.
- (2) L_i consists of the vertices of degree 1 in H^i .
- (3) S_i contains a central vertex c_i of H^i .
- (4) $S_{i+1} \subsetneq S_i$ if S_i is not a \subseteq -minimal vertex separator of G .
- (5) S_i separates G such that L_i is not contained in a single component of $G - S_i$.

We already found suitable sets and the vertex c_i for $i = 1$ above. Now suppose S_i, c_i, H^i and L_i have already been defined for some $i \geq 1$. If there exists a vertex $s \in S_i$ such that $S_i \setminus \{s\}$ still satisfies property (5), then set $S_{i+1} := S_i \setminus \{s\}$, the graph $H^{i+1} := H^i$, the vertex $c_{i+1} := c_i$ and $L_{i+1} := L_i$. In this way all properties (1)-(5) are still maintained.

So suppose there does not exist such a vertex $s \in S_i$ and S_i is not a \subseteq -minimal vertex separator of G . Let K_1, \dots, K_k be the components of $G - S_i$ for some $k \in \mathbb{N}$. Since S_i is not a \subseteq -minimal vertex separator of G , there exists a vertex $x \in S_i$ such that $S_i \setminus \{x\}$ is still a vertex separator of G . Hence, x is not adjacent to any vertex of K_j , for some $j \in \mathbb{N}$ with $j \leq k$, say $j = 1$. Since $S_i \setminus \{x\}$ does not satisfy property (5) with respect to L_i , there exist two distinct components of $G - S_i$ which are both different from K_1 and contain vertices from L_i , say these

are K_2 and K_3 . Furthermore, x must be adjacent to vertices in K_2 and K_3 . Now pick a $v \in S_i$ which is adjacent to some vertex x_1 in K_1 . Because $S_i \setminus \{v\}$ does not satisfy property (5) either, we know that v is adjacent to vertices $x_2 \in V(K_2)$ and $x_3 \in V(K_3)$ as well. Therefore, $H^{i+1} := G[v, x_1, x_2, x_3]$ is an induced claw of G . Set L^{i+1} to be the vertices of degree 1 in H^{i+1} , the vertex $c_{i+1} := v$ and $S_{i+1} := S_i \setminus \{x\}$. Now $S_{i+1}, c_{i+1}, H^{i+1}$ and L_{i+1} satisfy properties (1)-(5), which completes the recursive definition.

Throughout the whole article we shall stick to the following conventions. If $a, b, c, d \in V(G)$ induce a claw in G , then by using the notation $G[a, b, c, d]$ we shall always put the vertex of degree 3 in the claw first within the brackets, i.e. a has degree 3 within the claw $G[a, b, c, d]$. If $a, b, c, d, e \in V(G)$ induce a bull in G , then by using the notation $G[a, b, c, d, e]$ we shall always put the vertex of degree 2 in the bull first within the brackets, followed by the two vertices of degree 3 in the bull, followed by their respective neighbouring vertices of degree 1. Finally, if $a, b, c, d, e, f \in V(G)$ induce a net in G , then by using the notation $G[a, b, c, d, e, f]$ we shall always put the vertices of degree 3 in the net first within the brackets, followed by their respective neighbouring vertices of degree 1.

Let $\ell \in \mathbb{N}$ such that S_ℓ is an \subseteq -minimal vertex separator of G . By the properties (1) and (3) we know that S_ℓ contains a central vertex c_ℓ of H^ℓ such that $H^\ell \subseteq G - (S_\ell \setminus \{c_\ell\})$. In both cases, whether H^ℓ is an induced claw or an induced net of G , this contradicts that $G - (S_\ell \setminus \{c_\ell\})$ is distance-2-complete centered at c_ℓ , which holds by assumption.

Now let us prove the converse. Assume that G is a locally finite connected claw- and net-free graph. Take an arbitrary finite minimal separator S of G and fix some $v \in S$. Due to Proposition D.3.1, we get that $G - S$ has exactly two components C_1 and C_2 . For each $i \in \{1, 2\}$ fix a finite connected subgraph $F_i \subseteq C_i$ which contains $N(S) \cap V(C_i)$. Since S is a minimal separator, there exists a shortest v - f path P in $G - (S \setminus \{v\}) \cap C_i$ for every $f \in F_i$ and every $i \in \{1, 2\}$. Let $n \in \mathbb{N}$ be the maximum length of all such shortest paths P for all $f \in F_i$ and every $i \in \{1, 2\}$. Now set

$$G_0 := G[S \cup \bigcup_{i=0}^n N_i(v)].$$

We claim that G_0 fulfills the antecedent of Theorem D.1.5. Clearly G_0 is connected and finite, as G is locally finite. Furthermore, since G_0 is an induced

subgraph of the claw-free and net-free graph G , we know that G_0 is claw-free and net-free as well. By definition, $N(S) \cup S$ is contained in G_0 . Hence S is again a finite minimal vertex separator in G_0 . So Theorem D.1.5 implies that G_0 is distance-2-complete centered at v .

Now suppose we have already defined G_i for some $i \in \mathbb{N}$. Then we recursively define $G_{i+1} := G[V(G_i) \cup N(V(G_i))]$. By definition, G_{i+1} is a finite connected graph without induced claws and without induced nets, and S is a minimal separator in G_{i+1} . So again by Theorem D.1.5 we know that G_{i+1} is distance-2-complete centered at v as well. But now we get that for every $k \in \mathbb{N}$ each distance class $N_k(v)$ induces a clique in each component of $G - S$ as witnessed by G_k . \square

Although we shall not need it in order to prove our main results, we would like to add the following structural result for locally finite claw-free and bull-free graphs. It tells us that finite minimal separators in such graphs always induce cliques. Note that such separators always exist, since G is locally finite. Since bull-free graphs must be net-free as well, the following result combined with Theorem D.1.6 give us very much information about the structure of infinite locally finite claw-free and bull-free graphs.

Lemma D.3.2. *In an infinite locally finite claw-free and bull-free graph every finite \subseteq -minimal vertex separator induces a clique.*

Proof. Let G be a graph as in the statement of the lemma and let S be a finite \subseteq -minimal vertex separator of G . Suppose for a contradiction that S contains two distinct vertices u, v such that $uv \notin E(G)$. Since G is claw-free, we know by Proposition D.3.1 that $G - S$ has precisely two components, call them C and C' . Furthermore, we know by Proposition D.2.2 that G contains a ray. So G has an end ω . Since S is finite, every ω -ray has a tail in either C or C' , say in C . Let $R = r_0 r_1 \dots$ be a distance increasing ray with respect to $\{u\}$ in $G[V(C) \cup \{u\}]$. Hence, $r_0 = u$. Such a ray exists due to Lemma D.2.3. Furthermore, let P be a $v - R$ path in $G[V(C) \cup \{u, v\}]$ that is shortest possible and, among all such paths, has the additional property that its end vertex on R lies as close to u on R as possible. Let $i \in \mathbb{N}$ such that r_i is the endvertex of P on R and let p denote the neighbour of r_i on P . In case p is an inner vertex of P , we denote by p^- the neighbour of p on P that is different from r_i . Otherwise, $p = v$ and we denote by

p^- some arbitrary neighbour of v in C' , which exists since S is minimal. Finally, let r_{-1} denote some arbitrary neighbour of u in C' .

Now we claim that $pr_{i-1} \notin E(G)$. For $r_i \neq u$ this follows immediately from the definition of P . In case $r_i = u$ and $p \in V(C)$, we know that S separates r_{-1} and p . So $pr_{-1} \notin E(G)$ holds as well. Note that the case $r_i = u$ and $p = v$ does not occur as $uv \notin E(G)$ by assumption. Furthermore, we know that $r_{i-1}r_{i+1} \notin E(G)$ as R is distance increasing with respect to $\{u\}$. Next consider the graph $G[r_i, r_{i-1}, r_{i+1}, p]$. Since G is claw-free and $r_{i-1}r_{i+1}, pr_{i-1} \notin E(G)$, we now know that $pr_{i+1} \in E(G)$.

We continue by proving $p^-r_j \notin E(G)$ for all $j \in \{i, i+1, i+2\}$. For $p = v$, this follows since $p^- \in V(C')$ and $r_i, r_{i+1}, r_{i+2} \in V(C)$ as $r_i \neq u$, and so S separates p^- from r_i, r_{i+1} and r_{i+2} . In case $p \neq v$, the definition of P ensures that p^- is not a neighbour of any vertex on R . Again using that R is distance increasing with respect to $\{u\}$, we know that $r_i r_{i+2} \notin E(G)$. By considering $G[r_i, p, r_{i+1}, p^-, r_{i+2}]$ and using that G is bull-free, we get that $pr_{i+2} \in E(G)$ must hold.

Now we derived a contradiction since we already argued that $p^-r_i, p^-r_{i+2}, r_i r_{i+2} \notin E(G)$, and so $G[p, p^-, r_i, r_{i+2}]$ is an induced claw. \square

Next let us prove Theorem D.1.8. To ease the readability of the paper, let us recall the theorem here.

Theorem D.1.8. *Let G be a locally finite connected claw-free graph such that the horns of every induced bull have a common neighbour. If G contains an induced bull, then G is a finite graph.*

Proof. Suppose for a contradiction that G is infinite and contains an induced bull B . Let b_1 and b_2 denote the horns of B , and let z denote the vertex of degree 2 in B (cf. Figure D.1). Furthermore, let a_1 and a_2 denote the other two vertices of B such that $a_i b_i \in E(B)$ for every $i \in \{1, 2\}$. Since the horns of B have a common neighbour in G , let us fix such a common neighbour c of b_1 and b_2 , which then clearly has to lie in $G - B$. As G is infinite, locally finite and connected, there exists some distance increasing ray $R = r_0 r_1 \dots$ with respect to $V(B) \cup \{c\}$ in G . We shall distinguish four possible cases according to where R might start, and derive contradictions for each case.

Case 1. $r_0 = z$.

For this case note first that $r_1b_i \notin E(G)$ for every $i \in \{1, 2\}$ as otherwise $G[r_1, z, r_2, b_i]$ would be an induced claw. Next let us verify that $r_1c \notin E(G)$. Suppose for a contradiction the edge r_1c exists. Then $G[c, b_1, b_2, r_1]$ is an induced claw as $b_1b_2 \notin E(G)$ and $r_1b_i \notin E(G)$ by the argument above; a contradiction. In particular, this implies $c \neq r_1$.

Furthermore, we can assume without loss of generality that $G[r_i, a_1, a_2, b_1, b_2]$ is not an induced bull for every $i \geq 1$. To see this note that $r_i \notin N(V(B))$ for every $i \geq 2$ as R is distance increasing with respect to $V(B) \cup \{c\}$. So if $B' = G[r_1, a_1, a_2, b_1, b_2]$ is an induced bull, then $r_1r_2 \dots$ is a distance increasing ray with respect to $V(B') \cup \{c\}$ which starts at the vertex of degree 2 of B' as well. Then we would consider B' instead of B . By the previous argument we know that not both of the edges r_1a_1 and r_1a_2 can exist. Now suppose for a contradiction that exactly one of these edges exists, say r_1a_1 . Then $G[a_1, r_1, a_2, b_1]$ is an induced claw, since $a_2b_1 \notin E(G)$ as B is an induced bull, $r_1b_1 \notin E(G)$ by the argument above and $r_1a_2 \notin E(G)$ by assumption. Since the analysis for the edge r_1a_2 works analogously, we know that $r_1a_1, r_1a_2 \notin E(G)$.

Now we can conclude that $B'' = G[a_2, a_1, z, b_1, r_1]$ is an induced bull with horns r_1 and b_1 . So there exists some $c' \in V(G - B'')$ which is a common neighbour of r_1 and b_1 . As r_1 is neither adjacent to c nor to b_2 , we know that $c' \neq c, b_2$ and since R is distance increasing with respect to $V(B) \cup \{c\}$ we get $c' \neq r_i$ for all $i \in \mathbb{N}$. Because $G[r_1, r_2, z, c']$ is not an induced claw and $zr_2 \notin E(G)$, there are two options how this can be avoided:

$$c'r_2 \in E(G). \tag{D.1}$$

and

$$zc' \in E(G) \tag{D.2}$$

Note that **D.2** and **D.1** cannot both hold because then $G[c', b_1, z, r_1]$ would be an induced claw.

Let us first deal with the case that **D.1** holds. Consider the graph $B_1 = G[c', r_1, r_2, z, r_3]$. If B_1 is not an induced bull, then this can only happen because the edge zc' exists, which means **D.2** holds as well; a contradiction. So $B_1 = G[c', r_1, r_2, z, r_3]$ is an induced bull. Then, however, its horns z and r_3 would need to have a common neighbour contradicting the property of R being distance increasing with respect to $V(B) \cup \{c\}$.

So we are left with the situation where [D.2](#) holds. Consider $B_2 = G[z, c', r_1, b_1, r_2]$. The only edge which can prevent B_2 from being an induced bull would be $c'r_2$, which cannot exist because [D.2](#) holds. So B_2 is an induced bull, whose horns b_1 and r_2 need to have a common neighbour $c'' \in V(G - B_2)$. Now consider $G[r_2, r_1, r_3, c'']$, which is not allowed to be an induced claw. Hence, the edge $c''r_1$ exists. Finally consider $B''' = G[c'', r_1, r_2, z, r_3]$, which cannot be an induced bull since then z and r_3 would need to have a common neighbour contradicting the property of R being distance increasing with respect to $V(B) \cup \{c\}$. But the only edge which can prevent B''' from being an induced bull is zc'' . This leads to the contradiction that $G[c'', b_1, z, r_2]$ is an induced claw and completes Case 1.

Case 2. $r_0 = a_i$ for some $i \in \{1, 2\}$.

Say, without loss of generality, $r_0 = a_2$ and consider $G[a_2, z, b_2, r_1]$, which cannot be an induced claw. The edge zb_2 cannot exist because B is an induced bull. The edge zr_1 cannot exist because then $zr_1r_2\dots$ would be a distance increasing ray with respect to $V(B) \cup \{c\}$ starting in z , which leads to a contradiction as in Case 1. Hence, the edge b_2r_1 needs to exist. Now consider $B_2 = G[b_2, a_2, r_1, z, r_2]$, which is an induced bull. So the horns z and r_2 of B_2 have a common neighbour $c' \in V(G - B_2)$. If $c' \in V(B)$, then we get a contradiction to R being distance increasing with respect to $V(B) \cup \{c\}$. Otherwise, however, we obtain a distance increasing ray $zc'r_2r_3\dots$ with respect to $V(B) \cup \{c\}$ which starts in z . This leads to a contradiction as in Case 1. So we have completed our consideration of Case 2.

Case 3. $r_0 = c$.

Since $b_1b_2 \notin E(G)$ and $G[c, b_1, b_2, r_1]$ is not an induced claw, we know one of the edges b_1r_1 or b_2r_1 must exist, say without loss of generality b_2r_1 . If $B_3 = G[r_1, c, b_2, b_1, a_2]$ were an induced bull, its horns b_1 and a_2 would need to have a common neighbour c'' in $V(G - B_3)$. Note that $c'' \neq r_i$ and $c''r_{i+3} \notin E(G)$ for every $i \in \mathbb{N}$ as R is distance increasing with respect to $V(B) \cup \{c\}$. Now, however, $r_1r_2\dots$ is a distance increasing ray with respect to $V(B_3) \cup \{c''\}$ which starts at r_1 . This is the same situation as in Case 1 and, therefore, leads towards a contradiction.

So B_3 is not an induced bull and only three edges could possibly witness this, namely r_1b_1 , r_1a_2 or ca_2 . First, if $r_1b_1 \in E(G)$, then consider r_1 instead of c as the common neighbour of b_1 and b_2 outside of B and $r_1r_2\dots$ as the distance increasing

ray with respect to $V(B) \cup \{r_1\}$. Now we are again in the situation of Case 3 but know that $r_2b_1, r_2b_2 \notin E(G)$, implying that $G[r_1, b_1, b_2, r_2]$ is an induced claw; a contradiction. Hence, we conclude that $r_1b_1 \notin E(G)$.

Second, suppose that $r_1a_2 \in E(G)$. Then $a_2r_1r_2 \dots$ would be distance increasing ray with respect to $V(B) \cup \{c\}$ starting at a_2 , which leads to a contradiction as in Case 2.

Third, suppose $ca_2 \in E(G)$, but $r_1b_1, r_1a_2 \notin E(G)$. Then $G[c, b_1, a_2, r_1]$ is an induced claw. This contradiction completes the analysis of Case 3.

Case 4. $r_0 = b_i$ for some $i \in \{1, 2\}$.

Let, without loss of generality, $r_0 = b_2$ and consider $G[b_2, a_2, c, r_1]$, which cannot be an induced claw. By Case 2 and Case 3 we know that $r_1a_2, r_1c \notin E(G)$. So the edge ca_2 must exist. Consider the bull $B'_3 = G[b_2, c, a_2, b_1, z]$. If B'_3 is induced, then R is distance increasing with respect to $V(B'_3) \cup \{a_1\} = V(B) \cup \{c\}$ and a_1 is a common neighbour of the horns b_1 and z of B'_3 . This puts us again in the situation of Case 1 and leads to a contradiction.

So let us finally consider the case that B'_3 is not induced. The only reason for this is the existence of the edge cz . Then, however, $G[c, b_1, b_2, z]$ is an induced claw; a contradiction. \square

We continue by showing that every locally finite connected claw-free graph with at least three ends contains a net, and therefore also a bull, as an induced subgraph.

Lemma D.3.3. *Every locally finite, connected claw-free and net-free graph has at most two ends.*

Proof. Suppose for a contradiction that G is a locally finite connected, claw-free, net-free graph with at least three different ends ω_1, ω_2 and ω_3 . Let D be a geodesic double ray containing an ω_2 -ray and an ω_3 -ray, which exists due to Lemma D.2.4. Let S be a finite vertex set which is at least at distance 2 from D and which separates D from ω_1 , i.e. every ω_1 -ray with start vertex in D meets S . To see that such a vertex set exists, first pick a finite vertex set $S' \subseteq V(G)$ which pairwise separates ω_1, ω_2 and ω_3 . Hence, only a finite set F of vertices of $V(D)$ is contained in $V(C(S', \omega_1))$. Now set $S = N_2(S' \cup F) \cap C(S' \cup F, \omega_1)$, which is still a finite

set since S' as well as F are finite and G is locally finite. Furthermore, S now separates D from ω_1 as desired.

Now consider all shortest S - D -paths. Among such shortest paths let P be one that meets D closest to ω_3 , say at vertex d , i.e. there exists no S - D -path with an endvertex $d' \neq d$ on the ω_3 -ray that is contained in D and starts in d . Choosing d in such a maximal way is possible since S is a finite set and due to the local finiteness each distance class starting from S is a finite set as well. Let b be the neighbour of d that lies on the ω_2 -ray R_2 that is contained in D and starts at d . Further denote the neighbour of b on R_2 that is different from d by y . Let c be the neighbour of d that lies on the ω_3 -ray R_3 which is contained in D and starts at d . Finally, let a be the neighbour of d on P and z be the neighbour of a on P different from d .

We first note that there is no edge ac since this would yield a path of the same length as P ending closer to ω_3 . Furthermore, there is no edge bc , otherwise D would not be geodesic. Since $G[d, a, b, c]$ cannot be an induced claw, we know that $ab \in E(G)$. This situation would look like depicted in Figure D.2.




Figure D.2.: The graph $G[a, b, d, z, y, c]$ forms an induced net.

We furthermore note that z is not adjacent to any vertex in $\{y, b, d, c\}$ since this would yield an S - D path shorter than P . Since D was geodesic, c cannot be adjacent to y or b . It remains to show that $ay \notin E(G)$. Suppose for a contradiction that $ay \in E(G)$, then $G[a, d, y, z]$ is an induced claw, contradicting our assumption.

Hence, we proved that $G[a, b, d, z, y, c]$ is an induced net, which contradicts our assumption on G . \square

Note that the proof of Lemma [D.3.3](#) shows also the following.

Corollary D.3.4. *Let G be a locally finite, connected claw-free and net-free graph with two ends. Let D be a geodesic double ray containing rays to the two ends of G . Then every vertex of G has distance at most 1 from D .* \square

Now let us deduce two corollaries with respect to bulls.

Corollary D.3.5. *Every locally finite, connected claw-free and bull-free graph has at most two ends.*

Proof. Since every net contains an induced bull, the statement follows immediately from Lemma [D.3.3](#). \square

Corollary D.3.6. *Let G be a locally finite, connected claw-free graph. If the horns of every induced bull B have a common neighbour in G , then G has at most two ends.*

Proof. Since every graph as in the premise of the statement is either finite or already bull-free by Theorem [D.1.8](#), the statement follows immediately from Corollary [D.3.5](#). \square

While Lemma [D.3.3](#) and its corollaries limit the variety of possible graphs we are considering in terms of the number of ends they can have, we shall now show that these classes are non-trivial. Before we make explicit constructions, we need to state a definition.

Given a graph G and some $k \in \mathbb{N}$, we call a graph G' a *k-blow-up* of G if we obtain G' from G by replacing each vertex of G by a clique of size k , where two vertices of G' are adjacent if and only if they are either both from a common such clique, or the original corresponding vertices were adjacent in G .

Example D.3.7. For some $k \in \mathbb{N}$ consider the k -blow-up of a ray, yielding a graph with one end, and the k -blow-up of a double ray for an example with two ends. Next we check that these graphs are claw-free. Suppose for a contradiction there exists an induced claw C . Then the vertex c of degree 3 in C is contained in a clique corresponding to a vertex v of the ray (or double ray). Now, however, two

non-adjacent neighbours of c could only lie in two cliques which correspond to two neighbours of vertex v of the ray (or double ray). The third vertex of degree 1 in C cannot be contained in any clique without causing a contradiction to C being an induced subgraph.

These graphs are also bull-free and, therefore, net-free as well. Suppose there exists an induced bull B , consisting of a triangle $K = G[a_1, a_2, z]$ and the two horns b_1 and b_2 where $a_i b_i \in E(B)$ for every $i \in \{1, 2\}$. First note that one edge xy of K has to lie in clique corresponding to a vertex of the ray (or double ray). By the structure of a bull, x or y is adjacent to a horn h of B . Then, however, by the structure of the whole graph, h is must be adjacent to both, x and y , contradicting that B is an induced subgraph.

Finally, let us mention that Lemma D.3.3 only holds for locally finite graphs. In the following example we exhibit countable claw-free and net-free graphs, which are not locally finite, but have $k \geq 3$ or countably many ends:

Example D.3.8. Fix for every $i \in \mathbb{Z}$ a clique $K_{\aleph_0}^i$ of size \aleph_0 such that the vertex sets of all these cliques and the set \mathbb{Z} are pairwise disjoint. Furthermore, fix two distinct vertices i^+ and i^- in each $K_{\aleph_0}^i$. For any set $I \subseteq \mathbb{Z}$ we now define the graph D_I as follows. Set $V(D_I) = (\mathbb{Z} \setminus I) \cup \bigcup_{i \in I} V(K_{\aleph_0}^i)$. If $i, (i+1) \in \mathbb{Z} \setminus I$, set $i(i+1) \in E(D_I)$. For $i \in I$ and $(i+1) \notin I$, set $i^+(i+1) \in E(D_I)$, and for $i \in I$ and $i-1 \notin I$, set $i^-(i-1) \in E(D_I)$. Finally, for $i, (i+1) \in I$ set $i^+i^- \in E(D_I)$. This completes the definition of D_I .

It is easy to see that that D_I has precisely $|I| + 2 \in \mathbb{N} \cup \{\infty\}$ many ends. Now suppose D_I contains a claw or a net, call it H . Note that each vertex of H with degree 3 in H has to lie in some $K_{\aleph_0}^i$ with $i \in I$. From this point on it is easy to see that no matter where the vertices of degree 1 in H are located in D_I , the graph H cannot be an induced subgraph of D_I .

D.4. Hamiltonicity results

The main goal of this section is to extend Theorem D.1.1 to locally finite graphs. Let us briefly recall the theorem here.

Theorem D.1.1. [98, Thm. 2.9] *Let G be a finite graph. If G is claw-free and net-free, then*

- (1) G is connected implies G has a Hamilton path.
- (2) G is 2-connected implies G is Hamiltonian.
- (3) For $k \geq 2$, G is $(k + 1)$ -connected or $G = K_{k+1}$ if and only if G is k -leaf-connected.

Next we prove Theorem D.1.3, which is an extension of statement (2) of Theorem D.1.1 to locally finite graphs. Our key tool to prove this result is the characterisation of infinite locally finite claw-free and net-free graphs in terms of distance-2-completeness, which is Theorem D.1.6.

Theorem D.1.3. *Every locally finite, 2-connected claw-free and net-free graph is Hamiltonian.*

Proof. Let G be a graph as in the statement of this theorem. By statement (2) of Theorem D.1.1 we can assume G to be infinite. By Lemma D.3.3 we know that G has at most two ends. We only write the proof of this theorem for the case when G has precisely two ends, say ω_1 and ω_2 . The case that G has only one end works analogously, but is slightly easier.

Let $S \subseteq V(G)$ be any finite minimal vertex separator of G , which exists since G is locally finite. Furthermore, let us fix some $v \in S$. By Proposition D.3.1 we know that $G - S$ has exactly two components, call them L and R . By Theorem D.1.6 we know that the graph $G - (S \setminus \{v\})$ is distance-2-complete centered at v . Within the graph $G - (S \setminus \{v\})$ let L_i and R_i denote the i -th distance classes of v in the components L and R , respectively. Let $\ell \in \mathbb{N}$ be the maximum number with the property that S has a neighbour in L_ℓ or R_ℓ . Now we define a hierarchy of subgraphs starting with

$$G_0 := G \left[S \cup \bigcup_{i=1}^{\ell} L_i \cup \bigcup_{i=1}^{\ell} R_i \right].$$

Furthermore, we define:

$$G_{i+1} := G[V(G_i) \cup L_{i+1} \cup R_{i+1}].$$

Note that $G_0 = G_j$ for every $j \in \mathbb{N}$ with $j \leq \ell$, but $G_i \subsetneq G_{i+1}$ for every $i \in \mathbb{N}$ with $i \geq \ell$ since G is infinite.

Since each G_i is an induced subgraph of G , it is claw-free and net-free as well. Furthermore, each G_i is finite since G is locally finite. Using that each subgraph $G[L_i]$ and $G[R_i]$ is complete because $G - (S \setminus \{v\})$ is distance-2-complete centered at v , it follows easily from the definition of G_0 and from G being 2-connected that each G_i is 2-connected as well. Hence we get from statement (2) of Theorem D.1.1 that each G_i contains a Hamilton cycle.

We now prove that there is a Hamilton cycle C'_n in G_n for every $n > \ell$ such that $|E(C'_n) \cap \delta(R_n)| = 2$ and $|E(C'_n) \cap \delta(L_n)| = 2$ holds. We start by fixing an arbitrary Hamilton cycle C_n of G_n and fix an orientation of C_n . Starting from v this cycle has to meet R_n at some point the first time, say in vertex w_1 , via the edge v_1w_1 with $v_1 \in R_{n-1}$ (cf. Figure D.3). Beginning from v , say the first time C_n leaves R_n happens at vertex w_2 via the edge w_2v_2 for some $v_2 \in R_{n-1}$. To define the desired Hamilton cycle C'_n we follow C_n from v till w_1 and collect all vertices from R_n ending in w_2 , which we can do since each $G[R_n]$ complete. We now return to v_2 via the edge w_2v_2 . Next we follow C_n but whenever the cycle goes from R_{n-1} to R_n , say via some edge v_kv_k , and comes back from R_n to R_{n-1} the next time, say via an edge $v_{k+1}w_{k+1}$, we replace this segment of the cycle by the edge v_kv_{k+1} , which exists since $G[R_{n-1}]$ is complete.




Figure D.3.: Modified cycle C'_n meeting $E(R_{n-1}, R_n)$ only twice by replacing grey edges of C_n by black dashed edges.

Similarly, we can modify the cycle to incorporate L_n in this manner. Thus we have found the desired Hamilton cycle C'_n in G_n . Note that in this process we only altered the initial Hamilton cycle C_n at L_n , L_{n-1} , R_{n-1} and R_n .

We now show that for any $n > \ell$ the cycle C'_n can be extended without altering edges in G_{n-1} to a Hamilton cycle D_{n+1} of G_{n+1} such that $|\delta(R_n) \cap E(D_{n+1})| = 4$ and $|\delta(R_{n+1}) \cap E(D_{n+1})| = 2$ holds as well as analogue statements for L_n and L_{n+1} . We shall give the argument only for R_{n+1} as the modification for L_{n+1} works analogously. We know that there are at least two edges $a_1b_1, a_2b_2 \in E(R_n, R_{n+1})$ which do not share a common endvertex and where $a_i \in R_n$ and $b_i \in R_{n+1}$ for any $i \in \{1, 2\}$, since G_{n+1} is 2-connected. If $|R_n| = 2$, we can easily use these two edges to get the desired extension D_{n+1} of C'_n . So we may assume that R_n has at least 3 vertices. Say without loss of generality that a_1 lies before a_2 on C'_n . Furthermore, say that C'_n meets R_n the first time (starting from v) in w_1 via the edge v_1w_1 with $v_1 \in R_{n-1}$ and leaves R_n the last time from w_2 via the edge w_2v_2 where $v_2 \in R_{n-1}$. We now have to consider two cases.

Case 1. $|\{w_1, a_1, a_2, w_2\}| \geq 3$.

Without loss of generality let $a_1 \neq w_1$. In this case we follow C'_n until w_1 , then collect all vertices from R_n but w_2 and a_2 such that we end in a_1 (see Figure D.4), which we can do since $G[R_n]$ is a clique. Next we use the edge a_1b_1 , collect all vertices in R_{n+1} while ending in b_2 , return to R_n via the edge b_2a_2 . If $a_2 = w_2$, we can immediately follow C'_n to close a cycle. Otherwise we use the edge a_2w_2 and then proceed with C'_n to close a cycle. Doing the same with L_{n+1} yields the desired D_{n+1} . This completes the argument in Case 1.




Figure D.4.: The situation in Case 1.

Case 2. $a_1 = w_1$ and $a_2 = w_2$.

Let x be an arbitrary vertex from $R_n - \{a_1, a_2\}$. Since $G[w_1, v_1, x, b_1]$ is not an induced claw, one of the edges v_1x or xb_1 exists (cf. Figure D.5). If v_1x exists, we can operate as in Case 1 by just switching the roles of x and w_1 . Should xb_1 exist, we can proceed as in Case 1 as well, this time by switching the roles of x and a_1 . This completes the argument for Case 2.




Figure D.5.: The situation in Case 2: the dashed edges prevent the graph $G[w_1, v_1, x, b_1]$ from being an induced claw.

This shows that we can always extend C'_n to the desired cycle D_{n+1} . Since D_{n+1} is also a valid candidate for C'_{n+1} and $D_{n+1} \cap G_{n-1} = C'_n \cap G_{n-1}$, we can inductively extend C'_n through all R_n and L_n with $n > \ell$ and obtain a well-defined subgraph C as limit object via its edge set:

$$E(C) := \left\{ e \in E(G) ; \exists k \in \mathbb{N} : e \in \bigcap_{i \geq k} E(D_i) \right\}.$$

The rest of this proof consist of verifying that \overline{C} is a Hamilton circle of G . By the definition of C we immediately get that every vertex of G is contained in C . It remains to check that \overline{C} is a circle in $|G|$. From the definition of all the D_i and C we immediately get that every vertex of G has degree 2 in C . By Lemma D.2.9 it remains to prove that C is topologically connected and that every end of G has degree 2 in \overline{C} .

In order to prove that \overline{C} is topologically connected, it is enough to show that C meets every finite cut of G by Lemma D.2.6. This, however, holds since each finite cut F of G is eventually contained in G_m for all $m > m_0$ where m_0 is some sufficiently large integer, which means that each Hamilton cycle D_m of G_m for $m > m_0$ meets F in the same set of edges and, hence, so does C . We even get that each finite cut F of G is met in an even number of edges by C , since the intersection of a cycle and a cut is always even.

Note first that \overline{C} contains every end ω of G because C contains every vertex of G . Now using that being topologically connected and being arc-connected is equivalent for closed subspaces of $|G|$ by Lemma D.2.7, we know that for every end ω of G there exists at least one arc in \overline{C} with ω as its endpoint. So each end of G has degree at least 1 in \overline{C} .

Next let us prove that each end of G has degree at most 2 in \overline{C} . For this let us define $R_{\geq n} = \bigcup_{i=n}^{\infty} R_i$ and $L_{\geq n} = \bigcup_{i=n}^{\infty} L_i$ for every $n > \ell$. Since each R_n and each L_n separates the two ends of G if $n > \ell$, say every ω_1 -ray has a tail in $R_{\geq n}$ and every ω_2 -ray has a tail in $L_{\geq n}$ for every $n > \ell$. By definition of C we get that $|\delta(R_{\geq n}) \cap E(C)| = |E(R_{n-1}, R_n) \cap E(C)| = 2$ holds for every $n > \ell$. Now suppose for a contradiction that \overline{C} contains at least three vertex disjoint arcs all of which have ω_1 as their endpoint. By choosing $m \in \mathbb{N}$ big enough, we can guarantee that each of the three arcs contains a vertex in $G - R_{\geq m}$. The partition $(V(G - R_{\geq m}), V(R_{\geq m}))$ of $V(G)$ induces the finite cut $\delta(R_{\geq m})$ and both of its sides are met by \overline{C} due to the choice of m and since ω_1 is contained in $\overline{R_{\geq m}}$. By Lemma D.2.5 we now know that each of the three arcs must use an edge of $\delta(R_{\geq m})$. However, the three arcs being vertex disjoint now contradicts $|\delta(R_{\geq n}) \cap E(C)| = 2$. So ω_1 has degree at most 2 in \overline{C} . An analogue argument shows that ω_2 has degree at most 2 in \overline{C} as well.

Finally, we prove that each end of G has degree at least 2 in \overline{C} . Theorem D.2.8 tells us that each end of G has an even degree in \overline{C} if C meets every finite cut of G in an even number of edges. This holds as already proven above. So we can conclude that both ends of G have degree precisely 2 in \overline{C} , which completes the proof that \overline{C} is a Hamilton circle of G . \square

We move on by proving Theorem D.1.4, which is an extension of statement (3) of Theorem D.1.1 to locally finite graphs. Recall that we call a locally finite graph G *topologically k -leaf-connected* where $k \in \mathbb{N}$ if $|V(G)| > k$ and given any set $S \subseteq V(G) \cup \Omega(G)$ with $|S| = k$, then G has a topological spanning tree \overline{T} whose set of leaves is precisely S .

Theorem D.1.4. *Let G be a locally finite, connected, claw-free and net-free graph, and let $k \in \mathbb{N}$ satisfy $k \geq 2$. Then G is $(k + 1)$ -connected or $G = K_{k+1}$ if and only if G is topologically k -leaf-connected.*

Proof. By statement (3) of Theorem D.1.1 we may assume for both implications

that G is infinite. Let us first assume that G is an infinite, but locally finite graph that is topologically k -leaf-connected. We show that G is $(k + 1)$ -connected. Assume for a contradiction that G has a vertex separator $S \subseteq V(G)$ of size at most k . Let $S' \supseteq S$ be a superset of S such that S' still separates G , which is possible since G is infinite, and $|S'| = k$. By the topologically k -leaf-connectedness there exists a subgraph T of G such that \bar{T} is a topological spanning tree of G , whose set of leaves is exactly S' . Since \bar{T} contains vertices from two components of $G - S'$, it must traverse S by Lemma D.2.5, yielding a vertex of from S that has degree at least 2 in T ; a contradiction.

Suppose for the other implication that G is an infinite, but locally finite claw-free and net-free graph which is $(k + 1)$ -connected for $k \geq 2$. We show that G is topologically k -leaf-connected. By Theorem D.1.6 we know that for every finite minimal vertex separator $S \subseteq V(G)$ of G and every $v \in S$, the graph $G - (S \setminus \{v\})$ is distance-2-complete centered at v . So let us fix such an S and some $v \in S$. Let G_i, R, R_i, L and L_i be defined as in the proof of Theorem D.1.3.

Let us fix some $B = \{l_1, \dots, l_k\} \subseteq V(G) \cup \Omega(G)$ for the rest of the proof. We have to show that a topological spanning tree of G exists whose set of leaves is exactly B . By Theorem D.1.6 we know that G has at most two ends. We shall give the proof only in the case when G has precisely one end. For the case that G contains two ends the argument can easily be adapted. A consequence of assuming G to have only one end is that L or R is finite, say L . We shall distinguish two cases, namely whether B contains one of the two ends of G or not.

First let us assume that B contains no ends of G . Let $\ell \in \mathbb{N}$ be sufficiently large such that B, L and a finite connected subgraph in R containing $N(S) \cap R$ is contained in $G_{\ell-1}$. Similar to our proof of Theorem D.1.3 we now show that there exists a spanning tree $T_{\ell+1}$ of $G_{\ell+1}$ with precisely B as its set of leaves such $|\delta(R_{\ell+1}) \cap E(T_{\ell+1})| = 2$ holds in $G_{\ell+1}$.

To prove this, first we verify that $G_{\ell+1}$ is also $(k + 1)$ -connected. Otherwise, there exists a separator S_k of size at most k , separating two vertices x and y in $G_{\ell+1}$, but not in G . Hence, there exists an x - y -path P in G disjoint to S_k . Since P does not exist in $G_{\ell+1}$, it must pass through $R_{\ell+1}$ to $R_{\ell+2}$. By shorten P on $G[R_{\ell+1}]$, which is a clique, we obtain an x - y -path in $G_{\ell+1}$ which is disjoint to S_k ; a contradiction.

Since $G_{\ell+1}$ is $(k + 1)$ -connected and, as an induced subgraph of G , also claw-free

and net-free, there exists a spanning tree $T'_{\ell+1}$ of $G_{\ell+1}$ whose set of leaves is B by statement (3) of Theorem D.1.1. Next we modify this tree to obtain the desired tree $T_{\ell+1}$. First we root $T'_{\ell+1}$ in v and orient its edges away from the root. Now we get $T_{\ell+1}$ from $T'_{\ell+1}$ by shortening all but one of the directed paths P starting and ending in R_ℓ and otherwise only using vertices from $R_{\ell+1}$ by an edge from the start vertex to the endvertex of P , which exists since $G_{\ell+1}$ is a clique. Note that if we do not need to do this replacement, then $T'_{\ell+1}$ is already as desired. We now modify the remaining of such a path in $R_{\ell+1}$ to one containing all vertices of $R_{\ell+1}$, but with the same start and endvertex. The resulting graph is our $T_{\ell+1}$, which is indeed a tree since it is connected and every cycle in $T_{\ell+1}$ would yield a cycle in $T'_{\ell+1}$ either directly or by replacing edges from $E(T_{\ell+1}) \setminus E(T'_{\ell+1})$ by the corresponding paths in $R_{\ell+1}$.

We can now extend $T_{\ell+1}$ to a topological spanning tree \bar{T} of G , where T is a corresponding subgraph of G . For this we extend the unique branch of $T_{\ell+1}$ starting in v and ending in some leaf l_i of $T_{\ell+1}$ that contains edges from $G[R_{\ell+1}]$. We modify this branch at an edge in $G[R_{\ell+1}]$ to an arc via the end of G starting in v , ending in l_i and containing all remaining vertices of G , which are precisely those in $\bigcup_{i=\ell+2}^{\infty} R_i$. This modification can be done similarly to our extension of the Hamilton cycles in the proof of Theorem D.1.3. To see that the resulting standard subspace \bar{T} of G is indeed a topological spanning tree of G , we have to check that it is topologically connected and does not contain a circle from $|G|$. Similarly as in the proof of Theorem D.1.3, it is easy to check that T intersects every finite cut of G . Hence, Lemma D.2.6 implies that \bar{T} is topologically connected. To see that \bar{T} does not contain any circle from $|G|$, note that any circle within \bar{T} which corresponds to a finite cycle C in T would also imply that C is already contained in some T_n for sufficiently large $n \in \mathbb{N}$, which contradicts that T_n is a tree. In the case that \bar{T} contains a circle which does not correspond to a finite cycle in T , this circle would induce a cycle in some T_n for sufficiently large $n \in \mathbb{N}$. This is done by replacing an arc of the circle which uses the end of G and whose endpoints are two distinct vertices $a, b \in R_n$ by the edge $ab \in E(G[R_n])$. Hence, we again obtain a contradiction to T_n being a tree. This completes the proof for the first case.

Now let us assume that $B = \{l_1, \dots, l_{k-1}, \omega\}$ contains the end ω of G . As in the first case choose a sufficiently large $\ell \in \mathbb{N}$ such that $B - \{\omega\}$, L and a finite connected subgraph in R containing $N(S) \cap R$ is contained in $G_{\ell-1}$. Let w be a

vertex in R_ℓ with a neighbour in $R_{\ell+1}$. By statement (3) of Theorem D.1.1 there is a spanning tree T_ℓ in G_ℓ with $B - \{\omega\} \cup \{w\}$ as its set of leaves. Let us pick $v \in S$ as the root of T_ℓ . But now we can extend the branch ending in w by an ω -ray Q starting in w with $V(Q) = \{w\} \cup \bigcup_{i=\ell+1}^{\infty} R_i$. Similar as in the first case it is easy to verify that the closure \overline{T} of the resulting subgraph T yields a topological spanning tree of G whose set of leaves is precisely B . \square

Finally, we prove Theorem D.1.2, which forms an extension of statement (1) of Theorem D.1.1 to locally finite graphs.

Theorem D.1.2. *For an infinite locally finite connected graph G that is claw-free and net-free, precisely one of the following statements holds:*

- (1) G has only one end and admits a spanning ray.
- (2) G has only two ends and admits a spanning double ray.

Proof. We shall distinguish three cases with respect to the connectivity of G . Suppose first that G is not 2-connected. Then let v be a cut vertex of G . By Theorem D.1.6 we know that G is distance-2-complete centered at v . So there are precisely two components of $G - v$, both of which are infinite if G has two ends, and just one of them is infinite in case G has only one end. Using the structure of distance-2-complete graphs we easily find either a spanning double ray if G has two ends, or a spanning ray if G has only one end.

For the second case let us assume that G is 3-connected. By Theorem D.1.4 we know that G is 2-leaf-connected. So we can find for any two distinct $x, y \in V(G) \cup \Omega(G)$ a Hamilton arc of G with x and y as endpoints. In the case that G has precisely two ends, we can find a Hamilton arc of G with these ends as its endpoints. Since G has no further ends, we immediately get that this Hamilton arc induces a double ray in G .

Similarly, in the case when G has just one end ω we can find a Hamilton arc of G whose endvertices are some arbitrary vertex and ω . Also similarly as before, as G has only one end, this Hamilton arc is a desired spanning ray of G , which completes the second case.

It remains to prove the statement under the assumption that G is 2-connected, but not 3-connected. Hence, there is a minimal separator $\{u, v\}$ of G with u and v being distinct. By Theorem D.1.6 we know that $G - u$ and $G - v$ are

distance-2-complete centered at v and u , respectively. Let R and L be the two components of $G - \{u, v\}$.

Since $\{u, v\}$ is a minimal vertex separator of G , we know that u has at least one neighbour in R as well as in L . Furthermore, any two neighbours of u in R (or L) must be adjacent due to the claw-freeness of G . Hence, any two neighbours of u in R either lie in some common distance class of v within R or in two successive distance classes of v within R . An analogue statement holds for neighbours in L .

Now suppose u has two distinct neighbours in L or R , say R . As before, let R_i and L_i denote the i -th distance classes of v in the components R and L , respectively. Furthermore, let $R_0 = L_0 = \{v\}$. Let $n \in \mathbb{N}$ be minimal such that u has no neighbour in R_{n+1} . Now let u_1, u_2 be two distinct neighbours of u in R such that $u_2 \in R_n$. Hence, u_1 either lies in R_n with $n \geq 1$ or in R_{n-1} with $n \geq 2$. By the choice of n , we know that u has no neighbour in R_m for $m > n$. Hence, each R_m is a separator in G for $m \geq n$ if $R_{m+1} \neq \emptyset$. Furthermore, in case $R_{m+1} \neq \emptyset$, the 2-connectedness of G implies that there are at least two edges g_m, h_m each with one endvertex in R_m and the other in R_{m+1} for all $m \geq n$ such that g_m and h_m do not have a common endvertex in R_m . Next we prove the existence of a spanning ray or double ray T in $G - u$, depending whether G has only one or two ends, that uses the edge u_1u_2 . We shall distinguish two cases.

Case 1. $u_1, u_2 \in R_n$ with $n \geq 1$.

Let $x \in R_n$ be distinct from u_1 and u_2 if it exists, otherwise let $x = u_1$. We first pick a finite path P from x to v whose vertex set consists precisely of $\{x\} \cup \bigcup_{i=0}^{n-1} R_i$. Such a path exists since each R_i induces a clique and each vertex in R_i has a neighbour in R_{i-1} for every $i > 0$. Next we extend P within $G[R_n]$ to a path P' which contains all vertices of R_n as well as the edge u_1u_2 and which has an endvertex incident with g_n or h_n in case $R_{n+1} \neq \emptyset$. Due to the existence of g_m and h_m for all $m \geq n$ where $R_{m+1} \neq \emptyset$, we can extend P' to a finite path or a ray T' , depending whether R is finite or infinite, that starts in v , uses the edge u_1u_2 and whose vertex set equals $V(R) \cup \{v\}$. Similarly as before, we finally extend T' to also contain all vertices of $V(L)$, which yields the desired T .

Case 2. $u_1 \in R_{n-1}$ with $n \geq 2$.

Let $x \in R_{n-1}$ be distinct from u_1 if it exists, otherwise let $x = u_1$. Similarly as in Case 1, pick a finite path P from x to v whose vertex set consists precisely of

$\{x\} \cup \bigcup_{i=0}^{n-2} R_i$. Now extend P within $G[R_{n-1}]$ to a path P' which also contains all vertices of R_{n-1} and which has u_1 as its endvertex. By making use of the existence of g_m and h_m for all $m \geq n$ where $R_{m+1} \neq \emptyset$, extend P' via the edge u_1u_2 to a finite path or ray T' that starts in v , uses the edge u_1u_2 and whose vertex set equals $V(R) \cup \{v\}$. Finally, extend T' to also contain all vertices of $V(L)$, which yields the desired T .

Hence, we can find the desired spanning ray or double ray T in $G - u$ that uses the edge u_1u_2 . Similarly as before, we now incorporate u by replacing the edge u_1u_2 in T by the path u_1uu_2 , which yields a desired spanning ray or double ray in G .

It remains to prove the existence of the desired spanning ray or double ray under the assumption that u has precisely one neighbour in each of R and L . Since G has at least one end, one of L or R must be infinite, say R . Let u^* denote the neighbour of u in R . Now let $n \in \mathbb{N}$ be such that $u^* \in R_n$. Let us first deal with the case that L contains precisely one vertex. Then we can find a spanning ray T in $G - u$ since v is a cutvertex of $G - u$. The unique vertex in L , call it v_L , must be the start vertex of T and v must be the second vertex of T because of v being a cutvertex of $G - u$. Since G is 2-connected, v_L must be adjacent to u in G . So we get our desired spanning ray of G by taking T and adding the edge uv_L . For the case that L contains at least two vertices we shall derive a contradiction. So suppose for a contradiction that L contains at least two vertices. We may assume that u and v are not adjacent in G . To see this let us assume to the contrary that $uv \in E(G)$. Since L contains at least two vertices, G is 2-connected, but u has only one neighbour in L , there exists a vertex $w_L \in V(L)$ which is adjacent to v but not to u . Similarly, there exists a vertex w_R in R which is adjacent to v , but not to u . Now, however, we have a contradiction since the graph $G[v, u, w_L, w_R]$ is an induced claw in G .

Since R_n is a separator in G , there exists a vertex $x \in R_n$ which is distinct from u^* and adjacent to some vertex in R_{n+1} . Let x^+ denote some neighbour of x in R_{n+1} and x^- denote a neighbour of x in R_{n-1} . Furthermore, let x^{--} denote a neighbour of x^- in R_{n-2} if $n \geq 2$, and otherwise let x^{--} be a neighbour of v in L which is not adjacent to u . Note that the latter is possible since for $n = 1$, we have that $x^- = v$, and G being 2-connected together with L containing at least two vertices guarantees the existence of the desired x^{--} under the assumption that u

has precisely one neighbour in L . Now we prove that the edge u^*x^+ cannot exist. Suppose for a contradiction that $u^*x^+ \in E(G)$ and let u^- denote a neighbour of u^* in R_{n-1} . Since the only neighbour of u in L is u^* , we get that u is not adjacent to x^+ . In case $n \geq 2$, we also see by the same argument that u is not adjacent to u^- . In case $n = 1$, we get that $u^- = v$, and hence by assumption that u is again not adjacent to u^- . Finally, u^- is not adjacent to x^+ because R_n is a separator in G . Hence, $G[u^*, x^+, u^-, u]$ is an induced claw in G , which is a contradiction. Next we consider the graph $G[x, u^*, x^-, x^+]$. We know that x^- and x^+ are not adjacent because R_n is a separator in G . As argued before, u^* and x^+ are not adjacent as well. Hence, x^- and u^* must be adjacent to prevent $G[x, u^*, x^-, x^+]$ from being an induced claw in G . Now we derive a contradiction because the graph $G[u^*, x, x^-, u, x^+, x^{--}]$ is an induced net.

□

E. Forcing Hamiltonicity in locally finite graphs via forbidden induced subgraphs II: paws

E.1. Introduction

In this second paper out of a series we extend another sufficient condition for Hamiltonicity in finite graphs to locally finite ones. For this we consider, given a locally finite connected graph G , the topological space $|G|$ [24, 25], known as the Freudenthal compactification of G . Beside the graph G , seen as a 1-complex, the space $|G|$ also contains additional points, namely the *ends* of G , which are equivalence classes of one-way infinite paths of G under the relation of being inseparable by finitely many vertices. Following the topological approach from [27, 28], we use *circles*, i.e. homeomorphic images of the unit circle $S^1 \subseteq \mathbb{R}^2$ in $|G|$, to extend the notion of cycles and allowing infinite ones. Then we call G *Hamiltonian* if there is a circle in $|G|$ containing all vertices of G .

This series of articles focuses on extending certain local conditions that guarantee the existence of a Hamilton cycle in finite graphs, namely such in terms of forbidden induced subgraphs. In this paper we focus on a condition involving precisely two graphs: the *claw*, i.e. $K_{1,3}$, and the *paw*, which is the graph obtained from a triangle and an additional vertex which is adjacent to precisely one vertex of the triangle (cf. Figure E.1). For the rest of this paper, we shall denote the vertex of degree 1 in a paw by a_1 and those two vertices non-adjacent to a_1 by b_1 and b_2 . The remaining vertex of a paw will always be called a_0 , as depicted in Figure E.1.

The following theorem is probably the first Hamiltonicity result for finite graphs in terms of forbidden induced subgraphs. Note that, given any two graphs G and H , we call G a *H-free* graph if G does not contain any induced subgraph isomorphic to H .




Figure E.1.: The subgraphs we focus on in this paper.

Theorem E.1.1. [39, Thm. 4] *Every finite 2-connected claw-free and paw-free graph is Hamiltonian.*

Even for finite graphs the condition in Theorem E.1.1 is very restricting: the only graphs satisfying this condition are cycles, cliques and cliques with a matching removed. We shall see in Section E.3 that there do not exist any infinite locally finite 2-connected claw-free and paw-free graphs.

However, we shall study a variant of the condition in Theorem E.1.1 where the paw-freeness is relaxed. We focus on the following Hamiltonicity result due to Broersma and Veldmann. In order to state it we have to give two further definitions. A graph is called *pancyclic* if it contains a cycle of every possible length. Let H be an induced subgraph of a graph G and $v, w \in V(H)$. We shall write $\phi_H(v, w)$ for the property that v and w have a common neighbour in G outside of $V(H)$. We shall simply write $\phi(v, w)$ if the context makes it clear to which subgraph H we are referring to.

Theorem E.1.2. [16, Thm. 2] *Let G be a finite, 2-connected, claw-free graph. If every induced paw of G satisfies $\phi(a_1, b_i)$ for some $i \in \{1, 2\}$, then either G is pancyclic or G is a cycle.*

Obviously being pancyclic implies Hamiltonicity for finite graphs. We shall generalise Theorem E.1.2 to locally finite graphs, where we focus on verifying Hamiltonicity. We do this since we probably have no meaningful length parameter for distinguishing different infinite cycles, but we have a meaningful notion for Hamiltonicity. More precisely, we will prove the following:

Theorem E.1.3. *Let G be a locally finite, 2-connected, claw-free graph. If every induced paw of G satisfies $\phi(a_1, b_i)$ for some $i \in \{1, 2\}$, then G is Hamiltonian.*

Different from the graphs considered in our first paper of this series [57], we shall state examples of graphs fulfilling the premise of Theorem E.1.3 with arbitrarily, but finitely many ends, with \aleph_0 many and with 2^{\aleph_0} many ends.

The structure of this paper is as follows. In Section E.2 we introduce the needed definitions and notation. Then we conclude that section by introducing the tools we need for the proof of Theorem E.1.3. In Section E.3 we show that no infinite locally finite graphs exist that meet the criteria of Theorem E.1.1, except for being finite. Afterwards we give examples of infinite locally finite graphs fulfilling the conditions of Theorem E.1.3. Finally, we prove our main result, Theorem E.1.2, in Section E.4 and we start that section with a sketch of our proof.

E.2. Preliminaries

In general we follow the graph theoretical notation from [24]. Especially regarding the topological notions for locally finite graphs, we refer to [24, Ch. 8.5]. To see a wider survey regarding topological infinite graph theory, see [25].

E.2.1. Basic notions

All graphs considered in this paper are undirected and simple. Generally, we do not assume a graph to be finite. We call a graph *locally finite* if every vertex has finite degree.

For the rest of this section let G denote some graph. Later in this section, however, we shall make further assumptions on G .

Let X be a vertex set of G . We denote by $G[X]$ the induced subgraph of G with vertex set X . For small vertex sets, we sometimes omit the set brackets, i.e. we write $G[a, b, c]$ as a short form for $G[\{a, b, c\}]$. We write $G - X$ for the graph $G[V \setminus X]$. If H is a subgraph of G we shall write $G - H$ instead of $G - V(H)$. Again we omit set brackets around small vertex sets, especially for singleton sets. We briefly denote the cut $E(X, V \setminus X)$ by $\delta(X)$. For any $i \in \mathbb{N}$ we denote by $N_i(X)$ and $N_i(v)$ the set of vertices of distance at most i in G from the vertex set X or from a vertex $v \in V(G)$. We denote by $\partial(X)$ the set of vertices v of X with $N(v) \not\subseteq X$.

Let H be a subgraph of G and $v, w \in V(H)$. We denote by $\varphi_H(v, w)$ the

property that v and w have a common neighbour in $G - H$. We shall drop the subscript when it is clear to which subgraph H we refer to.

Let C be a cycle of G and u be a vertex of C . We implicitly fix an orientation of the cycle and we write u^+ and u^- for the neighbour of u in C in positive and negative, respectively, direction of C using a fixed orientation of C . Later on we will not always mention that we fix an orientation for the considered cycle using this notation. For two vertices v and w on C , we denote by vCw the v - w path in C that follows the orientation from v to w .

If G is a finite graph containing a cycle of length s for every $s \in \{3, 4, \dots, |V(G)|\}$, we call G *pancyclic*.

If v and w are vertices of a tree T , then we denote by vTw the unique v - w path in T .

A one-way infinite path R in G is called a *ray* of G and a two-way infinite path in G is called a *double ray* of G . A subgraph of a ray R that is itself a ray is called a *tail* of R . The unique vertex of degree 1 of R is called the *start vertex* of R . For a vertex r on a ray R , we denote the tail of R with start vertex r by rR .

An equivalence relation can be defined on the set of all rays of G by saying that two rays in G are *equivalent* if there are infinitely many disjoint paths in G joining these two rays. It is easy to check that this defines in fact an equivalence relation. The corresponding equivalence classes of rays under this relation are called the *ends* of G . We denote the set of ends of a graph G by $\Omega(G)$. If $R \in \omega$ for some end $\omega \in \Omega(G)$, then we briefly call R an ω -ray.

Note that for any end ω of G and any finite vertex set $S \subseteq V(G)$ there exists a unique component $C(S, \omega)$ that contains tails of all ω -rays. We say that a finite vertex set $S \subseteq V(G)$ *separates* two ends ω_1 and ω_2 of G if $C(S, \omega_1) \neq C(S, \omega_2)$. Note that any two different ends can be separated by a finite vertex set.

Let R be a ray in G and $X \subseteq V(G)$ be finite. We call R *distance increasing w.r.t. X* if $|V(R) \cap N_i(X)| = 1$ for every $i \in \mathbb{N}$. Note that a distance increasing ray w.r.t. X has its start vertex in X .

E.2.2. Topological notions

We assume G to be locally finite and connected for the rest of this section. The graph G together with its ends can be endowed with a certain topology, yielding

the space $|G|$ referred to as *Freudenthal compactification* of G . Note that within $|G|$, every ray of G converges to the end of G it is contained in. For a precise definition of $|G|$, see [24, Ch. 8.5]. See [36] for Freudenthal's paper about the Freudenthal compactification, and see [29] about the connection to $|G|$.

Given a point set X in $|G|$, we denote its closure in $|G|$ by \overline{X} .

We call the image of a homeomorphism which maps from the unit circle $S^1 \subseteq \mathbb{R}^2$ to $|G|$ a *circle* of G . We call G *Hamiltonian* if there is a circle in $|G|$ containing all vertices of G , and thus also all ends of G due the closedness of circles. Such a circle is called a *Hamilton circle* of G . Note that this notion coincides with the usual notion of Hamiltonicity for finite graphs.

E.2.3. Tools

In this subsection we introduce some basic lemmas we shall use to prove our results. We begin with a brief lemma about the existence of distance increasing rays with respect to finite vertex sets. The proof of this lemma works via a very easy compactness argument and we omit it here. In case a proof is desired, see for example [57, Lemma 2.3].

Lemma E.2.1. *Let G be an infinite locally finite connected graph and $X \subseteq V(G)$ be finite. Then there exists a distance increasing ray w.r.t. X .*

The following statements are all about claw-free graphs. The first is a very easy observation and probably folklore, so we do not prove it here. However, in case a proof is desired, consider for example [56, Prop. 3.7.].

Proposition E.2.2. *Let G be a connected claw-free graph and S be a minimal vertex separator in G . Then $G - S$ has exactly two components.*

Since we have to extend cycles very carefully in the proof of our main result, Theorem E.1.3, the following lemma will be very helpful for us. Again, that result is probably folklore and the proof of that lemma is very easy, but it can be found for example in [56, Lemma 3.8.].

Lemma E.2.3. *Let G be a connected claw-free graph and S be a minimal vertex separator in G . For every vertex $s \in S$ and every component K of $G - S$, the graph $G[N(s) \cap V(K)]$ is complete.*

The next lemma is a structural result for locally finite claw-free graphs about vertex sets separating some finite vertex set from all ends of the graph. This result forms the backbone for the proof of the main result of this article.

Lemma E.2.4. [56, Lemma 3.10] *Let G be an infinite, locally finite, connected, claw-free graph and X be a finite vertex set of G such that $G[X]$ is connected. Furthermore, let $\mathfrak{S} \subseteq V(G)$ be a finite minimal vertex set such that $\mathfrak{S} \cap X = \emptyset$ and every ray starting in X has to meet \mathfrak{S} . Then the following holds:*

- (1) $G - \mathfrak{S}$ has $k \geq 1$ infinite components K_1, \dots, K_k and the set \mathfrak{S} is the disjoint union of minimal vertex separators S_1, \dots, S_k in G such that for every i with $1 \leq i \leq k$ each vertex in S_i has a neighbour in K_j if and only if $j = i$.
- (2) $G - \mathfrak{S}$ has precisely one finite component K_0 . This component contains all vertices of X and every vertex of \mathfrak{S} has a neighbour in K_0 .

Given a graph G , a finite vertex set X and a set \mathfrak{S} all as in Lemma E.2.4 we shall call \mathfrak{S} an X -umbrella.

The following, last lemma of this section is the tool we use to verify Hamiltonicity for locally finite graphs in this paper.

Lemma E.2.5. [56, Lemma 3.11] *Let G be an infinite, locally finite, connected graph and $(C^i)_{i \in \mathbb{N}}$ be a sequence of cycles of G . Now G is Hamiltonian if there exists an integer $k_i \geq 1$ for every $i \geq 1$ and vertex sets $M_j^i \subseteq V(G)$ for every $i \geq 1$ and j with $1 \leq j \leq k_i$ such that the following is true:*

- (1) For every vertex v of G , there exists an integer $j \geq 0$ such that $v \in V(C^i)$ holds for every $i \geq j$.
- (2) For every $i \geq 1$ and j with $1 \leq j \leq k_i$, the cut $\delta(M_j^i)$ is finite.
- (3) For every end ω of G , there is a function $f : \mathbb{N} \setminus \{0\} \rightarrow \mathbb{N}$ such that the inclusion $M_{f(j)}^j \subseteq M_{f(i)}^i$ holds for all integers i, j with $1 \leq i \leq j$ and the equation $M_\omega := \bigcap_{i=1}^{\infty} \overline{M_{f(i)}^i} = \{\omega\}$ is true.
- (4) $E(C^i) \cap E(C^j) \subseteq E(C^{j+1})$ holds for all integers i and j with $0 \leq i < j$.
- (5) The equations $E(C^i) \cap \delta(M_j^p) = E(C^p) \cap \delta(M_j^p)$ and $|E(C^i) \cap \delta(M_j^p)| = 2$ hold for each triple (i, p, j) which satisfies $1 \leq p \leq i$ and $1 \leq j \leq k_p$.

E.3. Examples of graph meeting the criteria of

Theorem E.1.3

In this section we state examples of infinite locally finite 2-connected claw-free graphs where each induced paw satisfies $\varphi(a_1, b_i)$ for some $i \in \{1, 2\}$. While the class of claw-free and net-free graphs, which we considered in the first paper of this series [57], allows only graphs with at most two ends, the graphs in this paper have a bigger variety. We shall give examples of graphs with an arbitrary, but finite number of ends, with \aleph_0 many and with 2^{\aleph_0} many ends.

However, before we focus on these examples, we prove another proposition. This result tells us that we cannot try to extend Hamiltonicity results about locally finite 2-connected claw-free and paw-free graphs as such graphs do not exist.

Proposition E.3.1. *Every infinite locally finite connected claw-free graph containing a cycle, also contains an induced paw.*

Proof. Let G be a graph as in the statement and let C be a cycle of G . Pick a ray $R = r_0 r_1 r_2 \dots$ that is distance increasing w.r.t. $V(C)$, which exists by Lemma E.2.1. Recall that $r_0 \in V(C)$ holds. Now $G[r_0, r_0^+, r_0^-, r_1]$ is no induced claw since G is claw-free. If $r_0^- r_1 \in E(G)$ or $r_0^+ r_1 \in E(G)$, then this would yield a K_3 , say $G[r_0^+, r_0, r_1]$. Since R is distance increasing w.r.t. $V(C)$, we know that $G[r_0^+, r_0, r_1, r_2]$ is an induced paw.

The only other possibility to avoid $G[r_0, r_0^+, r_0^-, r_1]$ being an induced claw is $r_0^- r_0^+ \in E(G)$ but $r_0^- r_1, r_0^+ r_1 \notin E(G)$. Then $G[r_0^-, r_0^+, r_0, r_1]$ forms an induced paw. \square

Before we come to the examples, in which the k -blow-up operation is involved, let us recall the definition of a k -blow-up. Given a graph G and some $k \in \mathbb{N}$, we call a graph G' a k -blow-up of G if we obtain G' from G by replacing each vertex of G by a clique of size k , where two vertices of G' are adjacent if and only if they are either both from a common such clique, or the original corresponding vertices were adjacent in G .

In three examples we now state graphs that meet the criteria of Theorem E.1.3 by describing an initial graph, from which we then take the line graph and then a k -blow-up. Before we state the first example, let us fix some notation. Let $n \in \mathbb{N}$

and let S_n denote the infinite tree where each vertex but one has degree 2 and the other vertex has degree n .

Example E.3.2. To find a graph with an arbitrary, but finite number of ends, consider the k -blow-up of the line graph of S_n where $k \geq 2$ and $n \geq 3$. To briefly describe this graph in other words: It is the k -blow-up of the graph formed by a complete graph on n vertices $V(K_n) = \{v_1, \dots, v_n\}$ with pairwise disjoint rays R_i starting at the v_i , see Figure E.2.

It is immediate that such graphs are 2-connected and claw-free. To check that every induced paw satisfies $\varphi(a_1, b_i)$ for some $i \in \{1, 2\}$ is also straightforward, so we leave this to the reader.




Figure E.2.: A graph whose k -blow-up for $k \geq 2$ has precisely $n \in \mathbb{N}$ ends and meets the conditions of Theorem E.1.3.

Obviously, we cannot use the construction from Example E.3.2 to obtain suitable graphs with infinitely many ends while staying locally finite. However, we can extend the idea from the previous example to get such graphs with infinitely many ends. For every $n \in \mathbb{N}$ let D_n denote the infinite tree where all vertices have degree 2 except for a set of vertices that induces in D_n a double ray all of whose vertices have degree n within D_n .

Example E.3.3. Consider the k -blow-up of the line graph of D_n for some $k \geq 2$ and $n \geq 3$ (cf. Figure E.3). Again it is easy to verify that such graphs satisfy the conditions of Theorem E.1.3. Furthermore, it is immediate that D_n has precisely \aleph_0 many ends for $n \geq 3$ and, hence, so has its line graph and the k -blow-up of the line graph.

Let us proceed to the third example describing graphs that meet the conditions of Theorem E.1.3 and have precisely 2^{\aleph_0} many ends. Again we fix some notation




Figure E.3.: The line graph of D_3 , whose k -blow-up for $k \geq 2$ has precisely \aleph_0 many ends and meets the conditions of Theorem E.1.3.

before. For every $n \in \mathbb{N}$ let T_n denote the infinite tree where each vertex has degree n .

Example E.3.4. Consider the k -blow-up of the line graph T_n for some $k \geq 2$ and $n \geq 3$ (cf. Figure E.4). As before, verifying that such graphs satisfy the conditions of Theorem E.1.3 is easy. Furthermore, T_n has precisely 2^{\aleph_0} many ends for $n \geq 3$. Therefore, its line graph and the k -blow-up of the line graph have that many ends as well.




Figure E.4.: The line graph of T_3 , whose k -blow-up for $k \geq 2$ has precisely 2^{\aleph_0} many ends and meets the conditions of Theorem E.1.3.

E.4. Proof of Theorem E.1.3

In this section we shall prove our main result, Theorem E.1.3, which is an extension of Theorem E.1.2 to locally finite graphs. Let us describe the general idea for the proof, which has already been successful for other Hamiltonicity results for locally finite claw-free graphs involving local conditions [55, 56]. In the end we want to apply Lemma E.2.5. However, in order to do that, we have to carefully construct suitable cycles and cuts as described in that lemma. The engine of our proof is the condition about induced paws as in Theorem E.1.3. This allows us to extend any cycle by one or two vertices neighbouring the cycle in a very controlled way keeping most edges of the initial cycle untouched, which is captured in Lemma E.4.1.

Then the first big step towards the proof of Theorem E.1.3 is to extend an arbitrary cycle C with respect to a $V(C)$ -umbrella \mathfrak{S} to contain all vertices of K_0 (as defined in Lemma E.2.4), without containing anything from \mathfrak{S} . This happens in Lemma E.4.2. The next step is to carefully extend the cycle into each component K_i while containing all vertices up to the third neighbourhood of S_i , but only precisely two vertices of each S_i (cf. Lemma E.2.4). This step is rather crucial and achieved in Lemma E.4.4.

From this point on we shall not only keep track of a cycle, but also of suitable cuts, each more or less resembling $\delta(V(K_i))$, which our cycle intersects precisely twice. We shall refer to this as the (\star) - *condition*. In the remaining lemma, Lemma E.4.6, we incorporate all remaining vertices of \mathfrak{S} while maintaining the (\star) - condition. The key idea here is to dynamically change the cuts as we extend the cycle. By iterating this whole procedure, we shall get a sequence of cycles and cuts which allows us to apply Lemma E.2.5.

Now let us start with the lemma which allows us to extend any cycle C within a graph as in Theorem E.1.3 to incorporate a vertex $v \in N(C)$ without altering many edges of C .

Lemma E.4.1. *Let G be a 2-connected, locally finite claw-free graph such that every induced paw of G satisfies $\phi(a_1, b_i)$ for some $i \in \{1, 2\}$. Let C be a cycle in G and v be a vertex in $N(C)$. Then there exists a cycle C' and a vertex $w \in N(C)$ such that $V(C) \cup \{v\} \subseteq V(C') \subseteq V(C) \cup \{v, w\}$ holds. Furthermore, if $xy \in E(C) \Delta E(C')$, then $x \in N_2(v)$ or $y \in N_2(v)$ holds.*

Proof. Let G be a graph as in the statement of the lemma and C be a cycle in G .

Let v be a vertex in $N(C)$. We prove the statement by a short case distinction. The three cases are depicted in Figure E.5:

Case 1. *There is a vertex $u \in V(C) \cap N(v)$ such that $u^+v \in E(G)$ or $u^-v \in E(G)$.*

Without loss of generality let us say $u^+v \in E(G)$. We obtain C' simply by exchanging uu^+ with uv and vu^+ , which completes Case 1.

Case 2. *For all vertices $u \in V(C) \cap N(v)$ we have that $u^+v, u^-v \notin E(G)$.*

By the claw-freeness we know that $u^-u^+ \in E(G)$, but this means $G[u^-, u, u^+, v]$ is an induced paw. Hence we get that u^- or u^+ shares a neighbour w with v in $V(G) \setminus \{u^-, u, u^+, v\}$, say without loss of generality u^+ . Now we distinguish two subcases:

Subcase 2.1. *There exists such a neighbour $w \notin V(C)$.*

In this situation we simply obtain C' by replacing uu^+ by uv, vw and wu^+ , completing Subcase 2.1.

Subcase 2.2. *All such w lie on $V(C)$.*

Since we are in the second case, we get that $w^-w^+ \in E(G)$. Hence we can get a new cycle C' by replacing the segment uCw^+ of C by $uvwu^+Cw^-w^+$. This completes our case distinction.

Finally, note that in each case changing C to C' does not alter edges whose endvertices have distance at most 2 from v . \square

Given the notation of Lemma E.4.1, we call the cycle C' a v -extension of the cycle C of type (1) if C' is formed as in Case 1. Similarly, we call C' a v -extension of C of type (2.1) or of type (2.2) if C' is formed as in Subcase 2.1 or 2.2, respectively. For a v -extension we also call v the *target* and u its *base*. We call a cycle D an *extension* of a cycle C if we obtain D from successively performing v -extensions (with possibly several different targets v) of C of any type. Whenever we talk about a v -extension of type (2.1) or (2.2), we shall denote by w the same vertex as in the proof of Lemma E.4.1. We call the edge we exclude from the cycle C by forming a v -extension that has the base as one of its endvertices the *foundation* (of the extension), see Figure E.5.

Given a cycle C and a $V(C)$ -umbrella \mathfrak{S} (cf. Lemma E.2.4), we now show that we can extend C to contain all of K_0 , but nothing from \mathfrak{S} .




Figure E.5.: The three types of v -extensions of the cycle C by replacing grey edges by dashed ones as occurring in the proof of Lemma E.4.1.

Lemma E.4.2. *Let G be an infinite, locally finite, 2-connected, claw-free graph such that every induced paw of G satisfies $\phi(a_1, b_i)$ for some $i \in \{1, 2\}$. Let C be a cycle and \mathfrak{S} a $V(C)$ -umbrella. Then there exists a cycle C' which is an extension of C such that $V(C') = V(K_0)$ and for each $e = xy \in E(C)$ with $x, y \in V(K_0) \setminus N_3(N(C))$ we have that $e \in E(C')$.*

Proof. We successively perform v -extensions of extensions of C with targets in $V(K_0)$ where, for a fixed target v , we always try perform a v -extension of type (1). If that is not possible we try to use one of type (2.1), and if that is not possible either we use one of type (2.2). The only case which might lead to a problem for the desired equality $V(C') = V(K_0)$ is by using v -extensions of type (2.1) since we not only incorporate the target v but one further vertex $w \in N(C)$. So let us assume we perform a v -extension of type (2.1) where w lies in \mathfrak{S} . Thus, w lies in a separator S_j as defined in Lemma E.2.4. This means that v and y , where uy is the foundation of the v -extension, are connected since the neighbourhood of w in the component of $G - S_j$ containing K_0 forms a clique, due to Proposition E.2.3. Hence we did not need to incorporate v via an extension of type (2.1), but could have done it via type (1). So our process terminates and yields cycle C' such that $V(C') = V(K_0)$ since K_0 is finite and connected due to Lemma E.2.4.

Note that for each $e = xy \in E(C)$ such that $x, y \in V(K_0) \setminus N_3(N(C))$ we have that $e \in E(C')$ holds due to Lemma E.4.1. Hence, the cycle C' is as desired. \square

Before we move on, we first give a definition used to capture how to carefully extend the cycle obtained from Lemma E.4.2 further into the infinite components

K_1, \dots, K_k as defined in Lemma E.2.4 in a convenient way.

Definition E.4.3. Let G be an infinite, locally finite, connected, claw-free graph, C be a cycle of G and \mathfrak{S} be a $V(C)$ -umbrella. Furthermore, let k , S_j and K_j be defined as in Lemma E.2.4. Now we call a tuple (D, M_1, \dots, M_k) *promising* if the following hold for every $j \in \{1, \dots, k\}$:

- (1) D is a cycle.
- (2) $M_j := S_j \cup V(K_j)$.
- (3) $V(K_0) \cup \bigcup_{1 \leq i \leq k} (N_3(S_i) \cap V(K_i)) \subseteq V(D)$
- (4) $|E(D) \cap \delta(M_j)| = 2$. ((\star) - condition)

Note that the definition of a promising tuple is defined relative to a certain umbrella. We shall not mention to which if the context makes this clear.

In the next two lemmas we show how to carefully extend a cycle in two steps to incorporate at least $N_3(\mathfrak{S})$ while respecting the (\star) - condition. Now we first prove that promising tuples exist.

Lemma E.4.4. *Let G be an infinite locally finite, connected, claw-free graph such that every induced paw of G satisfies $\phi(a_1, b_i)$ for some $i \in \{1, 2\}$. Furthermore, let C be a cycle of G and \mathfrak{S} be a $V(C)$ -umbrella. Also, let k , S_j and K_j be defined as in Lemma E.2.4.*

Then there is a promising tuple (D, M_1, \dots, M_k) such that for all $e = xy \in E(C)$ with $x, y \in V(K_0) \setminus N_3(N(C))$ we have $e \in E(D)$.

Proof. Let C' be an extension of C such that $V(C') = V(K_0)$, which exists by Lemma E.4.2. We shall successively form v -extensions with targets in \mathfrak{S} or replace edges ab in some S_j with a - b paths whose inner vertices lie in K_i until we have a cycle C'' containing some vertex of each K_i , precisely two vertices of each S_i , but no edge from any $G[S_i]$, while respecting the (\star) - condition. Let us first prove that having such a cycle C'' suffices to prove this lemma.

Assume we have such a cycle C'' , which fulfills the (\star) - condition. We now show that we can include any finite vertex set X from an arbitrary K_i for $i \in \{1, \dots, k\}$, hence especially $N_m(S_i) \cap V(K_i)$ for any $m \in \mathbb{N}$ while maintaining the (\star) - condition. For this we only target $v \in V(K_i)$ along a finite connected subgraph

within K_i containing X . Hence, the bases of our extensions will always lie in M_i . Whenever possible we include the target via a type (1) extension. If this is not possible, we try to incorporate the target by a type (2.1) extension and if this is also not possible, we use a type (2.2) extension.

In case we perform a v -extension of type (1), no altered edge has an endvertex in $V(K_0)$. Thus, this does not affect the (\star) - condition. If, for a v -extension of type (2.1), the foundation lies in M_i , then again we do not alter any edge with an endvertex in $V(K_0)$. The other case is a foundation $ur \in \delta(M_i)$ for $r \in \{u^+, u^-\}$. But this means that we exclude ur from the cut $\delta(M_i)$ and add rw to it. So again we maintain the (\star) - condition. Finally let us consider a v -extension of type (2.2). For $w \in V(K_i)$, we do not affect the (\star) - condition. So let us assume w lies in the separator S_i . As for an extension of type (2.1), the removal of the foundation ur for $r \in \{u^+, u^-\}$ and incorporating the edges uv, vw and wr does not change the size of the cut. It remains to check that excluding w^-w and ww^+ and adding w^-w^+ does not violate the (\star) - condition either. We check this by a short case distinction:

Case A. $w^-, w^+ \in M_i$.

In this case the edges w^-w, ww^+ and w^-w^+ all lie in M_i , so the (\star) - condition is maintained.

Case B. $w^-, w^+ \in V(G) \setminus M_i$.

This case cannot happen since we assumed that C'' contains a vertex of K_i and precisely two vertices from S_i . However, C'' would cross $\delta(M_i)$ via the edges w^-w, ww^+ without entering K_i . So the (\star) - condition would already be violated by C'' ; a contradiction.

Case C. $w^- \in M_i$ but $w^+ \notin M_i$ (or vice versa).

In this case our v -extension does not intersect $\delta(M_i)$ with the edge w^+w anymore, but with the edge w^-w^+ . Hence, the (\star) - condition is again maintained. This completes our case analysis and shows the existence of a desired tuple (D, M_1, \dots, M_k) if the cycle C'' exists. Note that for each $e = xy \in E(C)$ with $x, y \in V(K_0) \setminus N_3(N(C))$ we have $e \in E(D)$, as shown in Lemma E.4.1.

Hence it remains to show that a cycle C'' exist, i.e., a cycle containing all of $V(K_0)$, at least some vertex of K_i and precisely two vertices from each S_i , but no

edge from any $G[S_i]$, while respecting the (\star) - condition. Say we already have a cycle Z satisfying the following for a (possibly empty) subset $I \subseteq \{1, \dots, k\}$.

- (i) $V(K_0) \subseteq V(Z)$.
- (ii) $V(Z) \cap V(K_i) \neq \emptyset$ for every $i \in I$.
- (iii) $|V(Z) \cap S_i| = 2$, but $E(Z) \cap E(G[S_i]) = \emptyset$ for every $i \in I$.
- (iv) $|E(Z) \cap \delta(M_i)| = 2$ for every $i \in I$.
- (v) $V(Z) \cap M_j = \emptyset$ for every $j \in \{1, \dots, k\} \setminus I$.

Note that C' meets these five conditions with $I = \emptyset$. Next we show how to obtain a cycle meeting all these five conditions for a superset of I . We form a v -extension Z' of Z with target $v \in S_j$ for some $j \in \{1, \dots, k\} \setminus I$. Again we analyse the situation via a case distinction.

Case 1. Z' can be formed of type (1).

Without loss of generality say the foundation of Z' is uu^+ . If $uu^+ \in E(K_0)$, then condition (iv) from above is still maintained. If $uu^+ \in \delta(M_i)$ for some $i \in I$, the size of the intersection of the cycle with the cut $\delta(M_j)$ stays the same, but the edge uv meets it instead of uu^+ . In both cases Z' now fulfills condition (iv) for $I \cup \{j\}$. Next we form a v' -extension Z'' of Z' where $v' \in N(v) \cap V(K_j)$. Such a v' exists since S_j is a minimal vertex separator by Lemma E.2.4. But now v' can only be included as in type (2.1), since extensions of type (1) and (2.2) need two neighbours of the target on the cycle from which the extension is formed. Hence, w lies in S_j , again we do not change the size of any intersection $E(Z') \cap \delta(M_i)$ with $i \in \{1, \dots, k\}$ and we do not incorporate any edge within $G[S_j]$. So Z'' satisfies all five condition above with respect to $I \cup \{j\}$. This completes the analysis of the first case.

Note that in the situation where we cannot perform a v -extension of Z of type (1) we can, by the proof of Lemma E.4.1, fix a desired base $u \in N(v)$ on Z in advance.

Case 2. Z' cannot be formed of type (1), but of type (2.1) with base $u \in V(K_0)$.

Now we further distinguish two subcases.

Subcase 2.1. $w \in S_j$.

In this case we have incorporated two vertices of S_j , meet condition (iv) also for j , but we use the edge vw from $G[S_j]$. As S_j is a minimal vertex separator by Lemma E.2.4, both of these vertices have a neighbour in K_j . Using that K_j is connected, we can find a v - w path P whose inner vertices lie in K_j . Now we modify Z' by replacing vw with P . The resulting cycle is as desired.

Subcase 2.2. $w \in S_i$ for some $i \neq j$.

For the sake of clarity we again distinguish two further subcases depending on where the foundation of the v -extension lies. Let us denote the foundation by ur where $r \in \{u^-, u^+\}$.

Subcase 2.2.1. $r \notin S_i$.

This is not possible since the neighbourhood of w in each component of $G - S_i$ induces a clique. Since we demand for u as the base of Z' to lie in $V(K_0)$, we know that both u and r lie in the same component of $G - S_i$, namely the one different from K_i . Hence, $vr \in E(G)$ and we could have included v as in type (1) against our assumption.

Subcase 2.2.2. $r \in S_i$.

In this case we know that $i \in I$. So $ur \in \delta(M_i)$ holds. By condition (iii) we know that $|V(Z) \cap S_i| = 2$. Without loss of generality, say $V(Z) \cap S_i = \{r, s\}$ for some other vertex $s \in S_i$. As Z satisfies condition (i), (ii) and (iii), we know that Z contains precisely one path Q with endvertices in S_i and all inner vertices in K_i . Hence, Q must be an s - r -path. Now form the v -extension Z' of Z of type (2.1), then delete all inner vertices of Q and r from Z' and replace it by an s - w -path all whose inner vertices lie in K_i . The resulting cycle contains only v from S_j and we can proceed as in Case 1. This completes the analysis for the second case.

Case 3. Z' cannot be formed of type (1), but of type (2.2) with base $u \in V(K_0)$.

Note that in this case $w \in V(Z)$ and w does lie in S_i with $i \in I$ as due to condition (ii) of Z the edge w^-w^+ would cross the separator S_i ; a contradiction. Hence, $w \in V(K_0)$ holds. The only possibility that $E(Z') \cap \delta(M_\ell) \neq E(Z) \cap \delta(M_\ell)$ for some $\ell \in I$ is that one of w^+, w^- is contained in S_ℓ while the other vertex is contained in the component of $G - S_\ell$ different from K_ℓ . Say $w^+ \in S_\ell$ holds. Note

that not both, w^+ and w^- can lie in S_ℓ since $V(C) \subseteq V(K_0) \cap V(Z')$. Hence, $|E(Z') \cap \delta(M_i)| = 2$ holds for all $i \in I \cup \{j\}$ and Z' contains precisely v from S_j . To complete the argument for this case we can now proceed as in Case 1. This completes our case analysis and shows the existence of the desired cycle C'' .

Finally, note that the constructed promising tuple (D, M_1, \dots, M_k) satisfies that for each $e = xy \in E(C)$ with $x, y \in V(K_0) \setminus N_3(N(C))$ we have $e \in E(D)$. We either performed v -extensions, which cause no problems as checked in Lemma E.4.1. Apart from that we included or excluded paths from $V(K_i) \cup S_i$ in Subcase 2.1 and Subcase 2.2.2, which does not affect edges from C . \square

From Lemma E.4.4 we do not necessarily get a cycle that contains all vertices of \mathfrak{S} . However, in order to apply Lemma E.2.5 we have to incorporate all these vertices while maintaining constraints for suitable cuts $\delta(M_i)$. With the next lemma we shall achieve this. The key idea is to start from a promising tuple and incorporate the remaining vertices while dynamically changing the vertex sets M_i to maintain the constraints for the cuts $\delta(M_i)$. We shall encode our desired objects via the following definition before we move on to the next lemma.

Definition E.4.5. Let $G = (V, E)$ be an infinite locally finite, connected, claw-free graph, C be a cycle of G and \mathfrak{S} be a $V(C)$ -umbrella. Furthermore, let k , S_j and K_j be defined as in Lemma E.2.4. Now we call a tuple (D, M_1, \dots, M_k) *good* if the following properties hold for every $j \in \{1, \dots, k\}$:

- (1) D is a cycle of G which contains $V(K_0)$ and $S_j \cup (N_3(S_j) \cap V(K_j))$.
- (2) $V(K_j) \setminus N(S_j) \subseteq M_j \subseteq V(K_j) \cup \mathfrak{S} \cup N(\mathfrak{S})$.
- (3) $|E(D) \cap \delta(M_j)| = 2$. ((\star) - condition)

Lemma E.4.6. *Let G be an infinite locally finite, connected, claw-free graph, such that every induced paw of G satisfies $\phi(a_1, b_i)$ for some $i \in \{1, 2\}$. Furthermore, let C be a cycle of G and \mathfrak{S} be a $V(C)$ -umbrella, where k , S_j and K_j be defined as in Lemma E.2.4. Let (D, M_1, \dots, M_k) be a promising tuple for the $V(C)$ -umbrella \mathfrak{S} .*

Then there is a good tuple (D', N_1, \dots, N_k) (for the same umbrella \mathfrak{S}) and for each $e = xy \in E(C)$ such that $x, y \in V(K_0) \setminus N_3(N(C))$ we have $e \in E(D')$.

Proof. We prove the statement via a recursive construction always performing v -extensions with targets $v \in \mathfrak{S} \setminus V(D)$ until we eventually obtain the cycle D' of

our desired good tuple. We initialise this construction with the promising tuple (D, M_1, \dots, M_k) . Let us denote by pD the cycle obtained after p many performed v -extensions. During this process we also alter the sets M_i . Let pM_i denote the corresponding vertex set after having performed p many v -extensions for every $i \in \{1, \dots, k\}$. We note that this construction process will eventually terminate, say after $z \leq |\mathfrak{S} \setminus V(D)|$ many steps, since in each step we add at least one vertex from the finite set $\mathfrak{S} \setminus V(D)$ and do not exclude any vertices at all.

Whenever possible we include our target vertices via a type (1) extension. If this is not possible, we try to incorporate them via a type (2.1) extension and if this is also not possible, we use an extension of type (2.2).

Now suppose we have already constructed pD and pM_i for all $i \in \{1, \dots, k\}$ such that the (\star) - condition is maintained and let $v \in S_i \setminus V({}^pD)$ be our next target.

Case 1. *There exists a v -extension ${}^{p+1}D$ of pD of type (1).*

Let xy be the foundation of ${}^{p+1}D$. Recall that this edge gets substituted by xv and vy by forming the extension. We define for all $r \in \{1, \dots, k\}$ the set

$${}^{p+1}M_r = \begin{cases} {}^pM_r \setminus \{v\} & \text{if } x, y \notin {}^pM_r \\ {}^pM_r \cup \{v\} & \text{else} \end{cases}$$

This ensures that the (\star) - condition is still maintained for ${}^{p+1}D$ and each ${}^{p+1}M_r$.

Case 2. *There exists no v -extension of pD of type (1), but of type (2.1).*

Let ux be the foundation of ${}^{p+1}D$ where u is the base of the extension. Now we define for all $r \in \{1, \dots, k\}$:

$${}^{p+1}M_r = \begin{cases} {}^pM_r \cup \{v, w\} & \text{if } u, x \in {}^pM_r \\ {}^pM_r \setminus \{w\} \cup \{v\} & \text{if } u \in {}^pM_r \text{ and } x \notin {}^pM_r \\ {}^pM_r \setminus \{v\} \cup \{w\} & \text{if } u \notin {}^pM_r \text{ and } x \in {}^pM_r \\ {}^pM_r \setminus \{v, w\} & \text{if } u, x \notin {}^pM_r \end{cases}$$

Again we note that these cases respect the (\star) - condition: If u and x lie within one pM_r , the cycle ${}^{p+1}D$ meets $\delta({}^{p+1}M_r)$ still twice. If the $ux \in \delta({}^pM_s)$, we exclude ux from the corresponding intersection but add precisely vw to it.

Case 3. *There exists no v -extension of pD of type (1), but of type (2.2).*

Again let ux be the foundation of ${}^{p+1}D$. Here we define the new sets for $r \in \{1, \dots, k\}$ as in the case of an type (2.1) v -extension:

$${}^{p+1}M_r = \begin{cases} {}^pM_r \cup \{v, w\} & \text{if } u, x \in {}^pM_i \\ {}^pM_r \setminus \{w\} \cup \{v\} & \text{if } u \in {}^pM_i \text{ and } x \notin {}^pM_i \\ {}^pM_r \setminus \{v\} \cup \{w\} & \text{if } u \notin {}^pM_i \text{ and } x \in {}^pM_i \\ {}^pM_r \setminus \{v, w\} & \text{if } u, x \notin {}^pM_i \end{cases}$$

To verify that the new cuts $\delta({}^{p+1}M_r)$ together with ${}^{p+1}D$ still satisfy the (\star) -condition, we first note that we delete three edges from the cycle, namely: w^-w, ww^+ and ux . But we include the edges uv, vw, wx, w^-w^+ . Regarding the intersection of the cycle and the cuts, the exclusion of w^-w and ww^+ precisely cancels out the effect of adding w^-w^+ . The same holds for excluding ux and adding uv, vw and wx as we already in Case 2. This completes the recursive definition of the cuts for the good tuple.

We now define $D' := {}^zD$ and $N_r := {}^zM_r$ for every $r \in \{1, \dots, k\}$. It remains to check that (D', N_1, \dots, N_k) satisfies all properties of a good tuple. We already argued that the tuple respects the (\star) -condition.

Regarding property (1) note that we we started the recursive definition with a cycle D that is part of a promising cycle for the $V(C)$ -umbrella \mathfrak{S} . So D already contained all vertices from $V(K_0) \cup (N_3(S_j) \cap V(K_j))$ for every $j \in \{1, \dots, k\}$, and we just included the remaining vertices from \mathfrak{S} when forming D' .

Note for property (2) that we have only added or excluded vertices in $\mathfrak{S} \cup N(\mathfrak{S})$ in each step when changing the vertex sets for our cuts. Also note that the vertex set M_j we started with were defined as $M_j = S_j \cup V(K_j)$. \square

We now combine the previous lemmas to prove the main theorem. Let us restate the statement of theorem first.

Theorem E.1.3. *Let G be a locally finite, 2-connected, claw-free graph. If every induced paw of G satisfies $\phi(a_1, b_i)$ for some $i \in \{1, 2\}$, then G is Hamiltonian.*

Proof. Let G be a graph as in the statement of the theorem. We may assume G to be infinite by Theorem E.1.2. We shall recursively construct a sequence of good tuples $(C^i, M_1^i, \dots, M_{k(i)}^i)$ where each tuple $(C^{i+1}, M_1^{i+1}, \dots, M_{k(i+1)}^{i+1})$ is defined with respect to a $V(C^i)$ -umbrella for every $i \in \mathbb{N}$ as follows.

Start with an arbitrary cycle A in G . This is possible, since G is 2-connected. Let C^0 be an extension of A with $N_3(V(A)) \subseteq V(C^0)$.

Next suppose the cycle C^i has already been defined up to some $i \in \mathbb{N}$:

- Let \mathfrak{S}^{i+1} be a $V(C^i)$ -umbrella, K_i^{i+1} and S_i^{i+1} be defined as in Lemma E.2.4 and let $k : \mathbb{N} \rightarrow \mathbb{N}$ be the function such that \mathfrak{S}^{i+1} leaves precisely $k(i+1)$ infinite components.
- Let $(D^{i+1}, Y_1^{i+1}, \dots, Y_{k(i+1)}^{i+1})$ be a promising tuple we get by applying Lemma E.4.4 with the cycle C^i and the $V(C^i)$ -umbrella \mathfrak{S}^{i+1} .
- Then set $(C^{i+1}, M_1^{i+1}, \dots, M_{k(i+1)}^{i+1})$ to be a good tuple we get from Lemma E.4.6 applied with the promising tuple $(D^{i+1}, Y_1^{i+1}, \dots, Y_{k(i+1)}^{i+1})$, the cycle C^i and the $V(C^i)$ -umbrella \mathfrak{S}^{i+1} .

We now conclude the proof by verifying that we can apply Lemma E.2.5.

For condition (1) of Lemma E.2.5 we have to show that for every $v \in V(G)$ there is some $j \in \mathbb{N}$ such that $v \in V(C^i)$ for every $i \geq j$. This holds since every $v \in V(G)$ has finite distance to $V(C^0)$. So it follows from property (1) of good tuples.

For condition (2) of Lemma E.2.5 we need to prove that for every $i \geq 1$ and j with $1 \leq j \leq k(i)$, the cut $\delta(M_j^i)$ is finite. Since G is locally finite, it suffices to show that M_j^i has a finite neighbourhood. Due to property (2) of good tuples we know that $N(M_j^i) \subseteq \mathfrak{S}^i \cup N_2(\mathfrak{S}^i)$. Since \mathfrak{S}^i is a finite set and G is locally finite, we obtain that $N(M_j^i)$ is finite.

Regarding condition (3) of Lemma E.2.5 we need to prove for every end ω of G the existence of a function $f : \mathbb{N} \setminus \{0\} \rightarrow \mathbb{N}$ such that the $M_{f(j)}^j \subseteq M_{f(i)}^i$ holds for all integers i, j with $1 \leq i \leq j$ and that the equation $M_\omega := \bigcap_{i=1}^{\infty} \overline{M_{f(i)}^i} = \{\omega\}$ is true. To verify this let us fix an arbitrary end ω of G . We first define the desired function f . Note that for each $i \geq 1$ we know that K_0^i and \mathfrak{S}^i are finite by Lemma E.2.4. Hence each ω -ray has a tail in K_ℓ^i for some $\ell \in \{1, \dots, k(i)\}$. Now set $f(i) := \ell$.

Let us now check that $M_{f(j)}^j \subseteq M_{f(i)}^i$ holds for all $1 \leq i \leq j$. By properties (1) and (2) of a good tuple, it is easy to see that $V(K_{f(j)}^j) \subseteq V(K_{f(i)}^i)$ holds for all $1 \leq i \leq j$ and similarly $M_{f(j)}^j \subseteq M_{f(i)}^i$.

To verify condition (3) it remains to show that $M_\omega := \bigcap_{i=1}^{\infty} \overline{M_{f(i)}^i} = \{\omega\}$ is true. First note that no vertex can lie in this intersection since each vertex will eventually be contained in K_0^t for some sufficiently large $t \in \mathbb{N}$, and hence not in any set M_i^ℓ for all $\ell \geq t$. Furthermore, the definition of f already ensures $\omega \in M_\omega$. So let us consider some end ω' of G distinct from ω . Let S be a finite vertex set separating ω from ω' . Since S is finite, we know that $S \subseteq V(K_0^t)$ holds for some sufficiently large $t \in \mathbb{N}$. Hence $\omega' \notin \overline{K_{f(t)}^t}$ and similarly $\omega' \notin \overline{M_{f(t)}^t}$. So we obtain the desired conclusion $\omega' \notin M_\omega$.

Now we focus on condition (4) of Lemma E.2.5. There we need to verify the inclusion $E(C^i) \cap E(C^j) \subseteq E(C^{j+1})$ for all integers i and j with $0 \leq i < j$. This holds since we checked in Lemma E.4.2, Lemma E.4.4 and Lemma E.4.6 that whenever we change a cycle C to a cycle C' each edge $e = xy \in E(C)$ with $x, y \in V(K_0) \setminus N_3(N(C))$ lies also in $E(C')$. So by property (1) of good tuples the sequence of cycles $(C^i)_{i \in \mathbb{N}}$ satisfies condition (4) of Lemma E.2.5.

Finally, let us verify condition (5) of Lemma E.2.5. So we have to show that the equations $E(C^i) \cap \delta(M_j^p) = E(C^p) \cap \delta(M_j^p)$ and $|E(C^i) \cap \delta(M_j^p)| = 2$ hold for each triple (i, p, j) which satisfies $1 \leq p \leq i$ and $1 \leq j \leq k(p)$. This, however, immediately follows from the satisfied previous (4) and the fact that good tuples satisfy the (\star) - condition.

Hence, we can apply Lemma E.2.5, which proves the Hamiltonicity of G and concludes our proof. \square

F. Hamilton-laceable bi-powers of locally finite bipartite graphs

F.1. Introduction

Recently, Li [77] restricted the power operation for graphs to preserve bipartiteness by defining another operation, called the *bi-power* of a graph, which is defined as follows:

Let G be a graph and $k \in \mathbb{N}$. For two vertices $x, y \in V(G)$ let $\text{dist}_G(x, y)$ denote the distance in G between x and y . Then the k -th *bi-power* G_B^k of G is defined as:

$$V(G_B^k) := V(G),$$

$$E(G_B^k) := \{xy \mid \text{dist}_G(x, y) \text{ is odd and at most } k \text{ where } x, y \in V(G)\}.$$

Note that $G_B^2 = G_B^1 = G$. Also, the k -th bi-power of a bipartite graph is still bipartite.

Two classical Hamiltonicity results for finite graphs that involve the (usual) square and cube of a graph are the following ones by Fleischner and Sekanina.

Theorem F.1.1. [34] *The square of any finite 2-connected graph is Hamiltonian.*

Theorem F.1.2. [94] *The cube of any finite connected graph on at least 3 vertices is Hamiltonian.*

Both of these results have been extended to *locally finite* infinite graphs, i.e. graphs where every vertex has finite degree, by Georgakopoulos [38, Thm. 3, Thm. 5]. The crucial conceptual starting point of Georgakopoulos' work is the topological approach initiated by Diestel and Kühn [27, 28]. They defined infinite cycles as *circles*, i.e. homeomorphic images of the unit circle $S^1 \subseteq \mathbb{R}^2$ within the Freudenthal compactification $|G|$ [24, 25] of a locally finite connected graph G . Using this notation, a *Hamilton circle* of G is a circle in $|G|$ containing all vertices of G , and we shall call G *Hamiltonian* if such a circle exists.

Motivated by Theorem F.1.1 and Theorem F.1.2 as well as by their extensions to locally finite infinite graphs, Li [77] proved Hamiltonicity results for finite as well as locally finite bipartite graphs involving the bi-power. In order to state Li's result for finite graphs we have to introduce another notation.

A finite connected bipartite graph G is called *Hamilton-laceable* if for any two vertices $v, w \in V(G)$ from different bipartition classes of G there exists a Hamilton-path whose endvertices are v and w .

Theorem F.1.3. [77, Thm. 6] *If G is a finite connected bipartite graph that admits a perfect matching, then G_B^3 is Hamilton-laceable.*

In contrast to the usual power operation and to Theorem F.1.1 and Theorem F.1.2, the statement of Theorem F.1.3 becomes false if we omit the assumption of G admitting a perfect matching as shown by Li [77]. We shall briefly discuss Li's example and adapt it to also yield an example for infinite graphs in Section F.3.

Beside Theorem F.1.3, Li proved a related result for locally finite infinite graphs in the same article [77]. However, the conclusion of the result only yields the existence of a Hamilton circle and does not speak about an adapted topological version of Hamilton-laceability.

Theorem F.1.4. [77, Thm. 5] *If G is a locally finite infinite connected bipartite graph that admits a perfect matching, then G_B^3 is Hamiltonian.*

In this paper, we extend Theorem F.1.3 to locally finite graphs by defining and using a natural topological extension of the notion of Hamilton-laceability. In order to define this, we have to state other definitions first.

Let G be a locally finite connected graph. As an analogue of a path, we define an *arc* as a homeomorphic image of the unit interval $[0, 1] \subseteq \mathbb{R}$ in $|G|$. We call a point p of $|G|$ an *endpoint* of α if 0 or 1 is mapped to p by the homeomorphism defining α . For $p, q \in |G|$, we shall briefly call an arc α a p - q *arc* if p and q are endpoints of α . Furthermore, an arc α in $|G|$ is called a *Hamilton arc* of G if it contains all vertices of G .

Now we are able to state a topological analogue of Hamilton-laceability. We call a locally finite connected bipartite graph G *Hamilton-laceable* if for any two vertices $v, w \in V(G)$ from different bipartition classes of G there exists a Hamilton arc whose endpoints are v and w , i.e. a Hamilton v - w arc.

Now we are able to state the main result of this paper.

Theorem F.1.5. *If G is a locally finite connected bipartite graph that admits a perfect matching, then G_B^3 is Hamilton-laceable.*

Clearly, as for finite graphs, our topological notion of Hamilton-laceability is stronger than Hamiltonicity since demanding the existence of a Hamilton arc for the two endvertices of an edge in the considered graph immediately yields a Hamilton circle (unless $G \neq K_2$). Hence, Theorem F.1.5 is a proper extension of Theorem F.1.4.

The structure of this paper is as follows. In Section F.2 we introduce the necessary notation, definitions and tools for the rest of the paper. In Section F.3 we briefly discuss via finite and locally finite counterexamples how some potential strengthenings of the statement of Theorem F.1.5 fail. Section F.4 starts with a brief discussion of the differences of our proof method compared to the one used by Li [77]. Afterwards, we prove a key lemma, Lemma F.4.5, which then enables us to prove the main result, Theorem F.1.5.

F.2. Preliminaries

All graphs in this paper are simple and undirected ones. Generally, we follow the graph theoretical notation from [24]. Regarding topological notions for locally finite graphs, we especially refer to [24, Ch. 8.5], and for a wider survey about topological infinite graph theory we refer to [25].

Throughout this section let G denote an arbitrary, hence also potentially infinite, graph.

F.2.1. Basic notions and tools

For any positive integer k let $[k] := \{1, \dots, k\}$.

Let X be a vertex set of G . We denote by $G[X]$ the induced subgraph of G with vertex set X and write $G - X$ for the graph $G[V(G) \setminus X]$. If H is a subgraph of G we shall write $G - H$ instead of $G - V(H)$. For an edge set $E \subseteq E(G)$ we denote by $G - E$ the subgraph of G with vertex set $V(G - E) := V(G)$ and edge set $E(G - E) := E(G) \setminus E$. To ease notation in case E is a singleton set, i.e. $E = \{e\}$ for some edge $e \in E(G)$, we shall write $G - e$ instead of $G - \{e\}$.

If T is a spanning tree of G and $e = xy \in E(T)$, let us denote by T_x and T_y the two components of $T - e$ containing x or y , respectively. Now $D_e := E(V(T_x), V(T_y)) \subseteq E(G)$ defines a cut of G and we denote it as the *fundamental cut of e w.r.t. T* in G . Also we call a cut of G a *fundamental cut w.r.t. T* if it is a fundamental cut of some edge $e \in E(T)$ w.r.t. T in G .

A path P is called an X -*path* if its endvertices lie in X , but the set of interior vertices of P is disjoint from X . Similarly, for a subgraph $H \subseteq G$ we call a path a H -*path* if it is a $V(H)$ -path. Given two vertex sets $A, B \subseteq V(G)$, we call a path Q in G an A - B path if Q is an a - b path for some $a \in A$ and some $b \in B$ whose set of interior vertices is disjoint from $A \cup B$. As before, given two subgraphs H_1, H_2 of G , we shall call a path a H_1 - H_2 path if it is a $V(H_1)$ - $V(H_2)$ path. If u and v are vertices of a (potentially infinite) tree T , then we write uTv to denote the unique u - v path in T .

We call a one-way infinite path R in G a *ray* of G , and a subgraph of R that is itself a ray a *tail* of R . We define an equivalence relation on the set of all rays of G by calling two rays in G *equivalent* if they cannot be separated in G via any finite vertex set of G . It is straightforward to check that this actually defines an equivalence relation. For two rays R_1, R_2 in G we shall write $R_1 \sim_G R_2$ to denote that R_1 and R_2 are equivalent in G . We shall drop the subscript in case it is clear in which surrounding graph we are arguing. Note that the statement $R_1 \sim_G R_2$ is equivalent to saying that there exist infinitely many pairwise disjoint R_1 - R_2 paths in G . We call the corresponding equivalence classes of rays under this relation the *ends* of G .

A subgraph H of G is called *end-faithful* if the following two properties hold:

- (i) every end of G contains a ray of H ;
- (ii) any two rays of H belong to a common end of H if and only if they belong to a common end of G .

A (possibly infinite) rooted tree T within a graph G is called *normal* if the endvertices of every T -path of G are comparable in the tree-order of T . Note that in the case of T being a spanning tree, every T -path is just an edge.

The following theorem is due to Jung. Since the reference [62] is a paper written in German, we include another textbook reference for the proof of the theorem.

Theorem F.2.1. [24,62] *Every countable connected graph has a normal spanning tree.*

The following lemma has also been proved by Jung [62] written in German. As before we include an additional textbook reference for the proof.

Lemma F.2.2. [24,62] *Every normal spanning tree of a graph G is an end-faithful subgraph of G .*

The following lemma is a basic tool in infinite combinatorics and well-known under the name Star-Comb Lemma. In order to formulate it we need to state another definition first.

We define a *comb* as the union of a ray R with infinitely many disjoint finite paths each having precisely its first vertex on R . The ray R is called the *spine* of the comb and the last vertices of the paths are called the *teeth* of the comb.

Lemma F.2.3. [24, Lemma 8.2.2] *Let U be an infinite set of vertices in a connected graph G . Then G contains either a comb with all teeth in U or a subdivision of an infinite star with all leaves in U .*

The following lemma is an easy consequence of the Star-Comb Lemma and should be known. We include the proof for the sake of completeness.

Lemma F.2.4. *Let G be a locally finite connected graph and let T be an end-faithful spanning tree of G . Then every fundamental cut w.r.t. T in G is finite.*

Proof. Let $e = xy \in E(T)$ and $D_e \subseteq E(G)$ be the fundamental cut of e w.r.t. T in G . Suppose for a contradiction that D_e is infinite. Now we apply Lemma F.2.3 to the vertex set $X = (\bigcup D_e) \cap V(T_x)$ in T_x and to the vertex set $Y = (\bigcup D_e) \cap V(T_y)$ in T_y . Since G is locally finite, the application of Lemma F.2.3 yields combs C_x and C_y in T_x and T_y , respectively, where the teeth of C_x lie in X and the teeth of C_y lie in Y . Let S_x and S_y denote the spines of C_x and C_y , respectively. Now within the graph $C_x \cup C_y \cup G[\bigcup D_e] \subseteq G$ there exist infinitely many disjoint S_x - S_y paths, witnessing that $S_x \sim_G S_y$. As S_x and S_y are contained in T but $S_x \not\sim_T S_y$, we have derived a contradiction to T being an end-faithful spanning tree of G . \square

F.2.2. Topological notions and tools

For this subsection we assume G to be a locally finite connected graph. We can endow the 1-skeleton of G together with its ends with a certain topology, yielding the space $|G|$ referred to as *Freudenthal compactification* of G . For a precise definition of $|G|$, see [24, Ch. 8.5]. Furthermore, we refer to [36] for Freudenthal's paper about the Freudenthal compactification, and to [29] regarding the connection to $|G|$. Note that the definition of $|G|$ ensures that each edge of G corresponds to an individual copy of the real unit interval $[0, 1]$ within $|G|$ and for adjacent edges of G , appropriate endpoints of the corresponding unit intervals are identified.

We denote the *closure* of a point set $X \subseteq |G|$ in $|G|$ by \overline{X} . A subspace S of $|G|$ is called a *standard subspace* if $S = \overline{F}$ for some edge set $F \subseteq E(G)$.

The next lemma yields an important combinatorial property of arcs. In order to state the lemma, let $\overset{\circ}{F}$ denote the set of inner points of edges $e \in F$ in $|G|$ for an edge set $F \subseteq E(G)$.

Lemma F.2.5. [24, Lemma 8.5.3] *Let G be a locally finite connected graph and $F \subseteq E(G)$ be a cut with sides V_1 and V_2 . If F is finite, then $\overline{V_1} \cap \overline{V_2} = \emptyset$, and there is no arc in $|G| \setminus \overset{\circ}{F}$ with one endpoint in V_1 and the other in V_2 .*

The following lemma ensures that being connected or being arc-connected are equivalent for closed subspaces of $|G|$.

Lemma F.2.6. [26, Thm. 2.6] *If G is a locally finite connected graph, then every closed topologically connected subset of $|G|$ is arc-connected.*

The next lemma characterises the property of a standard subspace of being topologically connected, and due to Lemma F.2.6 also being arc-connected, in terms of a purely combinatorial condition, which we shall make use of later.

Lemma F.2.7. [24, Lemma 8.5.5] *If G is a locally finite connected graph, then a standard subspace of $|G|$ is topologically connected (equivalently: arc-connected) if and only if it contains an edge from every finite cut of G of which it meets both sides.*

F.3. Counterexamples for potential strengthenings of Theorem F.1.5

F.3.1. No perfect matching in G

In this first subsection we discuss why the statement of Theorem F.1.5 becomes false if we omit the assumption of G having a perfect matching, even if we additionally assume higher connectivity of the graph G and focus on higher bi-powers. For finite graphs, Li [77] gave an example by constructing for every $k, \ell \in \mathbb{N}$ a k -connected balanced bipartite graph G that does not admit a perfect matching and where G_B^ℓ is not Hamiltonian.

Let us now recall Li's example and afterwards slightly extend it to also yield an example for locally finite infinite graphs, of course except for the property of being balanced. Let $s \geq \ell$ be an even number. Fix disjoint vertex sets V_0, \dots, V_{s+1} where $|V_0| = |V_{s+1}| > sk/2$ and $|V_i| = k$ for every $i \in [s]$. Define the graph $L_{k,s}$ by setting $V(L_{k,s}) = \bigcup_{i=0}^{s+1} V_i$ and by adding all possible edges between V_i and V_{i+1} for all $i \in \{0, 1, \dots, s+1\}$.

Clearly, $L_{k,s}$ is a finite balanced k -connected bipartite graph. By construction, the vertex set $V_0 \cup V_{s+1}$ forms an independent set of size greater than $|V(L_{k,s})|/2$. Hence, neither does $L_{k,s}$ admit a perfect matching nor $(L_{k,s})_B^\ell$ a Hamilton cycle; consequently $(L_{k,s})_B^\ell$ cannot be Hamilton-laceable either.

For locally finite infinite graphs let us analogously define the graph $H_{k,\ell}$ as follows. Let V_i denote disjoint vertex sets for every $i \in \mathbb{N}$ where $|V_0| > \lfloor \frac{\ell}{2} \rfloor k$ and $|V_i| = k$ for every $i > 0$. Similarly as before, set $V(H_{k,\ell}) = \bigcup_{i \in \mathbb{N}} V_i$ and add all edges between V_i and V_{i+1} for all $i \in \mathbb{N}$.

Clearly, $H_{k,\ell}$ is a locally finite k -connected bipartite graph. Since V_0 is an independent set and $|V_0| > |N(V_0)| = |V_1|$, we also have that $H_{k,\ell}$ does not contain a perfect matching for any $\ell \in \mathbb{N}$. Furthermore, in $(H_{k,\ell})_B^\ell$ the set V_0 has $\lfloor \frac{\ell}{2} \rfloor k < |V_0|$ many neighbours. Hence, $(H_{k,\ell})_B^\ell$ cannot be Hamiltonian, and hence also not be Hamilton-laceable.

F.3.2. Perfect matching in G_B^3 instead of G

In [77] Li also briefly addresses the question whether the assumption within the statement of Theorem F.1.3 of G admitting a perfect matching might be weakened

to G_B^3 admitting a perfect matching. As Hamiltonicity for bipartite graphs implies the existence of a perfect matching, this is a necessary condition. However, Li gave a counterexample for this condition to also be sufficient for finite graphs in the context of Theorem F.1.3. We shall now briefly recall Li's example and then slightly extend it to also yield a counterexample with respect to locally finite infinite graphs and Theorem F.1.5.

For $k \geq 3$ start with a tree that has precisely two vertices of degree $k + 1$ while all other vertices have degree 1. Next subdivide each edge that is incident with a leaf twice and call the resulting graph L_k . As Li noted, it is easy to check that $(L_k)_B^3$ does admit a perfect matching, but no Hamilton cycle.

For a locally finite infinite graph we now state the following construction. Start with a star $K_{1,k}$ whose centre is c and that has precisely k leaves ℓ_1, \dots, ℓ_k for some $k \geq 3$. Next subdivide each edge precisely twice. Now take the disjoint union of the resulting graph with a ray $R = r_1 r_2 \dots$. Finally, add the edge cr_1 and k further vertices that are all only adjacent to r_1 . Let us call the resulting graph H_k .

As in Li's example, it is easy to check that $(H_k)_B^3$ admits a perfect matching. The key observation why $(H_k)_B^3$ is not Hamiltonian is also the same as in Li's example, which we briefly recall with respect to our example. Note that the vertices ℓ_1, \dots, ℓ_k have degree 2 in $(H_k)_B^3$ and they all share c as a common neighbour. Hence, any Hamilton circle \bar{C} of $(H_k)_B^3$, where C denotes some subgraph of $(H_k)_B^3$, would impose c to have degree $k \geq 3$ in C , which is not possible.

F.3.3. Hamilton-connectedness

Another question that might arise when considering the statements of Theorem F.1.3 and Theorem F.1.5 is whether even Hamilton-connectedness can be deduced. To recall: the notion of *Hamilton-connectedness* is analogously defined to the one of Hamilton-laceability for finite as well as for locally finite graphs but by considering all pairs of distinct vertices instead of just those that have odd distance from each other. For finite bipartite graphs and the statement of Theorem F.1.3 we can clearly not deduce Hamilton-connectedness as the assumption of having a perfect matching ensures our graph to be balanced, while a Hamilton path with endvertices in the same bipartition class would violate this.

For locally finite infinite bipartite graphs we can also negatively answer the

question via a counterexample very easily. Consider a ray $R = r_1 r_2 \dots$. The graph R_B^3 does not admit a Hamilton r_2 - r_{2k} arc for any $k > 1$. To see this, note first that r_1 has degree 2 in R_B^3 . Hence, any potential spanning arc starting in r_2 would have the path $r_2 r_1 r_4$ as initial segment. Now r_3 has degree 2 in $R_B^3 - r_2$, which forces the initial segment $r_2 r_1 r_4 r_3 r_6$. By iterating this argument, we see that any potential spanning arc for R_B^3 starting in r_2 would contain an r_2 - r_{2k} path as an initial segment, preventing the existence of a Hamilton r_2 - r_{2k} arc.

F.4. Proof of the main result

We start this section by very briefly sketching and discussing the rough methodological differences regarding how Li proved Theorem F.1.4 in [77] and how we prove Theorem F.1.5. Before we can relate the different approaches, we have to give some additional definitions first.

Let G be a locally finite infinite graph and ω be an end of G . We define the *degree* of ω in G to be the supremum of the number of vertex-disjoint rays in G which are contained in ω . Furthermore, we call a continuous image of $S^1 \subseteq \mathbb{R}^2$ within $|G|$ that contains all vertices of G a *Hamilton curve* of G . Note that in contrast to a Hamilton circle of G , a Hamilton curve of G may traverse an end of G several times.

Li's approach to verify Hamiltonicity of G_B^3 is to first start with an end-faithful spanning tree T of G which contains a perfect matching of G . Then he carefully extends T by edge sets of suitable cycles such that an end-faithful spanning subgraph G' of G_B^3 is obtained that admits a Hamilton curve and each of whose ends has degree at most 3. It is easy to see that every Hamilton curve in G' is actually already a Hamilton circle as the end degrees do not allow to traverse an end multiple times. Due to the end-faithfulness, that Hamilton circle of G' is also one of G_B^3 . Furthermore, Li chooses the set of suitable cycles in such a way that they also prove the existence of a Hamilton curve of G' . For this he makes use of the following characterisation for the existence of Hamilton curves.

Theorem F.4.1. [68] *A locally finite connected graph G has a Hamilton curve if and only if every finite vertex set of G is contained in some finite cycle of G .*

The way how Li precisely constructs the mentioned set of cycles goes back to

his proof of Lemma 1 in [77], which is basically the finite version of Lemma F.4.5.

Although using Theorem F.4.1 can be a powerful and convenient tool, it does not seem immediately helpful for the purpose of verifying Hamilton-laceability (or Hamilton-connectedness). Our way to prove Theorem F.1.5 is to mimic Li's proof of Theorem F.1.3 in finite graphs. Hence, we especially construct certain Hamilton arcs of the considered graph G directly within the third bi-power of an end-faithful spanning tree of G that contains a perfect matching of G . This part of our proof happens in Lemma F.4.5. Later in the proof of Theorem F.1.5 when we apply Lemma F.4.5, we combine Hamilton arcs of suitable subgraphs of G_B^3 along an inductive argument to yield the desired Hamilton arcs of G_B^3 . The general idea of this part is the same as in Li's proof of Theorem F.1.3. However, we have to build our induction on a different parameter, namely on distances between vertices, to ensure that the parameter is always finite. Similarly as Lemma 1 in [77] was the key lemma in Li's proof of Theorem F.1.3, now Lemma F.4.5 is our key lemma to prove Theorem F.1.5.

Now let us start preparing to prove Theorem F.1.5. The following lemma ensures that we can always extend a perfect matching of a countable graph G to an end-faithful spanning tree of G . Although this lemma can be deduced from a more general lemma in Li's article [77, Lemma 6], we decided to include a proof here for the sake of keeping this article self-contained and because our proof seems simpler due to the less technical setting.

Lemma F.4.2. *Let G be a countable connected graph and M be a perfect matching of G . Then there exists an end-faithful spanning tree of G that contains M .*

Proof. Let G and M be as in the statement of the theorem. Now we apply Theorem F.2.1 with the graph G/M , which is still a countable connected graph, guaranteeing us the existence of a normal spanning tree T' of G/M . Next we uncontract every edge of M in T' and form, within G , a spanning tree T of G . Note for this that M is a perfect matching of G . Hence, every vertex of T' corresponds to an edge $m \in M$, and so we shall also work with M as the vertex set of T' for the rest of the proof. In order to define T we pick for every edge $e_{T'} = m_1 m_2 \in E(T')$ an arbitrary edge $e_T \in E(G)$ that witnesses the existence of $m_1 m_2 \in E(T')$, i.e. $e_T = u_1 u_2$ where $u_1, u_2 \in V(G)$ and u_i is an endvertex of m_i for each $i \in \{1, 2\}$.

Now we define T as follows:

$$\begin{aligned} V(T) &:= V(G), \\ E(T) &:= M \cup \{e_T \mid e_{T'} \in E(T')\}. \end{aligned}$$

Obviously, T is still connected and does not contain a finite cycle. Hence, T is a spanning tree of G that contains M by definition.

Next we verify that T is an end-faithful subgraph of G . Let ω be any end of G and $R \in \omega$. Now R/M is a ray in G/M , and since T' is an end-faithful spanning tree of G/M , there exists a ray R' in T' such that $R' \sim_{G/M} R/M$. Let \mathcal{P}' be a set of infinitely many pairwise disjoint R' - R/M paths in G/M witnessing $R' \sim_{G/M} R/M$. Now let R_T be any ray in T obtained from $T \setminus \{e_T \mid e_{T'} \in E(R')\}$ by adding edges from M . By uncontracting edges from M , the path system \mathcal{P}' now gives rise to a set of infinitely many pairwise disjoint R_T - R paths in G . Hence, $R_T \subseteq T$ is a desired ray satisfying $R_T \sim_G R$.

Now let R_1, R_2 be two rays of T . If $R_1 \sim_T R_2$, then R_1 must be a tail of R_2 or vice versa since T is a tree. Hence, $R_1 \sim_G R_2$. Conversely, suppose for a contradiction that $R_1 \sim_G R_2$ but $R_1 \not\sim_T R_2$. First, note that $R_1/M \sim_{G/M} R_2/M$ holds. Using that T' is an end-faithful subgraph of G/M we also know $R_1/M \sim_{T'} R_2/M$. However, from $R_1 \not\sim_T R_2$ we know that a finite set $S \subseteq V(G)$ exists separating R_1 and R_2 in T . Now the set $S' := \{m \in M \mid m \cap S \neq \emptyset\}$ defines a vertex set in T' that separates R_1/M and R_2/M in T' ; a contradiction. \square

The following question arose while preparing this article. It basically asks whether we can also get a normal spanning tree to satisfy the conclusion of Lemma F.4.2 instead of just an end-faithful one. Although neither a positive nor a negative answer to this question would substantially affect or shorten the proof of the main result of this paper, the question seems to be of its own interest. Hence, it is included here.

Question F.4.3. Let G be a countable connected graph and let M be a perfect matching of G . Does G admit a normal spanning tree that contains M ?

Since the assumption of countability in Lemma F.4.2 was only used to ensure the existence of a normal spanning tree, the following question is a related, but more general one than Question F.4.3.

Question F.4.4. Let G be a connected graph and M be a perfect matching of G such that G/M admits a normal spanning tree. Does G admit a normal spanning tree that contains M ?

Note that both questions above have positive answers when we restrict them to finite graphs. We can easily include the desired perfect matching during the constructing of a depth-first search spanning tree, which in particular is a normal one.

During a discussion with Carsten Thomassen it turned out that Question F.4.3 has an easy counterexample, which then also negatively answers Question F.4.4. In the following lines we shall state the counterexample. We shall follow the convention that the set of natural numbers \mathbb{N} contains the number 0.

For $i = 1, 2$ let $V^i = \{v_0^i, v_1^i, \dots\}$ be two disjoint countably infinite vertex sets. Now we define our desired graph G as follows. Let $V(G) := V^1 \cup V^2$. Furthermore, for each $i = 1, 2$ we define edge sets $E^i := \{v_j^i v_{j+1}^i \mid j \in \mathbb{N}\}$. Finally, we define the edge set $E^* = \{v_{2k}^1 v_{2k}^2 \mid k \in \mathbb{N} \setminus \{0\}\}$, and set $E(G) := E^1 \cup E^2 \cup E^*$. Next we define a perfect matching M of G which cannot be included in any normal spanning tree of G , independent of the choice of the root vertex for the tree. Set $M := \{v_{2k}^i v_{2k+1}^i \mid i \in \{1, 2\}, k \in \mathbb{N}\}$. See Figure F.1 for a picture of the graph G together with its perfect matching M .




Figure F.1.: The graph G with the perfect matching M indicated by bold edges.

Next we prove why no normal spanning tree of G can contain M , independent of the choice for the root of the tree. Note first that any normal spanning tree T of G must contain infinitely many edges from E^* . This is because otherwise T would contain two disjoint rays in the end of G , which contradicts the fact that normal spanning trees are end-faithful and, therefore, contain only a unique ray in each end of G . See [62] for a proof of this fact. Now suppose for a contradiction that T is a normal spanning tree of G containing M . Let $r \in V(T)$ denote the root of T ,

and let $j \in \{1, 2\}$ and $\ell \in \mathbb{N}$ be such that $r = v_\ell^j$. Next let $p \in \mathbb{N}$ be the smallest index for which $p > \ell$ and $v_p^1 v_p^2 \in E(T)$ hold. Similarly, let $q \in \mathbb{N}$ be the second smallest index for which $q > \ell$ and $v_q^1 v_q^2 \in E(T)$ hold. Clearly, the paths $v_p^1 \dots v_q^1$ and $v_p^2 \dots v_q^2$ cannot both be contained in T since then T would contain a finite cycle. Without loss of generality say that $v_p^1 \dots v_q^1$ is not contained in T . Then let $s \in \mathbb{N}$ be the smallest index such that $p < s < q$ and $v_s^1 v_{s+1}^1 \notin E(T)$. Since M is contained in T , we know that s is an odd number. So v_s^1 has degree 2 in G and is a leaf in T . Since $\ell < s$, the vertices v_s^1 and v_{s+1}^1 lie on different branches in T . Hence, the edge $v_s^1 v_{s+1}^1$ contradicts the normality of T . Therefore, we can conclude that it is impossible for any normal spanning tree of G to contain the perfect matching M .

Now we continue with the key lemma for the proof of our main result.

Lemma F.4.5. *Let G be a locally finite connected bipartite graph and M be a perfect matching of G . Furthermore, let T be an end-faithful spanning tree of G containing M . Then for every edge $xy \in M$ there exists a Hamilton x - y arc of G_B^3 within $\overline{T_B^3} \subseteq |G_B^3|$*

Proof. Let G , T and M be as in the statement of the lemma. First note that, as G is bipartite, clearly $T_B^3 \subseteq G_B^3$. Now let $m_0^1 = x_0^1 y_0^1$ be an arbitrary edge from M . We recursively make the following definitions: Set $T_0 = T[\{x_0^1, y_0^1\}]$. Now for every $i \in \mathbb{N}$, set T_{i+1} to be the subtree of T induced by all vertices of T which are contained in some edge $m' \in M$ such that m' has one endvertex in distance at most 1 to T_i within T . Next we shall recursively define a Hamilton x_0^1 - y_0^1 path A_i in each $(T_i)_B^3$ for every $i \in \mathbb{N}$ where A_i contains each edge $m \in M$ which is contained in $T_i - V(T_{i-1})$, where $V(T_{-1}) := \emptyset$.

First set $E(A_0) := m_0^1$. Now suppose we have already defined A_i and want to define A_{i+1} . Let us enumerate the edges in $T_i - V(T_{i-1})$ that are contained in $E(A_i) \cap M$ by $m_i^1, \dots, m_i^{p_i}$ for some $p_i \in \mathbb{N}$. Furthermore, let us fix names for the endvertices of each such edge by writing $m_i^j = x_i^j y_i^j$ for every $j \in [p_i]$. By definition, either x_i^j or y_i^j is adjacent to some endvertex y_{i-1}^q , for some $q \in [p_{i-1}]$, of an edge $m_{i-1}^q \in M$ contained in T_{i-1} , w.l.o.g. say x_i^j . Let $T_{x_i^j}$ denote the component of $T_{i+1} - y_{i-1}^q x_i^j$ containing x_i^j if $i > 0$, and set $T_{x_0^j} := T_1$ for $i = 0$. Now we shall extend A_i further into each $T_{x_i^j}$ (unless $E(T_{x_i^j}) = m_i^j$) by replacing the edges m_i^j for all $j \in [p_i]$ in A_i by a Hamilton x_i^j - y_i^j path P_i^j of $(T_{x_i^j})_B^3$ to form A_{i+1} . We shall

distinguish three cases of how to do this for each $j \in [p_i]$.

Case 1. *Both components of $T_{x_i^j} - m_i^j$ are trivial.*

In this case we keep the edge m_i^j to also be an edge of A_{i+1} .

Case 2. *Precisely one component of $T_{x_i^j} - m_i^j$ is trivial, w.l.o.g. say the one containing x_i^j .*

In this situation let $N(y_i^j) = \{x_i^j, v_1, v_2, \dots, v_k\}$ for some $k \geq 1$. Since M is a perfect matching of G , each v_r must be contained in some edge from M , say $v_r w_r \in M$ for every $r \in [k]$. Now set $P_i^j := x_i^j w_1 v_1 w_2 v_2 \dots w_k v_k y_i^j$.

Case 3. *No component of $T_{x_i^j} - m_i^j$ is trivial.*

Now let $N(y_i^j) = \{x_i^j, v_1, v_2, \dots, v_k\}$ and $N(x_i^j) = \{y_{i-1}^q, y_i^j, a_1, a_2, \dots, a_\ell\}$ for some $k, \ell \geq 1$. Using as before that M is a perfect matching of G , each v_r and each a_s must be contained in some edge of M , say $v_r w_r \in M$ for every $r \in [k]$ and $a_s b_s \in M$ for every $s \in [\ell]$. Finally, set $P_i^j := x_i^j w_1 v_1 w_2 v_2 \dots w_k v_k a_1 b_1 a_2 b_2 \dots a_\ell b_\ell y_i^j$.

This completes the definition of A_{i+1} . Note that this definition ensures that an edge $e \in A_i$ is also contained in A_{i+1} except $e = m_i^j \in M$ for some $j \in [p_i]$ as above and at least one component of $T_{x_i^j} - m_i^j$ is non-trivial.

Next we define our desired Hamilton $x_0^1 - y_0^1$ arc in $\overline{T_B^3} \subseteq |G_B^3|$. For this we set

$$A := \left\{ e \in \bigcup_{n \in \mathbb{N}} E(A_n) \mid e \text{ is contained in all but finitely many } A_n \right\}.$$

Now we claim that \overline{A} is a desired $x_0^1 - y_0^1$ Hamilton arc of G_B^3 .

Note first that by definition $\overline{A} \subseteq \overline{T_B^3} \subseteq |G_B^3|$ and \overline{A} contains all vertices of T , and therefore all vertices of G because T is a spanning tree of G . It remains to verify that \overline{A} is an $x_0^1 - y_0^1$ arc. We shall do this by showing that \overline{A} contains an $x_0^1 - y_0^1$ arc but $\overline{A - e}$ does not for any edge $e \in A$, which implies that \overline{A} is an $x_0^1 - y_0^1$ arc.

To prove that \overline{A} contains an $x_0^1 - y_0^1$ arc in $|G_B^3|$ it is enough to show that A intersects every finite cut of G_B^3 due to Lemma F.2.7. Now let $F \subseteq E(G_B^3)$ be an arbitrary finite cut of G_B^3 . Hence, $F \subseteq E(G_B^3[V(T_n)])$ for some $n \in \mathbb{N}$. Since A_{n+1} is a Hamilton $x_0^1 - y_0^1$ path in $(T_{n+1})_B^3$ and $E(A_{n+1}) \cap E(G_B^3[V(T_n)]) = A \cap E(G_B^3[V(T_n)])$, we know that A intersects F .

For the remaining argument that $\overline{A - e}$ does not contain an $x_0^1 - y_0^1$ arc for any edge $e \in A$, we shall first find a finite cut F of G_B^3 that A intersects precisely in e . Let $n \in \mathbb{N}$ such that $e \in E(G_B^3[V(T_n)])$. Since A_n is a Hamilton $x_0^1 - y_0^1$ path in $G_B^3[V(T_n)]$, there exists a cut $F_n = E(L_n, R_n)$ of $G_B^3[V(T_n)]$ such that $x_0^1 \in L_n$ and $y_0^1 \in R_n$ and $E(A_n) \cap F_n = \{e\}$. Let \mathcal{C}_L and \mathcal{C}_R be the sets of all components of $T - E(T_n)$ that intersect L_n and R_n , respectively. Next we extend the bipartition (L_n, R_n) of $V(T_n)$ to a bipartition (L, R) of $V(T)$. We set $L := \bigcup\{V(C) \mid C \in \mathcal{C}_L\}$ and $R := \bigcup\{V(C) \mid C \in \mathcal{C}_R\}$. Especially, this yields $x_0^1 \in L_n \subseteq L$ and $y_0^1 \in R_n \subseteq R$. Furthermore, T intersects the cut $F_G := E(L, R)$ of G in the same edges as T_n intersects F_n . Hence, T intersects F_G in only finitely many edges. Next note that any edge $f = uv \in F_G \setminus E(T)$ lies in the fundamental cut D_g of G w.r.t. T for every edge g that lies on the $u-v$ path in T . Especially, f lies in $D_{g'}$ for some of the finitely many edges $g' \in F_G \cap E(T)$. As every fundamental cut of G w.r.t. T is finite by Lemma F.2.4, this implies that F_G is a finite cut of G . Since G is locally finite and by definition of G_B^3 , we furthermore get that the bipartition (L, R) of $V(G) = V(G_B^3)$ also yields a finite cut F of G_B^3 . By the definition of the A_i 's we know that every A_m for $m \geq n$ also satisfies $E(A_m) \cap F = \{e\}$. Hence, $A \cap F = \{e\}$.

To complete the argument, note that every $x_0^1 - y_0^1$ arc in $|G_B^3|$ must intersect F by Lemma F.2.5. This, however, implies that $\overline{A - e}$ cannot not contain an $x_0^1 - y_0^1$ arc. \square

Now we are able to prove our main result.

Proof of Theorem F.1.5. Let u and v be two vertices from different bipartition classes of G and let M be a perfect matching of G . By Lemma F.4.2 there exists an end-faithful spanning tree T of G that contains M . Since G is bipartite, the distance between u and v in T is also odd. We now prove the statement of the theorem by induction on $d := |E(uTv) \setminus M|$. Since $|E(uTv)|$ is odd and $|E(uTv) \cap M| \leq \left\lfloor \frac{|E(uTv)|}{2} \right\rfloor$ as M is a perfect matching of G , we know that $d = 0$ holds precisely when $E(uTv) = \{uv\}$ and $uv \in M$. Now the statement follows from Lemma F.4.5.

Next let us verify the statement for $d > 0$ while assuming we have verified it for all smaller values for d . There must exist an edge $xy \in E(uTv) \setminus M$, say without loss of generality $x \in E(uTy)$. As T is end-faithful, the fundamental cut D_{xy}

w.r.t. T in G is finite by Lemma F.2.4. Note that T is also an end-faithful spanning tree of G_B^3 . To see this observe first that given any ray R in G_B^3 we obtain by applying Lemma F.2.3 to $V(R)$ within $T \subseteq G_B^3$, and due to T being locally finite, a comb whose spine is equivalent to R in T , and hence also in G_B^3 . Second, let two non-equivalent rays R_1, R_2 in T be given. As T is an end-faithful spanning tree of G , there exists a finite vertex set $S \subseteq V(G)$ such that $R_1 - S$ and $R_2 - S$ lie in different components of $G - S$. By definition of G_B^3 and due to the locally finiteness of G , we get that $S \cup N(S)$ is a finite vertex set separating R_1 and R_2 in G_B^3 . Now since G is locally finite and T is also an end-faithful spanning tree of G_B^3 , we know that the fundamental cut D_{xy}^B w.r.t. T in G_B^3 is finite as well. For ease of notation set $H := G[T_x]$ and $K := G[T_y]$. Due to the finite fundamental cut D_{xy} , we know that the spaces $|H|$ and $|K|$ are homeomorphic to the subspaces of $|G|$ induced by the closures \overline{H} and \overline{K} , respectively. Furthermore, $\overline{H} \cap \overline{K} = \emptyset$ by Lemma F.2.5. The same observations hold for H_B^3 and K_B^3 since T is also an end-faithful spanning tree of G_B^3 implying, as note before, that the fundamental cut D_{xy}^B w.r.t. T in G_B^3 is finite. Next we shall make a case distinction of how to apply our induction hypothesis. Note that since $|E(uTv)|$ is odd, either $|E(uTx)|$ and $|E(yTv)|$ are both odd or they are both even.

Case 1. $|E(uTx)|$ and $|E(yTv)|$ are odd.

In this case we apply our induction hypothesis with the graphs H and K , their perfect matchings $M_x := M \cap E(T_x)$ and $M_y := M \cap E(T_y)$, the end-faithful spanning trees T_x and T_y , which contain M_x and M_y respectively, and the pairs of vertices (u, x) and (y, v) . Hence we obtain a Hamilton $u-x$ arc A_x of H_B^3 within within $(T_x)_B^3$ and a Hamilton $y-v$ arc A_y of K_B^3 within within $(T_y)_B^3$. Since $\overline{H}_B^3 \cap \overline{K}_B^3 = \emptyset$ holds within $|G_B^3|$, we obtain a Hamilton $u-v$ arc of G_B^3 by joining A_x and A_y via the edge xy .

Case 2. $|E(uTx)|$ and $|E(yTv)|$ are even.

First note for this case that there exist edges $xx', yy' \in M$ since M is a perfect matching. As $M \subseteq E(T)$ by assumption, we know that $xx' \in E(T_x)$ and $yy' \in E(T_y)$. Also note that $|E(uTv \setminus M)| = |E(uTx' \setminus M)| + |E(y'Tv \setminus M)| + 1$, whether x' or y' are contained in uTv or not, and that $|E(uTx')|$ and $|E(y'Tv)|$ are both odd. Furthermore, $\text{dist}_T(x', y') = 3$, and, therefore, $\text{dist}_G(x', y') \in \{1, 3\}$. Hence, $x'y' \in E(T_B^3) \subseteq E(G_B^3)$. Due to these observations we can apply our induction

hypothesis as in Case 1 but with x' and y' instead of x and y , yielding again the desired Hamilton u - v arc of G_B^3 . \square

Chapter III.

Dijoins of digraphs

G. Disjoint dijoins for classes of dicuts in finite and infinite digraphs

G.1. Introduction

In this paper we consider directed graphs, which we briefly denote as *digraphs*. A *dicut* in a digraph is a cut for which all of its edges are directed to a common side of the cut. A famous theorem of Lucchesi and Younger [78] states that in every finite digraph the least size of a set of edges meeting every non-empty dicut equals the maximum number of disjoint dicuts in that digraph. Such sets of edges are called *dijoins*.

Woodall conjectured the following in some sense dual statement, where the roles of the minimum and maximum are reversed.

Conjecture G.1.1 (Woodall 1976 [111]). *The size of a smallest non-empty dicut in a finite digraph D is equal to the size of the largest set of disjoint dijoins of D .*

This conjecture is a long-standing open question in this area and is included in a list of important conjectures compiled by Cornuéjols [22].

Not much is known in general about this conjecture. It is easy to find two disjoint dijoins in a bridgeless weakly connected digraph D : just consider an orientation of the underlying undirected multigraph that yields a strongly connected digraph, which exists by a theorem of Robbins [87]; then the two desired dijoins of D are the set of edges which agrees with this orientation and the set of edges which disagrees with this orientation. This observation was noted by Seymour and DeVos [114].

But beyond that, there is no known bound on the size of a smallest dicut that can guarantee the existence of even three disjoint dijoins [113, 114].

There are several partial results restricting the attention to digraphs with certain properties. Lee and Wakabayashi [73] verified Conjecture G.1.1 for digraphs whose underlying multigraph is series-parallel, which was later improved by Lee and

Williams [74] proving the conjecture for digraphs whose underlying multigraph is planar and does not contain a triangular prism $K_3 \square K_2$ as a minor*. Schrijver [92] and independently Feofiloff and Younger [33] verified the conjecture for *source-sink connected* digraphs, i.e. digraphs where from every source there exists a directed path to every sink. For additional partial results see [81].

Thomassen [105] showed with some tournament on 15 vertices that a dual version regarding directed cycles and disjoint *feedback arc sets*, i.e. sets of edges meeting every directed cycle, fails for non-planar digraphs. For planar digraphs, these questions are obviously equivalent and still open.

A capacitated version of Woodall's Conjecture (cf. Section G.3), conjectured by Edmonds and Giles [32] was proven to be false by Schrijver [93]. Although false in general, the conjecture of Edmonds and Giles has been verified in some special cases. In particular, the works by Lee and Wakabayashi [73], Lee and Williams [74], and Feofiloff and Younger [33] are actually about the conjecture of Edmonds and Giles and obtain corresponding results about Woodall's Conjecture as corollaries. For more research regarding this line of work, including a study of the structure of possible counterexamples, see [21, 23, 99, 110].

Instead of focusing on specific classes of digraphs, one other possible avenue to explore Conjecture G.1.1 is to restrict the attention not to all dicuts, but to some specific classes of dicuts of a digraph. In [40, 43], Gollin and Heuer considered a similar approach regarding classes of dicuts in their attempt of generalising the theorem of Lucchesi and Younger to infinite digraphs.

For a digraph D and a non-empty class of non-empty dicuts \mathfrak{B} of D , a set of edges meeting every dicut in \mathfrak{B} is a \mathfrak{B} -*dijoin*. Using this terminology, a natural modification of Conjecture G.1.1 is the following question.

Question G.1.2. For which digraphs D and classes of dicuts \mathfrak{B} of D is the size of a smallest dicut in \mathfrak{B} equal to the size of a largest set of disjoint \mathfrak{B} -dijoins of D ?

A *dibond* of D is a minimal non-empty dicut. We say a class \mathfrak{B} of dicuts of a digraph D is *dibond-closed* if every dibond which is contained in some dicut in \mathfrak{B} is contained in \mathfrak{B} as well. Note that whenever \mathfrak{B} is a dibond-closed class of dicuts,

*Note that they proved the dual statement about feedback arc sets in planar digraphs without a minor isomorphic to $K_5 - e$, the graph obtained from K_5 by deleting one of its edges, which is the planar dual of the triangular prism $K_3 \square K_2$.

then Question G.1.2 for this class is equivalent to the question for the class of all dibonds contained in \mathfrak{B} . Hence in the setting of Question G.1.2, Conjecture G.1.1 translates into the question for the class $\mathfrak{B}_{\text{dibond}}$ of all dibonds of D .

For classes of dicuts which are not dibond-closed in digraphs that are not necessarily weakly connected, this question relates to the capacitated version of Question G.1.2 for classes of dibonds. We will investigate this connection in Section G.3.

In this paper, we will give positive answers for Question G.1.2 for several classes of dicuts. Two dicuts are *nested* if some side of one of them is a subset of some side of the other. A set of dicuts is *nested* if the dicuts in that set are pairwise nested. We prove a result for nested classes of finite dicuts, using the machinery developed by Berge [9] for transversal packings in balanced hypergraphs.

Theorem G.1.3. *Let D be a digraph and \mathfrak{B} be a nested class of finite dicuts of D . Then the size of a smallest dicut in \mathfrak{B} is equal to the size of a largest set of disjoint \mathfrak{B} -dijoins of D .*

Another interesting class is the class $\mathfrak{B}_{\text{min}}$ of dicuts of minimum size. An inductive construction allows us to reduce the following theorem to Theorem G.1.3.

Theorem G.1.4. *The size of a smallest dicut in a digraph D that contains finite dicuts is equal to the size of largest set of disjoint $\mathfrak{B}_{\text{min}}$ -dijoins of D , where $\mathfrak{B}_{\text{min}}$ denotes the class of dicuts of D of minimum size.*

The parts of Theorems G.1.3 and G.1.4 regarding infinite digraphs are proved using the compactness principle in combinatorics. We will also use that technique to prove a finitary version of the results of Lee and Williams [74] and Feofiloff and Younger [33] for infinite digraphs only considering dicuts of finite size (capacity).

Finally, we verify a cardinality version for classes of infinite dibonds of Question G.1.2. This proof uses a transfinite recursion to construct the dijoins in the case where the dibonds have the same cardinality as the order of the digraph. To complete the proof, we use the concept of bond-faithful decompositions due to Laviolette [72].

Theorem G.1.5. *Let D be a digraph and \mathfrak{B} be a class of infinite dibonds of D . The size of a smallest dibond in \mathfrak{B} is equal to the size of largest set of disjoint \mathfrak{B} -dijoins of D .*

This paper is structured as follows. After introducing some terminology in Section G.2, we discuss the capacitated version of Question G.1.2 and its connection to classes of dicuts which are not dibond-closed in Section G.3. In Section G.4, we look at transversal packings of balanced hypergraphs and prove Theorem G.1.3. Section G.5 is dedicated to prove Theorem G.1.4 for finite digraphs. We turn our attention to infinite digraphs in Section G.6. After showing that a more structural generalisation of Conjecture G.1.1 to infinite digraphs fails (see Example G.6.1), in Subsection G.6.1 we use the compactness principle in combinatorics to finish the proof of Theorem G.1.4 for infinite digraphs and prove a finitary version of the results of Lee and Williams [74] and Feofiloff and Younger [33]. Finally, in Subsection G.6.2 we prove Theorem G.1.5 and give an example that a capacitated version of that theorem fails.

G.2. Preliminaries

For general facts and notation for graphs we refer the reader to [24], for digraphs in particular to [8], and for hypergraphs to [9]. For some set theoretic background, including ordinals, cardinals and transfinite induction, we refer the reader to [60].

G.2.1. Digraphs

Let D be a digraph with vertex set $V(D)$ and edge set $E(D)$. We allow D to have parallel edges, but may assume for most purposes in this paper that D does not contain any loops. We view the edges of D as ordered pairs (u, v) of vertices $u, v \in V(D)$ and shall write uv instead of (u, v) , although this might not uniquely determine an edge if D contains parallel edges. We say D is *simple* if it does not contain parallel edges (or loops). For an edge $uv \in E(D)$ we furthermore call the vertex u as the *tail* of uv and v as the *head* of uv .

A digraph D is *weakly connected* if its underlying undirected (multi-)graph is connected. The components of the underlying undirected (multi-)graph are the *weak components* of D .

For two sets $X, Y \subseteq V(D)$ of vertices we define $E_D(X, Y) \subseteq E(D)$ as the set of those edges that have their head in $X \setminus Y$ and their tail in $Y \setminus X$, or their head in $Y \setminus X$ and their tail in $X \setminus Y$.

We consider a *bipartition* of $V(D)$ to be an ordered pair (X, Y) for which $X \cap Y = \emptyset$ and $X \cup Y = V(D)$. We call (X, Y) *trivial* if either X or Y is empty.

A set B of edges of D is a *cut* of D if there is a non-trivial bipartition (X, Y) of $V(D)$ such that $B = E_D(X, Y)$. We call (X, Y) a *representation* of B (or say (X, Y) *represents* B), and we refer to X and Y as the *sides* of the representation (or the *sides* of the cut, if the representation is inferred from context). Note that a cut of a weakly connected digraph has up to the ordering of the pair a unique representation. For a set \mathcal{B} of cuts, a set \mathcal{R} of bipartitions of $V(D)$ *represents* \mathcal{B} if for each $B \in \mathcal{B}$ there is a bipartition $(X, Y) \in \mathcal{R}$ that represents B .

We define a partial order on the set of bipartitions of $V(D)$ by

$$(X, Y) \leq (X', Y') \quad \text{if and only if} \quad X \subseteq X' \text{ and } Y \supseteq Y'.$$

Two bipartitions (X, Y) and (X', Y') of $V(D)$ are *nested* if one of (X, Y) , (Y, X) is \leq -comparable with one of (X', Y') , (Y', X') . Note that $(X, Y) \leq (X', Y')$, if and only if $(Y', X') \leq (Y, X)$. A set \mathcal{R} of bipartitions of $V(D)$ is nested if its elements are pairwise nested.

Two cuts B_1 and B_2 are *nested* if they can be represented by nested bipartitions of $V(D)$. Moreover, we call a set (or sequence) \mathcal{B} of cuts of D nested if there is a nested set of bipartitions that represents \mathcal{B} . Note that in a digraph which is not weakly connected, a set of pairwise nested cuts is not necessarily nested itself. If two cuts of D are not nested, we call them *crossing*.

A *bond* is a minimal non-empty cut (with respect to the subset relation). Note that if D is weakly connected, then a cut B of D represented by (X, Y) is a bond, if and only if the induced subdigraphs $D[X]$ and $D[Y]$ are weakly connected digraphs.

We call a cut B *directed*, or briefly a *dicut*, if all its edges have their head in one common side of the cut. A bond that is also a dicut is called a *dibond*.

If a bipartition (X, Y) of $V(D)$ represents a non-empty dicut, we call the side of the representation that contains the heads of the edges of the dicut the *in-shore* $\text{in}_D(X, Y)$ of the representation, and we call the side of the representation that contains the tails of the edges of the dicut the *out-shore* $\text{out}_D(X, Y)$ of the representation. If the representation of a dicut B is clear from the context (for example if D is weakly connected) we will speak of the *in-shore* $\text{in}_D(B)$ of B , or *out-shore* $\text{out}_D(B)$ of B .

A cut or dicut *separates* two vertices (or two sets of vertices) if there is a representation for which they are contained in different sides. Similarly, B separates two sets of vertices, if they are contained in different sides of some representation of B .

Let \mathfrak{B} be a class of dibonds of D . A \mathfrak{B} -*dijoin* is a set of edges meeting every dibond in \mathfrak{B} . We say D is \mathfrak{B} -*Menger* if the size of a smallest dibond in \mathfrak{B} is equal to the size of a largest set of disjoint \mathfrak{B} -dijoins of D .

G.2.2. Hypergraphs

A *hypergraph* \mathcal{H} is a tuple (X, H) consisting of a set X and a set H of subsets of X . The set X is called the *ground set* of \mathcal{H} and the elements of H are called *hyperedges*.

A set $T \subseteq X$ is a *transversal* of \mathcal{H} if $T \cap h$ is non-empty for every hyperedge $h \in H$.

For a set $Y \subseteq X$ we define the *subhypergraph of \mathcal{H} induced on Y* as

$$\mathcal{H}[Y] := (Y, \{h \cap Y : h \cap Y \neq \emptyset, h \in H\}).$$

For a subset $H' \subseteq H$, the hypergraph $\mathcal{H}(H') = (X, H')$ is the *partial hypergraph of \mathcal{H} generated by H'* . A *partial subhypergraph of \mathcal{H}* is a partial hypergraph of an induced subhypergraph. Lastly, for a set $Y \subseteq X$ we call for $H \upharpoonright Y := \{h \in H : H \subseteq Y\}$ the partial subhypergraph $\mathcal{H} \upharpoonright Y := (Y, H \upharpoonright Y)$ the *restriction of \mathcal{H} to Y* .

For a positive integer k , a map $c: X \rightarrow [k]$ is a *k -colouring of \mathcal{H}* if no hyperedge of size at least 2 is monochromatic, i.e. $|c(h)| > 1$ for each $h \in H$ with $|h| > 1$. If there is a k -colouring of \mathcal{H} , then we say \mathcal{H} is *k -colourable*.

G.3. On the capacitated version

A map $c: E(D) \rightarrow K$, where K denotes a set of cardinals, is called a *capacity* of D . If $K = \mathbb{N}$ (the set of non-negative integers), then we say c is *finitary*. We say a capacity $c: E(D) \rightarrow \{0, 1\}$ is *simple*.

Given a capacity c and a subset $A \subseteq E(D)$, then by $c(A)$ we denote the cardinality of the disjoint union of $\{c(a) : a \in A\}$ (which is considered to be the cardinal sum of the cardinals). If this cardinal is finite, then we say A has *finite capacity*. For

a finitary capacity c and a finite $A \subseteq E(D)$ this coincides to the integer $\sum_{e \in A} c(e)$, and for a simple capacity c this coincides with the cardinal $|A \cap c^{-1}(1)|$.

Given a capacity c of D , a dibond $B \in \mathfrak{B}$ is c -cheapest in \mathfrak{B} if $c(B)$ is minimum among all dibonds in \mathfrak{B} , and a set of dijoins is c -disjoint if each edge $e \in E(D)$ is contained in at most $c(e)$ dijoins of that set. We say D is \mathfrak{B} -Menger with respect to c if the size of a c -cheapest dibond in \mathfrak{B} is equal to the size of a largest set of c -disjoint \mathfrak{B} -dijoins. This leads to the following capacitated version of Question G.1.2.

Question G.3.1. Which digraphs D with capacity c and classes of dicuts \mathfrak{B} are \mathfrak{B} -Menger with respect to c ?

In this setting, the conjecture of Edmonds and Giles [32] (which was proven to be false by Schrijver [93], cf. Figure G.1) translates to the statement that every finite digraph D with capacity c of D is $\mathfrak{B}_{\text{dibond}}$ -Menger with respect to c , where $\mathfrak{B}_{\text{dibond}}$ denotes the class of all finite dibonds of D .




Figure G.1.: Schrijver’s counterexample to the conjecture of Edmonds and Giles.

The dashed edges have capacity 0, and the solid edges have capacity 1.

Given a capacity c of a digraph D one can replace each edge of e of positive capacity by a set of $|c(e)|$ many distinct parallel edges. Defining the capacity of each of these newly created edges as 1 we obtain a digraph with a simple capacity, and it is not hard to see that the question whether for a class \mathfrak{B} of dicuts is \mathfrak{B} -Menger with respect to c is equivalent to the corresponding question for the ‘corresponding’ class of dicuts of D' . In fact, the key feature of any known counterexample of the conjecture of Edmonds and Giles is to assign capacity 0 to some edges. Going one step further by deleting the edges of capacity 0 allows us equivalently talk about Question G.3.1 in the setting of Question G.1.2, as the following construction and proposition shows.

Construction G.3.2. Given a digraph D with a capacity c of D we define a digraph \hat{D} .

Let \hat{D} be the digraph on $V(D)$ obtained by replacing each edge e of D with $c(D)$ many distinct edges $\{e_\alpha : 0 \leq \alpha < c(D)\}$, each of which has the same head and tail as e . Note that any edge of capacity 0 is deleted in this process.

For each dicut B of D represented by (X, Y) , we define a corresponding dicut \hat{B} of \hat{D} as $\hat{B} := E_{\hat{D}}(X, Y)$. It is not hard to see that this is indeed well-defined and that the capacity of B equals the size of \hat{B} .

Moreover, for a class \mathfrak{B} of dicuts of D we define $\hat{\mathfrak{B}}$ as $\{\hat{B} : B \in \mathfrak{B}\}$.

Proposition G.3.3. *Let D be a digraph, let c be a capacity of D , and let \mathfrak{B} be a class of dicuts of D . Then D is \mathfrak{B} -Menger with respect to c if and only if \hat{D} is $\hat{\mathfrak{B}}$ -Menger.*

Proof. For a set \hat{A} of edges of \hat{D} we call the set $\text{tr}(\hat{A}) := \{e \in E(D) : e_\alpha \in \hat{A}\}$ the *trace* of \hat{A} . For a set A of edges of D and a set \hat{A} of edges of \hat{D} , we say the pair (A, \hat{A}) is *compatible* if $A = \text{tr}(\hat{A})$. Note that for each set A of edges e of D such that $c(e) > 0$ for all $e \in A$, there is a set of edges of \hat{D} such that (A, \hat{A}) is compatible.

For a compatible pair (F, \hat{F}) it is easy to see that F is a \mathfrak{B} -dijoin of D if and only if \hat{F} is a $\hat{\mathfrak{B}}$ -dijoin of \hat{D} . Moreover, if $\hat{\mathcal{F}}$ is a set of disjoint $\hat{\mathfrak{B}}$ -dijoins of \hat{D} , then $\{\text{tr}(\hat{F}) : \hat{F} \in \hat{\mathcal{F}}\}$ is c -disjoint.

Given a c -disjoint set $\{F_\alpha : \alpha \in \kappa\}$ of \mathfrak{B} -dijoins of D for some cardinal κ we define a set of disjoint $\hat{\mathfrak{B}}$ -dijoins iteratively as follows. In step α we construct \hat{F}_α by taking for each $e \in F_\alpha$ the edge e_β for the smallest ordinal β such that $e_\beta \notin \bigcup \{\hat{F}_\gamma : \gamma < \alpha\}$. By the assumption that e is contained in at most $c(e)$ many of the dijoints, this construction is well-defined and $(F_\alpha, \hat{F}_\alpha)$ is compatible. Hence we have the desired equivalence. \square

Recall that we say a class \mathfrak{B} of dicuts of a digraph D is *dibond-closed* if every dibond which is contained in some dicut in \mathfrak{B} is contained in \mathfrak{B} as well. In the context of Question G.1.2 it is quite natural to consider classes which are dibond-closed. Indeed, whenever we consider Question G.1.2 for a dibond-closed class \mathfrak{B} we can equivalently consider the question for the class of dibonds in \mathfrak{B} . But by considering the construction from Proposition G.3.3 we may destroy weak connectivity by deleting the capacity 0 edges. Therefore, a dibond in \mathfrak{B} may

correspond to a dicut in $\hat{\mathfrak{B}}$ which is not a dibond. In this way, Question G.3.1 for dibonds can be thought of as a special case of Question G.1.2, albeit for a slightly modified digraph.

In fact, the reverse is also true, as the following construction shows.

Construction G.3.4. Given a digraph D we define a digraph \tilde{D} with a capacity \tilde{c} of \tilde{D} .

For every subset $S \subseteq V(D)$ we add a distinct vertex v_S and edges $v_S s$ for each $s \in S$. Now we define a capacity \tilde{c} on \tilde{D} by setting $c(e) = 1$ if e is an edge of D and 0 otherwise.

For each dicut B of D represented by (X, Y) , we define a corresponding dibond \tilde{B} of \tilde{D} as follows. We define $X_B \subseteq V(\tilde{D})$ as the union of $\text{in}_D(X, Y)$ with $\{v_S : S \subseteq \text{in}_D(X, Y)\}$. Then we define \tilde{B} as the set of ingoing edges of X_B in \tilde{D} , i.e. $E_D(V(D) \setminus X_B, X_B)$. It is not hard to see that this is indeed well-defined and that the size of B equals the capacity of \tilde{B} since the edges of capacity 1 that \tilde{B} contains are precisely the edges in B .

Moreover, for a class \mathfrak{B} of dicuts of D we define $\tilde{\mathfrak{B}}$ as $\{\tilde{B} : B \in \mathfrak{B}\}$.

Proposition G.3.5. *Let D be a digraph and let \mathfrak{B} be a class of dicuts of D . Then D is \mathfrak{B} -Menger if and only if \tilde{D} is $\tilde{\mathfrak{B}}$ -Menger with respect to \tilde{c} .*

Proof. Since by construction $B \subseteq \tilde{B}$ for every dicut B of D we get that every \mathfrak{B} -dijoin of D is a $\tilde{\mathfrak{B}}$ -dijoin of \tilde{D} . Vice versa, every $\tilde{\mathfrak{B}}$ -dijoin of \tilde{D} that does not contain any capacity 0 edges is a \mathfrak{B} -dijoin of D . Since the capacity \tilde{c} is simple, the proposition immediately follows. \square

Therefore Question G.3.1 for classes of dibonds can really be thought of Question G.1.2 for non dibond-closed classes of dicuts.

In fact this observation will yield capacitated versions of many of our main results as it is easy to observe that if a class \mathfrak{B} of dibonds of a digraph D is nested, then so are $\hat{\mathfrak{B}}$ and $\tilde{\mathfrak{B}}$.

G.4. The dicut hypergraph

Let D be a digraph and let \mathfrak{B} be a class of dicuts of D . Then $\mathcal{H}(D, \mathfrak{B}) := (E(D), \mathfrak{B})$ is the \mathfrak{B} -dicut hypergraph of D . Note that a transversal of $\mathcal{H}(D, \mathfrak{B})$ is a \mathfrak{B} -dijoin.

If \mathfrak{B} is the set of all dicuts of D , then we denote by $\mathcal{H}(D)$ the *dicut hypergraph* $\mathcal{H}(D, \mathfrak{B})$.

G.4.1. Transversal packings of hypergraphs

Let $\mathcal{H} = (X, H)$ be a hypergraph. A tuple $(x_1, h_1, x_2, h_2, \dots, x_n, h_n, x_{n+1})$ is a *Berge-cycle* of \mathcal{H} if

- (1) $x_1, \dots, x_n \in X$ are distinct;
- (2) $h_1, \dots, h_n \in H$ are distinct;
- (3) $x_{n+1} = x_1$; and
- (4) $x_i, x_{i+1} \in h_i \in H$ for all $i \in [n]$.

The *length* of the Berge-cycle is n . A Berge-cycle of odd length is called *odd*. We call a Berge-cycle *improper* if some hyperedge h_i contains some x_j for $j \notin \{i, i+1\}$.

We call \mathcal{H} *balanced* if every odd Berge-cycle is improper. Balanced hypergraphs are one type of generalisation of bipartite graphs. By a theorem of Berge, finite balanced hypergraphs contain k pairwise disjoint transversals for k being the minimum size of a hyperedge. For the sake of completeness we will include a proof of this theorem.

Theorem G.4.1 (Berge [9, Corollary 2 of Section 5.3]). *Every finite balanced hypergraph $\mathcal{H} = (X, H)$ contains $k := \min_{h \in H} |h|$ disjoint transversals of \mathcal{H} .*

In order to provide a proof of this theorem, we will use the following characterisation of balanced hypergraphs.

Theorem G.4.2 (Berge [9, Theorem 7 of Section 5.3]). *A finite hypergraph \mathcal{H} is balanced, if and only if every induced subhypergraph of \mathcal{H} is 2-colourable.*

Proof. If \mathcal{H} contains an odd Berge-cycle $(x_1, h_1, x_2, h_2, \dots, x_n, h_n, x_{n+1})$ which is not improper, then the hypergraph induced on $\{x_i : i \in [n]\}$ contains the edges of an odd cycle, and hence is not 2-colourable.

For the other direction it is enough to show that a finite balanced hypergraph is 2-colourable. Suppose for a contradiction that $\mathcal{H} = (X, H)$ is a counterexample with a ground set X of minimum size. From the minimality we can deduce

that each $x \in X$ is contained in at least two hyperedges of size 2 and hence the subgraph G of \mathcal{H} containing all hyperedges of size 2 has minimum degree 2. Since \mathcal{H} is balanced, G is bipartite. Let $x \in X$ be such that it is no cut-vertex of G . The hypergraph induced on $X \setminus \{x\}$ has a 2-colouring by the minimality assumption. Since G is bipartite and x is not a cut-vertex, the neighbourhood of x in G is monochromatic. Hence we can extend the 2-colouring to \mathcal{H} , contradicting that \mathcal{H} is a counterexample. \square

Given this theorem, we can prove Theorem G.4.1.

Proof of Theorem G.4.1. For $k = 1$ the statement is obvious and if $k = 2$, the statement follows directly from Theorem G.4.2 since each colour class of a proper 2-colouring is a transversal. If $k > 2$, consider a k -colouring c for which the sum $\sum_{h \in H} |c(h)|$ is as large as possible. Note that if this sum equals $k|H|$, then each colour class of c is a transversal. So suppose for a contradiction that the sum is smaller. Then there is a hyperedge h_0 with $|c(h_0)| < k$. Since by assumption every hyperedge has size at least k , there is a colour p appearing twice on h_0 as well as a colour q not appearing on h_0 . Consider the subhypergraph induced on the colour classes S_p and S_q of these two colours. By Theorem G.4.2, this hypergraph has a 2-colouring c' . But then

$$\hat{c}(x) := \begin{cases} c(x) & \text{if } x \notin S_p \cup S_q \\ c'(x) & \text{if } x \in S_p \cup S_q \end{cases}$$

defines a k -colouring for which $\sum_{h \in H} |\hat{c}(h)| > \sum_{h \in H} |c(h)|$, a contradiction. \square

The technique of the compactness principle in combinatorics allows us to push these results about finite hypergraphs to infinite hypergraphs of *finite character*, i.e. hypergraphs in which every hyperedge is finite. We omit stating the compactness principle here but refer to [24, Appendix A].

Lemma G.4.3. *Let $\mathcal{H} = (X, H)$ be a hypergraph of finite character such that for each finite $Y \subseteq X$ there is a finite $\bar{Y} \subseteq X$ containing Y such that $\mathcal{H}|_{\bar{Y}}$ contains $\min_{h \in H|_{\bar{Y}}} |h|$ disjoint transversals of $\mathcal{H}|_{\bar{Y}}$. Then \mathcal{H} contains $k := \min_{h \in H} |h|$ disjoint transversals of \mathcal{H} .*

Proof. We construct the k disjoint transversals of \mathcal{H} via compactness. Given sets Z and Y with $Z \subseteq Y \subseteq X$, note that for a transversal T of $\mathcal{H}|_Y$ the set $T \cap Z$

is a transversal of $\mathcal{H}|Z$. Hence, by the assumptions of this lemma and the compactness principle there is a set $\{T_i : i \in [k]\}$ of subsets of X such that for each finite $Y \subseteq X$ the set $\{T_i \cap \bar{Y} : i \in [k]\}$ is a set of k disjoint transversals of $\mathcal{H}|Y$. Hence $T_i \cap T_j = \emptyset$ for $i \neq j$. Moreover, each T_i meets any $h \in H$ since $T_i \cap \bar{h}$ is a transversal of $\mathcal{H}|\bar{h}$ which is a finite subhypergraph of \mathcal{H} since \mathcal{H} has finite character. Therefore, $\{T_i : i \in [k]\}$ is as desired. \square

Since every partial subhypergraph of a balanced hypergraph is again balanced (cf. [9, Proposition 1 of Section 5.3]), we obtain the following corollary of Lemma G.4.3 and Theorem G.4.1.

Corollary G.4.4. *Every balanced hypergraph $\mathcal{H} = (X, H)$ of finite character contains $k := \min_{h \in H} |h|$ disjoint transversals of \mathcal{H} .* \square

G.4.2. Applications to the dicut hypergraph

We now turn our attention back to the dicut hypergraph and will show that the dicut hypergraph for a nested class of dicuts is balanced. Note that the following lemma holds for infinite digraphs as well.

Lemma G.4.5. *For any digraph D , every Berge-cycle $(e_1, B_1, e_2, B_2, \dots, e_n, B_n, e_1)$ of the dicut hypergraph $\mathcal{H}(D)$ of D which is odd and for which the set $\{B_i : i \in [n]\}$ is nested, is improper.*

Proof. Let $(e_1, B_1, e_2, B_2, \dots, e_n, B_n, e_1)$ be an odd Berge-cycle and furthermore let $\{(X_i, Y_i) : i \in [n]\}$ be a nested set of bipartitions of $V(D)$ representing $\{B_i : i \in [n]\}$ such that (X_i, Y_i) represents B_i and $\text{in}_D(X_i, Y_i) = Y_i$ for all $i \in [n]$.

By setting $B_0 := B_n$, $X_0 := X_n$ and $Y_0 := Y_n$, since $e_{i+1} \in B_i \cap B_{i+1}$ for all $0 \leq i < n$, we get either $(X_i, Y_i) \leq (X_{i+1}, Y_{i+1})$ or $(X_{i+1}, Y_{i+1}) \leq (X_i, Y_i)$. While n is odd, these two possibilities cannot occur in an alternating fashion throughout the whole cycle. Hence we may assume without loss of generality that either $(X_n, Y_n) \leq (X_1, Y_1) \leq (X_2, Y_2)$ or $(X_2, Y_2) \leq (X_1, Y_1) \leq (X_n, Y_n)$. We continue the argument with the former inequality, the other case is symmetric.

Consider the set I of all $i \in [n]$ for which either $(X_1, Y_1) \leq (X_i, Y_i)$ or $(Y_i, X_i) \leq (X_1, Y_1)$. Since $2 \in I$ and $n \notin I$, there is an integer j with $2 \leq j < n$ with $j \in I$ and $j+1 \notin I$. And since $B_j \cap B_{j+1}$ is nonempty and hence (X_j, Y_j) and

(X_{j+1}, Y_{j+1}) are \leq -comparable, we get either $(Y_j, X_j) \leq (X_1, Y_1) \leq (Y_{j+1}, X_{j+1})$, or $(X_{j+1}, Y_{j+1}) \leq (X_1, Y_1) \leq (X_j, Y_j)$.

Note that the first case is not possible since e_{j+1} would be an edge with tail in Y_1 and head in X_1 , contradicting that Y_1 is the in-shore of B_1 . Hence, $(X_{j+1}, Y_{j+1}) \leq (X_1, Y_1) \leq (X_j, Y_j)$ and $e_{j+1} \in B_1 \cap B_j \cap B_{j+1}$, proving that the Berge-cycle $(e_1, B_1, e_2, B_2, \dots, e_n, B_n, e_1)$ is improper. \square

Hence, if \mathfrak{B} is a nested class of dicuts of D , Lemma G.4.5 shows that $\mathcal{H}(D, \mathfrak{B})$ is balanced, and with Theorem G.4.1 we get Theorem G.1.3 for finite digraphs, and together with Lemma G.4.3 we can complete the proof.

Theorem G.1.3. *Let D be a digraph and \mathfrak{B} be a nested class of finite dicuts of D . Then D is \mathfrak{B} -Menger.*

Proof. Let k denote the size of a smallest dicut in \mathfrak{B} . By Lemma G.4.5, the \mathfrak{B} -dicut hypergraph $\mathcal{H} := \mathcal{H}(D, \mathfrak{B})$ is balanced and by assumption has finite character. Each finite restriction of \mathcal{H} whose set of hyperedges is non-empty is balanced and hence contains k disjoint transversals by Theorem G.4.1. The theorem follows from Lemma G.4.3. \square

With the observations of Section G.3, we obtain the capacitated version of this theorem.

Corollary G.4.6. *Let D be a digraph with a capacity c and let \mathfrak{B} be a nested class of dibonds of D of finite capacity. Then D is \mathfrak{B} -Menger with respect to c .* \square

A dicut of D is *atomic* if it has a representation in which one the sides contains only a single vertex, i.e. a source or a sink. It is easily verified that a set of atomic dicuts is nested. Hence, we get the following corollary.

Corollary G.4.7. *Let D be a digraph (with a capacity c) and let \mathfrak{B} be a class of atomic dicuts of D of finite size (capacity). Then D is \mathfrak{B} -Menger (with respect to c).* \square

G.5. Minimum size dicuts and disjoint mini-dijoins

Let D be a digraph and let \mathfrak{B} be a class of dicuts of D . We say \mathfrak{B} is *corner-closed* if for each non-empty $B, B' \in \mathfrak{B}$ which are crossing and represented by (X, Y)

and (X', Y') , respectively, the dicuts

$$B \wedge B' := E_D(\text{out}_D(X, Y) \cup \text{out}_D(X', Y'), \text{in}_D(X, Y) \cap \text{in}_D(X', Y'))$$

and

$$B \vee B' := E_D(\text{out}_D(X, Y) \cap \text{out}_D(X', Y'), \text{in}_D(X, Y) \cup \text{in}_D(X', Y'))$$

are in \mathfrak{B} . Note that it is easy to see that $B \wedge B'$ and $B \vee B'$ are indeed non-empty dicuts whose definition does not depend on the choice of the representations of B and B' . In particular, for any representations (X, Y) and (X', Y') of B and B' , respectively, with $\text{in}_D(X, Y) = Y$ and $\text{in}_D(X', Y') = Y'$, the bipartition $(X \cup X', Y \cap Y')$ represents $B \wedge B'$ and the bipartition $(X \cap X', Y \cup Y')$ represents $B \vee B'$.

Moreover, we observe that for the digraphs \hat{D} and \tilde{D} from Constructions [G.3.2](#) and [G.3.4](#), respectively, for dicuts B and B' of D , we obtain

$$\hat{B} \wedge \hat{B}' = \widehat{B \wedge B'}, \quad \hat{B} \vee \hat{B}' = \widehat{B \vee B'}, \quad \tilde{B} \wedge \tilde{B}' = \widetilde{B \wedge B'}, \quad \text{and} \quad \tilde{B} \vee \tilde{B}' = \widetilde{B \vee B'}.$$

In particular, a class \mathfrak{B} of dicuts of D is corner-closed, if and only if $\hat{\mathfrak{B}}$ is corner-closed, if and only if $\tilde{\mathfrak{B}}$ is corner-closed.

Remark G.5.1. Similarly to the proof of [Theorem G.4.1](#), we can consider a set $\{F_i : i \in [m]\}$ of disjoint \mathfrak{B} -dijoins of a digraph D as a partial colouring of the edges of D where an edge e is coloured with the colour i if and only if $e \in F_i$. We call such a colouring $f : \bigcup\{F_i : i \in [m]\} \rightarrow [m]$ a *\mathfrak{B} -Menger colouring of D* .

Since each F_i is a \mathfrak{B} -dijoin, we obtain that each dicut $B \in \mathfrak{B}$ is coloured with every colour. Note that if $|B| = m$, then B is necessarily *colourful*, i.e. B contains every colour.

Theorem G.5.2. *Let D be a finite digraph (with capacity c), let m be a positive integer and let \mathfrak{B} denote a corner-closed class of dicuts of D all of size m (capacity m). Then D is \mathfrak{B} -Menger (with respect to c).*

Proof. The capacitated version of this theorem follows from the non-capacitated version by the observations of both [Section G.3](#) and above. We prove the non-capacitated version by induction on the number of non-atomic dicuts in \mathfrak{B} . If \mathfrak{B} contains only atomic dicuts, then the statement follows from [Corollary G.4.7](#).

Otherwise, let $B \in \mathfrak{B}$ be non-atomic represented by (X, Y) with $\text{in}_D(X, Y) = Y$. Consider the digraph D_1 obtained by identifying all vertices in Y to a single vertex (and deleting loops, afterwards) with \mathfrak{B}_1 being the class of dicuts in \mathfrak{B} that are dicuts of D_1 , as well as the digraph D_2 obtained by identifying all vertices in X to a single vertex (and deleting loops, afterwards) with \mathfrak{B}_2 being the class of dicuts in \mathfrak{B} that are dicuts of D_2 . By construction, $E(D_1) \cap E(D_2) = B$. Note that in both \mathfrak{B}_1 and \mathfrak{B}_2 the number of non-atomic dicuts strictly decreased since B is atomic in both D_1 and D_2 and each non-atomic dicut of D_1 or D_2 is non-atomic in D as well. By induction, for $j \in \{1, 2\}$ there are sets $\{F_i^j : i \in [m]\}$ of disjoint \mathfrak{B}_j -dijoins of D_j . Note that since $|B| = m$, for each $e \in B$ there is a unique $i_e \in [m]$ and a unique $j_e \in [m]$ such that $e \in F_{i_e}^1 \cap F_{j_e}^2$. We claim that $\{F_{i_e}^1 \cup F_{j_e}^2 : e \in B\}$ is a set of disjoint \mathfrak{B} -dijoins. As in Remark G.5.1, we consider these edges sets as partial colourings of $E(D)$ with colours B .

The fact that the sets are pairwise disjoint follows from the observation that they are the union of disjoint dijoints of D_1 and D_2 respectively and that $E(D_1) \cap E(D_2) = B$.

Note that any bipartition (X', Y') of $V(D)$ which is nested with (X, Y) naturally defines a bipartition of $V(D_1)$ and $V(D_2)$, one of which is trivial and the other representing a cut of the respective digraph which equals $E_D(X', Y')$. Let $B' \in \mathfrak{B}$ be represented by (X', Y') with $\text{in}_D(X', Y') = Y'$ and assume $B' \notin \mathfrak{B}_1 \cup \mathfrak{B}_2$. Consider the corners $B \wedge B'$ and $B \vee B'$ represented by $(X \cup X', Y \cap Y')$ and $(X \cap X', Y \cup Y')$, respectively. Since $B' \notin \mathfrak{B}_1 \cup \mathfrak{B}_2$, both of the corners are non-empty, and since \mathfrak{B} is corner-closed, they are in \mathfrak{B} as well. Furthermore, since both $(X \cup X', Y \cap Y')$ and $(X \cap X', Y \cup Y')$ are nested with (X, Y) , we obtain that $B \vee B' \in \mathfrak{B}_1$ and $B \wedge B' \in \mathfrak{B}_2$. Note that since both $B \wedge B'$ and B are colourful by assumption, every colour that appears on $B \cap E(D[X'])$ also appears on $B' \cap E(D[Y'])$, and since $B \vee B'$ and B are colourful, every colour that appears on $B \cap E(D[Y'])$ also appears on $B' \cap E(D[X])$. Together with the colours appearing in $B \cap B'$ we deduce that B' is colourful as well, as desired. \square

Lemma G.5.3. *Let D be a digraph (with capacity c) that contains a dicut of finite size (capacity). The class \mathfrak{B}_{\min} of dicuts of D of minimum size (capacity) is corner-closed.*

Proof. Let $B_1, B_2 \in \mathfrak{B}_{\min}$ be crossing. Hence neither $B_1 \wedge B_2$ nor $B_1 \vee B_2$ are

empty. A simple double counting argument yields the equation $c(B_1) + c(B_2) = c(B_1 \wedge B_2) + c(B_1 \vee B_2)$. Therefore both $B_1 \wedge B_2$ and $B_1 \vee B_2$ are of capacity $c(B_1)$, since neither of them can be smaller. Hence they are dicuts of minimum capacity and hence in \mathfrak{B}_{\min} . \square

Now we can deduce the version of Theorem G.1.4 for finite digraphs as a direct corollary of Theorem G.5.2 and Lemma G.5.3.

Theorem G.5.4. *Every finite digraph D (with finitary capacity c) is \mathfrak{B}_{\min} -Menger (with respect to c), where \mathfrak{B}_{\min} denotes the class of dicuts of D of minimum size (capacity).* \square

G.6. Dijoins in infinite digraphs

First, we should ask what the ‘right’ generalisation of Conjecture G.1.1 for infinite digraphs should be. Often a more structural generalisation of such min-max theorems yields a more meaningful result than a simple generalisation using cardinalities. For example, Erdős conjectured and Aharoni and Berger [1] proved a structural generalisation of Menger’s theorem, where in every graph we can simultaneously find a set \mathcal{P} of disjoint paths between two sets A and B of vertices and a set S of vertices separating A and B such that each path in \mathcal{P} contains precisely one vertex in S and S contains no vertices not included in some path in \mathcal{P} .

However, such an Erdős-Menger-like structural generalisation of Conjecture G.1.1 fails in infinite digraphs, as we will illustrate in the following example.

Example G.6.1. We give an example of an infinite digraph with no pair a dicut B and a set of disjoint dijoins $\{F_e : e \in B\}$ such that $B \cap F_e = \{e\}$ for all $e \in B$.

Consider the digraph D with vertex set $V(D) := \mathbb{Z} \times \{1, -1\}$ and edge set

$$E(D) := \{(z, i)(z + i, i) : z \in \mathbb{Z}, i \in \{1, -1\}\} \cup \{(z, 1)(z, -1) : z \in \mathbb{Z}\},$$

as depicted in Figure G.2. We call the edges $(z, 1)(z, -1)$ of the second type the *rungs* of D .

Note that D has no finite dicut and for each infinite dicut B there is an $n \in \mathbb{Z}$ such that B contains all rungs $\{(z, 1)(z, -1) : z \leq n\}$ up to n . Moreover, whenever F_1




Figure G.2.: A counterexample to an Erdős-Menger-like structural generalisation of Conjecture G.1.1 to infinite digraphs.

and F_2 are two disjoint dijoins, then at least one of them contains infinitely many rungs of $\{(z, 1)(z, -1) : z \leq n\}$ up to any $n \in \mathbb{Z}$. Hence every dicut meets such a dijoin infinitely often. Moreover, note that in this example, the set of dicuts is nested.

In the light of this example, we can only hope for weaker generalisations to be possible.

In Subsection G.6.1, we will consider Question G.1.2 for classes of dicuts of finite size (capacity), where the structural generalisation as hinted above is still equivalent to a comparison of cardinals. In this setting, we can extend the respective results from the finite case using the compactness principle (or more precisely, Lemma G.4.3). Such an approach will not work for classes of dicuts that contain dicuts of both finite and infinite size (capacity).

In Subsection G.6.2, we will consider Question G.1.2 for classes of infinite dibonds, where we will prove a cardinality-version of this question for any class of infinite dibonds. We will also show that even a cardinality-version for classes of infinite dicuts (and hence a capacitated version, cf. Question G.3.1) can fail.

G.6.1. Classes of finite dicuts in infinite digraphs

Let D be a digraph. Given a set \mathcal{B} of dicuts of D , we define an equivalence relation on $V(D)$ by setting $v \equiv_{\mathcal{B}} w$ if and only if we cannot separate v from w by a dicut in \mathcal{B} .

It is easy to check that $\equiv_{\mathcal{B}}$ indeed defines an equivalence relation. Let $D/\equiv_{\mathcal{B}}$ denote the digraph which is obtained from D by identifying the vertices in the same equivalence class of $\equiv_{\mathcal{B}}$ and deleting loops. Note that $D/\equiv_{\mathcal{B}}$ does not contain any directed cycles. Given a capacity c of D , we call the restriction of c to $E(D/\equiv_{\mathcal{B}})$ the *capacity of $D/\equiv_{\mathcal{B}}$ induced by c* . We shall use the following observation from [40] about this digraph.

Proposition G.6.2. [40, Proposition 2.9(ii)] *Let D be a digraph and let \mathcal{B} be a set of dicuts of D . Then every dicut (or dibond, respectively) in \mathcal{B} of D is also a dicut (or dibond, respectively) of $D/\equiv_{\mathcal{B}}$.*

The main tool we use to extend results about Question G.1.2 from classes of finite graphs to finitary infinite versions is the following compactness-type lemma.

Lemma G.6.3. *Let D be a digraph (with capacity c) and let \mathfrak{B} be a class of dicuts of D of finite size (capacity). Suppose that for every finite set $\mathcal{B} \subseteq \mathfrak{B}$ there is a finite set $\overline{\mathcal{B}} \subseteq \mathfrak{B}$ containing \mathcal{B} such that the digraph $D/\equiv_{\overline{\mathcal{B}}}$ is $\overline{\mathcal{B}}$ -Menger (with respect to the capacity induced by c). Then D is \mathfrak{B} -Menger (with respect to c).*

Proof. Let k denote the capacity of a c -cheapest dicut in \mathfrak{B} and let $E' \subseteq E(D)$ denote the set of all edges of D of positive capacity. Consider the $\hat{\mathfrak{B}}$ -dicut hypergraph $\mathcal{H} := \mathcal{H}(\hat{D}, \hat{\mathfrak{B}})$ for \hat{D} and $\hat{\mathfrak{B}}$ as in Construction G.3.2. It is easy to observe that \mathcal{H} has finite character.

Moreover, note that for a finite set $\mathcal{B} \subseteq \mathfrak{B}$ containing a dicut of capacity k and for a $\overline{\mathcal{B}}$ as in the assumption, the hypergraph $\mathcal{H}_{\mathcal{B}} := \mathcal{H}(\hat{D}/\equiv_{\overline{\mathcal{B}}}, \hat{\mathcal{B}})$ has by Proposition G.3.3 and the assumption k disjoint transversals. By construction $\mathcal{H}_{\mathcal{B}}$ is a restriction of \mathcal{H} to a finite set. Moreover, for each restriction of \mathcal{H} to a finite set $Y \subseteq E(\hat{D})$ there is a finite set $\overline{Y} \subseteq E(\hat{D})$ containing Y such that $\mathcal{H}|_{\overline{Y}} = \mathcal{H}_{\mathcal{B}}$ for some finite $\mathcal{B} \subseteq \mathfrak{B}$. Hence the result follows from Lemma G.4.3 and again Proposition G.3.3. \square

Lemma G.6.3 together with Theorem G.5.2 yield the following corollary.

Corollary G.6.4. *Let D be a digraph (with capacity c), let m be a positive integer and let \mathfrak{B} denote a corner-closed class of dicuts of D all of size m (capacity m). Then D is \mathfrak{B} -Menger (with respect to c). \square*

Recall that by Lemma G.5.3 and the observations in Section G.5, applying Construction G.3.2 yields that $\hat{\mathfrak{B}}_{\min}$ is corner-closed. Hence we deduce the following corollary (and hence Theorem G.1.4).

Corollary G.6.5. *Let D be a digraph (with capacity c) that contains a dicut of finite size (capacity) and let \mathfrak{B}_{\min} be the set of dicuts of minimum size (capacity). Then D is \mathfrak{B}_{\min} -Menger. \square*

Lemma G.6.3 also yields the following corollary.

Corollary G.6.6. *If Conjecture G.1.1 is true for all weakly connected finite digraphs, then every weakly connected digraph D is $\mathfrak{B}_{\text{fin}}$ -Menger for the class $\mathfrak{B}_{\text{fin}}$ of finite dicuts of D . \square*

The *triangular prism* $K_3 \square K_2$ is the undirected graph with vertex set $V(K_3) \times V(K_2)$ and edges between (v_1, w_1) and (v_2, w_2) if and only if both $v_1 v_2 \in E(K_3)$ and $w_1 w_2 \in E(K_2)$. As mentioned in the introduction, the capacitated version Conjecture G.1.1 has been verified for planar digraphs with no minor isomorphic to the triangular prism $K_3 \square K_2$ by Lee and Williams.

Theorem G.6.7. [74] *Every finite weakly connected digraph D (with finitary capacity c) whose underlying multigraph is planar and contains no minor isomorphic to the triangular prism $K_3 \square K_2$ is \mathfrak{B} -Menger (with respect to c) for the class \mathfrak{B} of all dicuts of D .*

In order to extend Theorem G.6.7 to infinite digraphs, we begin by observing that each finite set \mathcal{B} of finite dicuts of D can be extended to a finite set $\overline{\mathcal{B}}$ of finite dicuts of D such that the auxiliary graphs $D/\equiv_{\overline{\mathcal{B}}}$ is a minor[†] of D . If instead \mathcal{B} is a finite set of dicuts of finite capacity, then we can still extend \mathcal{B} to a finite set $\overline{\mathcal{B}}$ and find a minor D^* of D which is \mathcal{B}^* -Menger with respect to the capacity obtained by restricting c to $E(D^*)$ for the set \mathcal{B}^* of all dicuts of D^* , as we will establish in the following lemma.

Lemma G.6.8. *Let D be a weakly connected digraph with capacity c , let \mathcal{B} be a finite set of dicuts of D of finite capacity. Then there is a finite set $\overline{\mathcal{B}}$ of dicuts of D of finite capacity with $\mathcal{B} \subseteq \overline{\mathcal{B}}$ and a finite minor D^* of D such that for the set \mathcal{B}^* of dicuts of D^* we have $\{B \setminus c^{-1}(0) : B \in \mathcal{B}^*\} = \{B \setminus c^{-1}(0) : B \in \overline{\mathcal{B}}\}$.*

Proof. Let $E' := \bigcup \mathcal{B}$. Let \mathcal{B}' be the set of all dicuts B of D with $B \subseteq E'$. Note that each dicut in \mathcal{B}' has finite capacity. Consider the digraph D' obtained from D by contracting $E(D) \setminus E'$. Note that the sets of dicuts of $D/\equiv_{\mathcal{B}'}$ and of D' coincide and both graphs have E' as their edge set. We claim that D' and $D/\equiv_{\mathcal{B}'}$ are

[†]We say a directed graph D is a minor of a directed graph D' , if D can be obtained from D' by an arbitrary sequence of vertex-deletions, edge-deletions and edge-contractions. Note that in the literature these type of minors are also known as *weak minors*.

isomorphic. Suppose for a contradiction that two distinct vertices v and w of D' cannot be separated by a dicut in \mathcal{B}' , and hence are contained in the same strong component of D' . Let P be a directed path from v to w in D' . As $D/\equiv_{\mathcal{B}'}$ contains no directed cycles, each edge of P is a loop of $D/\equiv_{\mathcal{B}'}$, contradicting its construction. Conversely, any vertices v and w which are not equivalent with respect to $\equiv_{\mathcal{B}'}$ do not lie in the same weak component of $D - E'$ and hence are not identified in D' .

If \mathcal{B}' is finite, then D' has only finitely many vertices and the result follows with $\bar{\mathcal{B}} := \mathcal{B}'$ and the digraph D^* obtained from D' by deleting for all pairs of vertices all but finitely many edges of capacity 0 between them. So let us assume \mathcal{B}' and hence D' is infinite.

Let W the set of all vertices of D' that are incident with some edge of positive capacity. If \mathcal{B}' contains a dicut with a representation (X, Y) for which $W \subseteq Y$, for some vertex $v_0 \in X$ we define $W' := W \cup \{v_0\}$, and $W' := W$ if no such dicut exists. For each non-empty proper subset $Z \subsetneq W$, let P_Z be a directed path from Z to $W \setminus Z$ in D' if such a path exist and let \mathcal{P} denote the set of all paths P_Z . Now let D^* be the digraph obtained from $D'[W'] \cup \mathcal{P}$ by deleting for each pair of vertices all but finitely many edges of capacity 0 between them.

Note that every component of D^* contains a vertex in W' and each edge of D^* is contained in a directed path between some vertices of W' . Consider a dicut B of D^* with a representation (X, Y) . As each component of D^* contains a vertex of W' , note that both $X \cap W'$ and $Y \cap W'$ are non-empty proper subsets of W' . If both $X \cap W$ and $Y \cap W$ are proper non-empty subsets of W , then by the choice of \mathcal{P} , there is a dicut B' of D' with a representation (X', Y') such that $X \cap W \subseteq X'$ and $Y \cap W \subseteq Y'$. Otherwise, one of X or Y is disjoint from W and is equal to $\{v_0\}$, thus B has capacity 0 and we can choose B' to be the dicut of D' separating $\{v_0\}$ from W . In particular, in both cases we obtain $B \setminus c^{-1}(0) = B' \setminus c^{-1}(0)$.

On the other hand, each dicut $B' \in \mathcal{B}'$ with a representation (X', Y') that separates W defines a dicut B of D^* represented by $(X' \cap V(D^*), Y' \cap V(D^*))$ for which we trivially obtain that $B \setminus c^{-1}(0) = B' \setminus c^{-1}(0)$. Lastly, if there is a dicut $B' \in \mathcal{B}'$ with a representation (X', Y') that does not separates W , then $c(B') = 0$ and $(\{v_0\}, V(D^*) \setminus \{v_0\})$ represents a dicut B with $B \setminus c^{-1}(0) = B' \setminus c^{-1}(0) = \emptyset$. Hence, the result follows with any finite $\bar{\mathcal{B}} \supseteq \mathcal{B}$ that for each $B \in \mathcal{B}'$ contains a $B' \in \mathcal{B}'$ with $B \setminus c^{-1}(0) = B' \setminus c^{-1}(0)$. \square

We now lift Theorem G.6.7 to infinite digraphs using Lemmas G.6.3 and G.6.8.

Corollary G.6.9. *Every weakly connected digraph D (with capacity c) whose underlying multigraph contains no minor isomorphic to either the triangular prism $K_3 \square K_2$, K_5 or $K_{3,3}$ is $\mathfrak{B}_{\text{fin}}$ -Menger (with respect to c) for the class $\mathfrak{B}_{\text{fin}}$ of dicuts of D of finite size (capacity).*

Proof. We may assume that D contains no dicuts of capacity 0 or else there is nothing to show. Consider a finite set \mathcal{B} of dicuts of D of finite capacity. By Lemma G.6.8 there is a finite set $\overline{\mathcal{B}}$ of dicuts with $\mathcal{B} \subseteq \overline{\mathcal{B}}$ of finite capacity and a finite minor D^* of D such that with \mathcal{B}^* denoting the set of dicuts of D^* and c^* denoting the capacity of D^* obtained from restricting c to $E(D^*)$, we have that $\{B \setminus c^{-1}(0) : B \in \overline{\mathcal{B}}\} = \{B \setminus c^{-1}(0) : B \in \mathcal{B}^*\}$. In particular, we conclude $D/\overline{\mathcal{B}}$ is $\overline{\mathcal{B}}$ -Menger with respect to c if and only if D^* is \mathcal{B}^* -Menger with respect to c^* .

Since $\overline{\mathcal{B}}$ contains no dicuts of capacity 0, we observe that $\emptyset \notin \mathcal{B}^*$ and hence that D^* is weakly connected. Since D does not contain a minor isomorphic to either $K_3 \square K_2$, K_5 or $K_{3,3}$, neither does D^* . Therefore, D^* is \mathcal{B}^* -Menger with respect to c^* by Theorem G.6.7. The result now follows from Lemma G.6.3. \square

Before we come to the next result we again have to introduce further notation.

A one-way infinite path is called a *ray* and the unique vertex of degree 1 in a ray is called its *start vertex*. An orientation of a ray R such that every vertex is oriented away from the start vertex of R is called a *forwards directed ray*, or briefly an *out-ray*. A *backwards directed ray*, or briefly a *back-ray*, is defined analogously.

For a weakly connected digraph D call a strongly connected component C of D a *source component* if no edge of D has its head in $V(C)$ and its tail in $V(D) \setminus V(C)$. A *sink component* of D is defined analogously. Furthermore, call a dicut B of D *sink-sided* (resp. *source-sided*) if $\text{out}_D(B)$ (resp. $\text{in}_D(B)$) contains neither a sink component (resp. source component) of D nor a out-ray (resp. back-ray) of D . A dicut of D that is either source-sided or sink-sided is called a *source-sink* dicut.

The following result is due to Feofiloff and Younger. We state it here adapted to our notation.

Theorem G.6.10. [33] *Every finite weakly connected digraph D (with finitary capacity c) is $\mathfrak{B}_{\text{s-s}}$ -Menger (with respect to c) for the class $\mathfrak{B}_{\text{s-s}}$ of all source-sink dicuts of D .*

Let us call a weakly connected digraph D *source-sink connected* if for every C^+ which is either a source component of D or a back-ray of D , and for every C^- which is either a sink component of D or a out-ray of D , there exists a directed path from C^+ to C^- in D .

With Theorem G.6.10, Feofiloff and Younger verified Conjecture G.1.1 for the class of finite source-sink connected digraphs since each dicut of a finite source-sink connected digraph is a source-sink dicut.

Now we shall lift Theorem G.6.10 to infinite graphs using Lemma G.6.3.

Corollary G.6.11. *Every weakly connected digraph D (with capacity c) is $\mathfrak{B}_{\text{s-s}}$ -Menger (with respect to c) for the class $\mathfrak{B}_{\text{s-s}}$ of all source-sink dicuts of D of finite size (capacity).*

Proof. Note that given a finite set $\mathcal{B} \subseteq \mathfrak{B}_{\text{s-s}}$ every $B \in \mathcal{B}$ is also a finite source-sink dicut of $D/\equiv_{\mathcal{B}}$. Hence, from Theorem G.6.10 we deduce that $D/\equiv_{\mathcal{B}}$ is \mathcal{B} -Menger with respect to the capacity obtained by restricting c to $E(D/\equiv_{\mathcal{B}})$, which indeed is finitary. The result now follows from Lemma G.6.3. \square

As for finite digraphs, this has an immediate consequence for source-sink connected digraphs regarding Conjecture G.1.1 and the class of all finite dicuts.

Corollary G.6.12. *Every weakly connected, source-sink connected digraph D (with capacity c) is $\mathfrak{B}_{\text{fin}}$ -Menger (with respect to c) for the class $\mathfrak{B}_{\text{fin}}$ of dicuts of D of finite size (capacity).*

Proof. The proof follows from Corollary G.6.11 and the observation that every dicut of D is a source-sink dicut. \square

G.6.2. Classes of infinite dibonds

In this subsection, we will prove Theorem G.1.5.

First we concentrate on the case where each dibond in the class has the same size as the digraph itself. Note that the following proof works for sets of bonds in undirected multigraphs as well.

Lemma G.6.13. *Let κ be an infinite cardinal, let D be a weakly connected digraph of size κ , and let \mathfrak{B} be a class of dibonds of D each of which has size κ . Then D is \mathfrak{B} -Menger.*

Proof. We build the dijoints inductively. For each $i < \kappa$ we start with empty sets F_i^0 . We fix an arbitrary enumeration $\{(i_\alpha, u_\alpha, v_\alpha) : \alpha < \kappa\}$ of the set $\kappa \times V(D) \times V(D)$.

Suppose for $\alpha < \kappa$ we already constructed a family of disjoint sets $(F_i^\alpha : i < \kappa)$ of edges such that $F^\alpha := \bigcup\{F_i^\alpha : i < \kappa\}$ has cardinality less than $|\alpha|^+ \cdot \aleph_0$. Let $X^\alpha \subseteq V(G)$ denote the set containing the end vertices of F^α as well as u_α and v_α . For each pair of distinct vertices $x, y \in X^\alpha$, let $P^\alpha(x, y)$ denote an undirected path between x and y in D which is edge disjoint to F^α if such a path exists, or let $P^\alpha(x, y) := \emptyset$ otherwise. We set

$$F_{i_\alpha}^{\alpha+1} := F_{i_\alpha}^\alpha \cup \bigcup\{P^\alpha(x, y) : x, y \in X^\alpha\} \text{ and } F_j^{\alpha+1} := F_j^\alpha \text{ for each } j \neq i_\alpha.$$

Note that $F^{\alpha+1} = \bigcup\{F_i^{\alpha+1} : i < \kappa\}$ has cardinality less than $|\alpha|^+ \cdot \aleph_0$, and hence we can continue the construction.

For a limit ordinal $\lambda \leq \kappa$, we set $F_i^\lambda := \bigcup\{F_i^\alpha : \alpha < \lambda\}$ for each $i < \kappa$. Note that by construction each F_i^λ has cardinality at most $|\lambda|$. Moreover, for all but at most λ many $i < \kappa$ the set F_i^λ is empty. Hence $F^\lambda := \bigcup\{F_i^\lambda : i < \kappa\}$ has cardinality at most $|\lambda|^2 = |\lambda| < |\lambda|^+$ and we can continue the construction as long as $\lambda < \kappa$.

Claim G.6.14. $F_{i_\alpha}^{\alpha+1}$ meets every dibond $B \in \mathfrak{B}$ separating u_α and v_α .

Proof of Claim G.6.14. By construction X^α meets both $\text{out}_D(B)$ and $\text{in}_D(B)$. Since B has size κ there is an edge in $B \setminus F^\alpha$, and since B separates some pair of vertices in X^α , there are vertices $x, y \in X^\alpha$ with $x \in \text{out}_D(B)$ and $y \in \text{in}_D(B)$ for which there is an undirected path between x and y which is edge disjoint to F^α . And since every such path meets B , so does $P^\alpha(x, y) \neq \emptyset$ and hence $F_{i_\alpha}^{\alpha+1}$. \square

With Claim G.6.14 we can deduce that the set $\{F_i^\kappa : i < \kappa\}$ is the desired set of disjoint \mathfrak{B} -dijoints. \square

A *decomposition* \mathcal{H} of a graph G is a set of subgraphs of G such that each edge of G is contained in a unique $H \in \mathcal{H}$. For an infinite cardinal κ , a decomposition of G is κ -*bond-faithful* if

- (1) each H has at most κ many edges;
- (2) any bond of size at most κ of G is a bond of some $H \in \mathcal{H}$; and
- (3) any bond of size less than κ of some $H \in \mathcal{H}$ is a bond of G .

Theorem G.6.15 (Laviolette [72, Theorem 3], Soukup [100, Theorem 6.3]). *For all infinite cardinals κ every graph has a κ -bond-faithful decomposition.*

Note that Laviolette originally only proved this theorem under the assumption of the generalised continuum hypothesis [72, Theorem 3]. This assumption was subsequently removed by Soukup using the technique of elementary submodels [100, Theorem 6.3].

Moreover, note that while this theorem was originally proven for simple graphs, it holds for multigraphs as well, and additionally we may assume that each graph in the decomposition is connected, as the following corollary summarises.

Corollary G.6.16. *For every infinite cardinal κ every multigraph has a κ -bond-faithful decomposition into connected graphs.*

Proof. For a multigraph G , consider a simple graph G' obtained by iteratively deleting parallel edges and loops. By Theorem G.6.15, this graph has a κ -bond-faithful decomposition \mathcal{H}' . For each $H \in \mathcal{H}'$ which is not connected, we replace it by its connected components to obtain a decomposition \mathcal{H}'' , which is again κ -bond-faithful as every bond is contained in a unique connected component. We construct a decomposition of G as follows. For any two vertices v and w such that there are at most κ many parallel edges between v and w , we add all of those edges to the unique simple graph $H \in \mathcal{H}''$ containing vw . Otherwise, we decompose the edges between v and w into sets of size κ , add one of those sets to the unique simple graph $H \in \mathcal{H}''$ containing vw , and for each other of those sets add a new graph to the decomposition consisting of precisely the edges in that set. Now it is easy to verify that the decomposition \mathcal{H} obtained in this manner is κ -bond-faithful. \square

We will use this concept to deduce Theorem G.1.5.

Proof of Theorem G.1.5. Let κ be the cardinality of a smallest dibond in \mathfrak{B} . Let \mathcal{H} be a κ -bond-faithful decomposition of the underlying multigraph as in Corollary G.6.16.

Every dibond B in \mathfrak{B} induces a dicut of size at most κ in some of the members of the κ -bond-faithful decomposition. This dicut cannot contain dibonds of size less than κ since such dibonds would be dibonds of D contained B , contradicting that B is a dibond.

To each $H \in \mathcal{H}$ we apply Lemma G.6.13 to the class \mathfrak{B}_H of those dibonds of H that are contained in some dicut of H that is induced by some dibond $B \in \mathfrak{B}$. Let $\{F_i^H : i < \kappa\}$ denote the set of disjoint \mathfrak{B}_H -dijoints of H . It is now easy to see that $\{\bigcup\{F_i^H : H \in \mathcal{H}\} : i < \kappa\}$ is a set of disjoint \mathfrak{B} -dijoints. \square

Finally, we show that a generalisation of this theorem to classes of dicuts fails. With the observations from Section G.3 we also see that a capacitated version of Theorem G.1.5 fails.

Example G.6.17. For any infinite cardinal κ consider the digraph D consisting of κ many pairwise non-incident edges e_α for all $\alpha < \kappa$, i.e. the edge set of D is

$$E := \{e_\alpha : \alpha < \kappa\}.$$

For every $I \subseteq \kappa$ let $B_I := \{e_\alpha : \alpha \in I\}$ denote the dicut consisting of the edges index by elements from I . Now consider the class of dicuts

$$\mathfrak{B} := \{B_I : I \subseteq \kappa, |I| = \kappa\}.$$

Note that any \mathfrak{B} -dijoin F of D has size at least κ since there are κ many disjoint dicuts in \mathfrak{B} . Moreover, note that $E \setminus F$ has size less than κ since $B_I \notin \mathfrak{B}$ for the set I for which $B_I = E \setminus F$. Hence, D does not contain two disjoint \mathfrak{B} -dijoints both contained in E .

However, we conjecture that the generalisation holds for nested classes of dicuts, which would yield with the observations from Section G.3 the capacitated version for nested classes of infinite dicuts.

Conjecture G.6.18. *Let D be a digraph and \mathfrak{B} be a nested class of infinite dicuts of D . Then D is \mathfrak{B} -Menger.*

H. Even circuits in oriented matroids

H.1. Introduction

Deciding whether a given digraph contains a directed cycle, briefly *dicycle*, of even length is a fundamental problem for digraphs and often referred to as the *even dicycle problem*. The computational complexity of this problem was unknown for a long time and several polynomial time equivalent problems have been found [64, 79, 80, 104]. The question about the computational complexity was resolved by Robertson, Seymour and Thomas [90] and independently by McCuaig [80] who stated polynomial time algorithms for one of the polynomially equivalent problems, and hence also for the even dicycle problem.

One of these polynomially equivalent problems makes use of the following definition.

Definition H.1.1 ([96]). Let D be a digraph. We call D *non-even*, if there exists a set J of directed edges in D such that every directed cycle C in D intersects J in an odd number of edges. If such a set does not exist, we call D *even*.

Seymour and Thomassen proved that the decision problem whether a given digraph is non-even, is polynomially equivalent to the even dicycle problem.

Theorem H.1.2 ([96]). *The problem of deciding whether a given digraph contains an even directed cycle, and the problem of deciding whether a given digraph is non-even, are polynomially equivalent.*

Furthermore, Seymour and Thomassen [96] characterised being non-even in terms of forbidden subgraphs. Their result can be stated more compactly by formulating it in terms of forbidden butterfly minors, which is a commonly used notion in directed graph structure theory [44, 61, 63], instead of forbidden subgraphs. Before we state their result, let us define the notion of butterfly minors and fix another notation.

Given a digraph D , an edge $e \in E(D)$ is called *butterfly-contractible* if it is not a loop and if it is either the unique edge emanating from its tail or the unique edge entering its head. A *butterfly minor* (sometimes also called *digraph minor* or just *minor*) of a digraph D is any digraph obtained from D by a finite sequence of edge-deletions, vertex-deletions and contractions of butterfly-contractible edges.

Note that the main idea behind the concept of a butterfly-contractible edge e within a digraph D is that every directed cycle in D/e either equals one in D or induces one in D by incorporating e . This property does not necessarily hold if arbitrary edges are contracted.

For every $k \geq 3$ let \overleftrightarrow{C}_k denote the symmetrically oriented cycle of length k (also called *bicycle*), i.e. the digraph obtained from C_k by replacing every edge by a pair of anti-parallel directed edges.

Now we can state the result of Seymour and Thomassen as follows.

Theorem H.1.3 ([96]). *A digraph D is non-even if and only if no butterfly minor of D is isomorphic to \overleftrightarrow{C}_k for some odd k .*

The main purpose of this work is to lift the even dicycle problem to oriented matroids, and to extend Theorem H.1.2 and partially Theorem H.1.3 to oriented matroids as well. Our main result (cf. Theorem H.1.9), subsumes Theorem H.1.3 together with a dual version in the setting of oriented matroids.

H.1.1. The Even Directed Circuit Problem in Oriented Matroids

In this paper we view a matroid as a tuple $M = (E, \mathcal{C})$ consisting of a finite *ground set* $E(M) := E$ containing the *elements* of M and the family \mathcal{C} of *circuits* of M .

In what follows we introduce a generalisation of the graph theoretic notion of being non-even to oriented matroids and state the main results of this work. For our purposes, the most important examples of matroids are *graphic matroids* and *cographic matroids*.

Let $G = (V, E)$ be a graph. The *cycle matroid* of G , denoted by $M(G)$, is the matroid (E, \mathcal{C}) where the set \mathcal{C} of circuits consists of all edge-sets of the cycles of G . Analogously, the *bond matroid* of G is $M^*(G) = (E, \mathcal{S})$ where \mathcal{S} is the set of bonds (or minimal non-empty edge cuts) of G . Note that $M(G)$ and $M^*(G)$ are the dual matroids of each another.

A matroid is called a *graphic matroid*, resp. a *cographic matroid* if it is, respectively, isomorphic to the cycle matroid or the bond matroid of some graph.

Digraphs can be seen as a special case of oriented matroids* in the sense that every digraph D has an associated oriented cycle matroid $M(D)$ whose signed circuits resemble the oriented cycles in the digraph D . In this spirit, it is natural to lift questions concerning cycles in directed graphs to more general problems on circuits in oriented matroids. The following algorithmic problem is the straight forward generalisation of the even dicycle problem to oriented matroids, and the main motivation of the paper at hand.

Problem H.1.4. Given an oriented matroid \vec{M} , decide whether there exists a directed circuit of even size in \vec{M} .

Our first contribution is to generalise the definition of non-even digraphs to oriented regular matroids in the following sense.

Definition H.1.5. Let \vec{M} be an oriented matroid. We call \vec{M} *non-even* if its underlying matroid is regular and there exists a set $J \subseteq E(\vec{M})$ of elements such that every directed circuit in \vec{M} intersects J in an odd number of elements. If such a set does not exist, we call \vec{M} *even*.

The reader might wonder why the preceding definition concerns only regular matroids. This has several reasons. The main reason is a classical result by Bland and Las Vergnas [12] which states that a binary matroid is orientable if and only if it is regular. Hence, if we were to extend the analysis of non-even oriented matroids beyond the regular case, we would have to deal with orientations of matroids which are not representable over \mathbb{F}_2 . This has several disadvantages, most importantly that cycle bases, which constitute an important tool in all of our results, are not guaranteed to exist any more. Furthermore, some of our proofs make use of the strong orthogonality property of oriented regular matroids[†], which fails for non-binary oriented matroids. Lastly, since Problem H.1.4 is an algorithmic question, oriented regular matroids have the additional advantage that they allow for a compact encoding in terms of totally unimodular matrices, which is not a given for general oriented matroids.

*For a formal and in-depth introduction of terms and notation used here please see Section H.2.

[†]For a definition we refer to Section H.2

The first result of this article is a generalisation of Theorem H.1.2 to oriented matroids as follows:

Theorem H.1.6. *The problems of deciding whether an oriented regular matroid represented by a totally unimodular matrix contains an even directed circuit, and the problem of recognising whether an oriented regular matroid given by a totally unimodular matrix is non-even, are polynomially equivalent.*

Theorem H.1.6 motivates a structural study of the class of non-even oriented matroids, as in many cases the design of a recognition algorithm for a class of objects is based on a good structural understanding of the class. In order to state our main result, which is a generalisation of Theorem H.1.3 to oriented graphic and cographic matroids, we have to introduce a new minor concept. We naturally generalise the concept of butterfly minors to regular oriented matroids, in the form of so-called *generalised butterfly minors*.

Definition H.1.7. Let \vec{M} be an orientation of a regular matroid M . An element $e \in E(\vec{M})$ is called *butterfly-contractible* if there exists a cocircuit S in M such that $(S \setminus \{e\}, \{e\})$ forms a signed cocircuit of \vec{M} .[‡] A *generalised butterfly minor* (*GB-minor for short*) of \vec{M} is any oriented matroid obtained from \vec{M} by a finite sequence of element deletions and contractions of butterfly-contractible elements.

Note that the order in which elements are deleted and butterfly-contractible elements are contracted can be modified as follows:

- If a butterfly-contractible element e is contracted and afterwards an element e' is deleted, then first deleting e' does not change the butterfly-contractibility of e .
- In case we first delete e and then contract the butterfly-contractible element e' , we may swap these operations if and only if e' is butterfly-contractible before the deletion of e .

Note that the generalised butterfly-contraction captures the same fundamental idea as the initial one for digraphs while being more general: Given a butterfly-contractible element e of a regular oriented matroid \vec{M} , we cannot have a directed

[‡]For a definition of a signed (co)circuits see Section H.2.

circuit C of \vec{M}/e such that $(C, \{e\})$ is a signed circuit of \vec{M} [§], and hence either C or $C \cup \{e\}$ must form a directed circuit of \vec{M} .

Replacing the notion of butterfly minors by GB-minors allows us to translate Theorem H.1.3 to the setting of oriented matroids in the following way:

Proposition H.1.8. *An oriented graphic matroid \vec{M} is non-even if and only if none of its GB-minors is isomorphic to $M(\overset{\leftrightarrow}{C}_k)$ for some odd $k \geq 3$.*

As our main result, we complement Proposition H.1.8 by determining the list of forbidden GB-minors for cographic non-even oriented matroids. We need the following notation: For integers $m, n \geq 1$ we denote by $\vec{K}_{m,n}$ the digraph obtained from the complete bipartite graph $K_{m,n}$ by orienting all edges from the partition set of size m towards the partition set of size n .

Theorem H.1.9. *An oriented cographic matroid \vec{M} is non-even if and only if none of its GB-minors is isomorphic to $M^*(\vec{K}_{m,n})$ for any $m, n \geq 2$ such that $m + n$ is odd.*

To prove Theorem H.1.9 we study those digraphs whose oriented bond matroids are non-even. Equivalently, these are the digraphs admitting an *odd dijoin*, which is an edge set hitting every directed bond an odd number of times. After translating GB-minors into a corresponding minor concept on directed graphs, which we call *cut minors*[¶], we show that the class of digraphs with an odd dijoin is described by two infinite families of forbidden cut minors (Theorem H.4.21). Finally, we translate this result to oriented cographic matroids in order to obtain a proof of Theorem H.1.9.

The structure of this paper is as follows. In Section H.2 we introduce the needed notation and basic facts about digraphs, matroids and oriented matroids for this paper. Furthermore, we prove that non-even oriented matroids are closed under GB-minors (Lemma H.2.5), which is then used to prove Proposition H.1.8 in the same section. We start Section H.3 by showing that the even directed circuit problem for general oriented matroids cannot be solved using only polynomially many calls to a signed circuit oracle (Proposition H.3.2). The remainder of

[§]In this case, $(C, \{e\})$ together with a signed cocircuit $(S \setminus \{e\}, \{e\})$ would contradict the orthogonality property (see Section H.2, $(*)$) for oriented matroids.

[¶]See the beginning of Section H.4 for a precise definition.

the section is devoted to the proof of Theorem [H.1.6](#). We also note that *odd* directed circuits can be detected in polynomial time in orientations of regular matroids (Proposition [H.3.15](#)). In Section [H.4](#) we characterise those digraphs that admit an odd dijoin (Theorem [H.4.21](#)) and use this to deduce our main result, Theorem [H.1.9](#).

H.2. Background

This section is dedicated to a formal introduction of basic terms and notation used throughout this paper. However, we assume basic familiarity with digraphs and matroid theory. For basic notation and facts about digraphs we refer the reader to [\[8\]](#). For missing terminology and basic facts from matroid theory not mentioned or mentioned without proof in the following, please consult the standard reading [\[84, 109\]](#).

For two sets X, Y we denote by $X \Delta Y := (X \cup Y) \setminus (X \cap Y)$ their *symmetric difference*. If X_1, \dots, X_k are several sets, then we denote by $\Delta_{i=1}^k X_i = X_1 \Delta X_2 \Delta \dots \Delta X_k$ the set of elements which appear in an odd number of the sets X_1, \dots, X_k . For $n \in \mathbb{N}$ we denote $[n] := \{1, 2, \dots, n\}$.

(Di)graphs

Graphs considered in this paper are multi-graphs and may include loops. Digraphs may have loops and multiple (parallel and anti-parallel) directed edges (sometimes called *edges*). Given a digraph D , we denote by $V(D)$ its vertex set and by $E(D)$ the set of directed edges. A directed edge with tail $u \in V(D)$ and head $v \in V(D)$ is denoted by (u, v) if this does not lead to confusion with potential parallel edges. By $U(D)$ we denote the *underlying multi-graph* of D , which is the undirected multi-graph obtained from D by forgetting the orientations of the edges. Given a digraph D and a partition (X, Y) of its vertex set, the set $D[X, Y]$ of edges with one endpoint in X and one endpoint in Y , if it is non-empty, is referred to as a *cut*. A cut of D is called *minimal* or a *bond*, if there is no other cut of D properly contained in it. It is well-known (cf. [\[24\]](#)) that if $U(D)$ is connected, then a cut $D[X, Y]$ is a bond if and only if both $D[X]$ and $D[Y]$ are weakly connected.

If there is no edge of D with head in X and tail in Y , the cut $D[X, Y]$ is called

directed and denoted by $\partial(X)$ (the set of edges leaving X). A *dijoin* in a digraph is a set of edges intersecting every directed cut (equivalently every directed bond).

Matroids

Matroids can be used to represent several algebraic and combinatorial structures of dependencies. The so-called *linear or representable* matroids are induced by vector configurations in linear spaces. Let $V = \mathbb{F}^n$ be a vector-space over a field \mathbb{F} and let $X = \{x_1, \dots, x_k\} \subseteq V$ for some $k \in \mathbb{N}$. Let A be the $n \times k$ -matrix over \mathbb{F} whose columns are x_1, \dots, x_k . Then we define the *column matroid induced by A* as $M[A] := (\{x_1, \dots, x_k\}, \mathcal{C}_A)$, where its set of circuits \mathcal{C}_A consists of the inclusion-wise minimal collections of linearly dependent vectors from $\{x_1, \dots, x_k\}$. It is a well-known fact that $M[A]$ is indeed a matroid for any choice of a matrix A . A matroid M is called \mathbb{F} -*linear* or *representable over the field \mathbb{F}* if there is a matrix A with entries in \mathbb{F} such that $M \simeq M[A]$. Graphic and cographic matroids, as introduced in Section H.1.1, form part of a larger class, the so-called *regular matroids*. A matroid M is called *regular* if it is \mathbb{F} -linear for every field \mathbb{F} . A fundamental property of regular matroids is that they are closed under element deletions and contractions (and hence matroid minors), cf. [84], Proposition 3.2.5. The following equivalent characterisation of regular matroids is useful for encoding purposes. A matrix with entries in \mathbb{R} is called *totally unimodular* if every square submatrix has determinant $-1, 0$ or 1 .

Theorem H.2.1 ([107]). *Let M be a matroid. Then M is regular if and only if $M \simeq M[A]$ for a totally unimodular real-valued matrix A . Furthermore, for any field \mathbb{F} , reinterpreting the $\{-1, 0, 1\}$ -entries of A as elements of \mathbb{F} , we obtain an \mathbb{F} -linear representation of M .*

Every graphic and every cographic matroid is regular, but not vice-versa. Regular matroids are in turn generalised by the *binary matroids*, which are the \mathbb{F}_2 -linear matroids.

Standard matroid notions used in our paper include *matroid minors* and the *deletion and contraction operations*, compare [84] for definitions. Throughout the paper, we use the following notation: Given a matroid M and an element $e \in E(M)$, we denote by $M \setminus e$ and M/e the matroids obtained from M by *deleting* and *contracting* e respectively.

These operations are consistent with deletions and contractions in graph theory in the following sense: If G is a graph and $e \in E(G)$, let us denote by G/e the graph obtained by contracting the edge e and by $G - e$ the graph obtained by deleting e . Then it holds that $M(G/e) \simeq M(G)/e$, $M(G - e) \simeq M(G) \setminus e$, $M^*(G - e) = M^*(G)/e$, and finally $M^*(G/e) \simeq M^*(G) \setminus e$.

Oriented Matroids

For missing terminology and basic facts from the theory of oriented matroids not mentioned or mentioned without proof in the following, please consult the standard reading [11].

An *oriented matroid* \vec{M} is a tuple (E, \mathcal{C}) consisting of a ground set E of elements and a collection \mathcal{C} of *signed subsets* of E , i.e. ordered partitions (C^+, C^-) of subsets C of E into *positive* and *negative* parts such that the following axioms are satisfied:

- $(\emptyset, \emptyset) \notin \mathcal{C}$
- If $(C^+, C^-) \in \mathcal{C}$, then $(C^-, C^+) \in \mathcal{C}$.
- If $(C_1^+, C_1^-), (C_2^+, C_2^-) \in \mathcal{C}$ such that $C_1^+ \cup C_1^- \subseteq C_2^+ \cup C_2^-$, then one of the equations $(C_1^+, C_1^-) = (C_2^+, C_2^-)$ or $(C_1^+, C_1^-) = (C_2^-, C_2^+)$ holds.
- Let $(C_1^+, C_1^-), (C_2^+, C_2^-) \in \mathcal{C}$ such that $(C_1^+, C_1^-) \neq (C_2^-, C_2^+)$, and let $e \in C_1^+ \cap C_2^-$. Then there exists some $(C^+, C^-) \in \mathcal{C}$ which satisfies $C^+ \subseteq (C_1^+ \cup C_2^+) \setminus \{e\}$ and $C^- \subseteq (C_1^- \cup C_2^-) \setminus \{e\}$.

In case these axioms are satisfied, the elements of \mathcal{C} are called *signed circuits*.

Two oriented matroids $\vec{M}_1 = (E_1, \mathcal{C}_1)$ and $\vec{M}_2 = (E_2, \mathcal{C}_2)$ are called *isomorphic* if there exists a bijection $\sigma : E_1 \rightarrow E_2$ such that $\{(\sigma(C^+), \sigma(C^-)) \mid (C^+, C^-) \in \mathcal{C}_1\} = \mathcal{C}_2$. For every oriented matroid $\vec{M} = (E, \mathcal{C})$ and a signed circuit $X = (C^+, C^-) \in \mathcal{C}$, we denote by $\underline{X} := C^+ \cup C^-$ the so-called *support* of X . From the axioms for signed circuits it follows that the set family $\underline{\mathcal{C}} := \{\underline{X} \mid X \in \mathcal{C}\}$ over the ground set E defines a matroid $M = (E, \underline{\mathcal{C}})$, which we refer to as the *underlying matroid* of \vec{M} , and vice versa, \vec{M} is called an *orientation* of M . A matroid is called *orientable* if it admits at least one orientation. A signed circuit (C^+, C^-) is called *directed* if either $C^+ = \emptyset$ or $C^- = \emptyset$. We use this definition also for the circuits of the underlying matroid M , i.e., a circuit of M is *directed* in \vec{M} if (C, \emptyset) (or equivalently

(\emptyset, C)) is a directed signed circuit of \vec{M} . We say that \vec{M} is *totally cyclic* if every element of M is contained in a directed circuit, and *acyclic* if there exists no directed circuit.

Classical examples of oriented matroids can be derived from vector configurations in real-valued vector spaces and, most importantly for the investigations in this paper, from directed graphs.

Given a configuration $x_1, x_2, \dots, x_k \in \mathbb{R}^n$ of vectors for some $k \in \mathbb{N}$, consider the matroid $M[A]$ with $A = (x_1, x_2, \dots, x_k) \in \mathbb{R}^{n \times k}$. Given a circuit $C = \{x_{i_1}, x_{i_2}, \dots, x_{i_\ell}\} \in \mathcal{C}$, then there are scalars $\alpha_1, \dots, \alpha_\ell \in \mathbb{R} \setminus \{0\}$ such that $\sum_{j=1}^{\ell} \alpha_j x_{i_j} = 0$, and the coefficients α_j are determined up to multiplication with a common scalar. It is therefore natural to assign two signed sets to the circuit as follows: $X(C) := (C^+, C^-)$ and $-X(C) := (C^-, C^+)$, where $C^+ := \{x_{i_j} \mid \alpha_{i_j} > 0\}$ and $C^- := \{x_{i_j} \mid \alpha_{i_j} < 0\}$. The *oriented matroid induced by A* is then defined as $\vec{M}[A] = (\{x_1, \dots, x_k\}, \{X(C), -X(C) \mid C \in \mathcal{C}_A\})$.

Given a digraph D we can, as in the undirected case, associate with it two different kinds of oriented matroids with ground set $E(D)$. Unsurprisingly, their underlying matroids are exactly the cycle matroid and the bond matroid of $U(D)$, respectively.

Definition H.2.2. Let D be a digraph.

- For every cycle C in D , let $(C^+, C^-), (C^-, C^+)$ be the two tuples describing a partition of $E(C)$ into sets of forward and backward edges, according to some choice of cyclical traversal of C . Then $\{(C^+, C^-), (C^-, C^+) \mid C \text{ cycle in } D\}$ forms the set of signed circuits of an orientation $M(D)$ of $M(U(D))$, called the *oriented cycle matroid induced by D*.
- For every bond $S = D[X, Y]$ in D , let S^+ be the set of edges in S with tail in X and head in Y , and let S^- contain those edges on S with tail in Y and head in X . Then $\{(S^+, S^-), (S^-, S^+) \mid S \text{ is a bond in } D\}$ forms the set of signed circuits of an orientation $M^*(D)$ of $M^*(U(D))$, called the *oriented bond matroid induced by D*.

Note that the directed circuits of an oriented cycle matroid are exactly the edge-sets of the directed cycles of the corresponding digraph D . Similarly, the directed circuits in an oriented bond matroid are the edge-sets of the directed

bonds in the corresponding digraph. An important fact to note at this point is that given a graphic (resp. cographic) matroid M , any orientation of M is necessarily isomorphic to the oriented cycle matroid (resp. oriented bond matroid) of a digraph D , we refer the reader to Corollary 6.2.8 in [12] for a proof of this fact.

Given an oriented matroid $\vec{M} = (E, \mathcal{C})$ and an element $e \in E$, we denote by $\vec{M} - e$ and \vec{M}/e the oriented matroids obtained from \vec{M} by *deleting* and *contracting* e , respectively. The signed circuits of these matroids are defined as follows:

$$\mathcal{C}(\vec{M} \setminus e) := \{(C^+, C^-) \in \mathcal{C} \mid e \notin C^+ \cup C^-\},$$

and the signed circuits of \vec{M}/e are the (inclusion-wise) support-minimal members of

$$\{(C^+ \setminus \{e\}, C^- \setminus \{e\}) \mid (C^+, C^-) \in \mathcal{C} \setminus \{(\emptyset, \emptyset)\}.$$

These definitions generalize to subsets $Z \subseteq E$, here we denote by $\vec{M} \setminus Z$ resp. \vec{M}/Z the oriented matroids obtained from \vec{M} by successively deleting (resp. contracting) all elements of Z (in arbitrary order[‡]).

Again, in the case of graphic and oriented cographic matroids, the deletion and contraction operations resemble the same operations in directed graphs: Given a digraph D and $e \in E(D)$, denote by D/e the digraph obtained by deleting e and identifying the endpoints of e . We then have $M(D) \setminus e \simeq M(D - e)$ and $M(D)/e \simeq M(D/e)$, whereas $M^*(D) \setminus e \simeq M^*(D/e)$ and $M^*(D)/e \simeq M^*(D - e)$.

For an oriented matroid \vec{M} with a collection \mathcal{C} of signed circuits, let $\hat{\mathcal{S}}$ be defined as the set of signed vectors (S^+, S^-) satisfying the following *orthogonality property* for every signed circuit $C = (C^+, C^-) \in \mathcal{C}$:

$$(S^+ \cap C^+) \cup (S^- \cap C^-) \neq \emptyset \iff (S^+ \cap C^-) \cup (S^- \cap C^+) \neq \emptyset. \quad (*)$$

Let \mathcal{S} denote the set of signed vectors of $\hat{\mathcal{S}} \setminus \{(\emptyset, \emptyset)\}$ with inclusion-wise minimal support. Then the members of \mathcal{S} are called the *signed cocircuits* of \vec{M} , compare the discussion after Theorem 2.2 in [12] for more background on the orthogonality property of oriented matroids.

The supports of the signed cocircuits form exactly the cocircuits of the underlying matroid M . A signed cocircuit (S^+, S^-) is called *directed* if $S^+ = \emptyset$ or $S^- = \emptyset$. If

[‡]It is well known that the order in which elements are deleted resp. contracted does not affect the outcome of the process.

the underlying matroid M of \vec{M} is regular, then the following stronger orthogonality holds for every signed circuit $(C^+, C^-) \in \mathcal{C}$, and every signed cocircuit $(S^+, S^-) \in \mathcal{S}$:

$$|C^+ \cap S^+| + |C^- \cap S^-| = |C^+ \cap S^-| + |C^- \cap S^+|. \quad (**)$$

A nice explanation of the strong orthogonality property and further background on orientations of regular matroids can be found in the paper [82] by Minty, as well as in chapter 6 from [12]. For any digraph D the signed cocircuits of $M(D)$ are the same as the signed circuits of $M^*(D)$, while the signed cocircuits of $M^*(D)$ are exactly the signed circuits of $M(D)$.

We conclude this first part of the preliminary section by stating a couple of important facts concerning orientations of (regular) matroids from the literature.

Theorem H.2.3 ([11]). *Let \vec{M} be an orientation of a regular matroid M . Then there exists a totally unimodular matrix A such that $\vec{M} \simeq \vec{M}[A]$ and $M \simeq M[A]$.*

We will also need the following matroidal version of the famous Farkas' Lemma:

Theorem H.2.4 ([11]). *Let \vec{M} be an oriented matroid and $e \in E(\vec{M})$. Then e is contained in a directed circuit of \vec{M} if and only if it is not contained in a directed cocircuit.*

H.2.1. Non-Evenness and GB-minors

Our main result, Theorem H.1.9, builds on the important fact that the non-even oriented matroids are closed under the GB-minor relation. In this subsection we present a proof of this fact and use it to derive Proposition H.1.8 from Theorem H.1.3.

Lemma H.2.5. *Every GB-minor of a non-even oriented matroid is non-even.*

Proof. It suffices to show the following two statements: For every non-even oriented matroid \vec{M} and every element $e \in E(\vec{M})$, the oriented matroid $\vec{M} \setminus e$ is non-even as well, and for every element $e \in E(\vec{M})$ which is butterfly-contractible, the oriented matroid \vec{M}/e is non-even as well. The claim then follows by repeatedly applying these two statements. Let us now fix a set $J \subseteq E(\vec{M})$ of elements intersecting every directed circuit in \vec{M} an odd number of times.

For the first claim, note that since the underlying matroid M of \vec{M} is regular, so is the underlying matroid of $\vec{M} \setminus e$. Then clearly the set $J \setminus \{e\}$ intersects every directed circuit in $\vec{M} \setminus e$ an odd number of times, proving that $\vec{M} \setminus e$ is non-even.

For the second claim, let $e \in E(\vec{M})$ be butterfly-contractible. Let S be a cocircuit of M such that $(S \setminus \{e\}, \{e\})$ forms a signed cocircuit of \vec{M} . Then the underlying matroid of \vec{M}/e is a matroid minor of the regular matroid M and is hence regular. Define $J' \subseteq E(\vec{M}) \setminus \{e\}$ via

$$J' := \begin{cases} J & \text{if } e \notin J \\ J \Delta S & \text{if } e \in J. \end{cases}$$

We claim that for every directed circuit C in \vec{M}/e , the intersection $C \cap J'$ is odd. Indeed, by definition either C is a directed circuit also in \vec{M} not containing e , or $C \cup \{e\}$ is a directed circuit in \vec{M} , or $(C, \{e\})$ is a signed circuit of \vec{M} . The last case however is impossible, as it would form a contradiction to the fact that e is a butterfly-contractible element of \vec{M} .

In the first case, since $e \notin C$, we must have $S \cap C = \emptyset$ as otherwise again C and the signed cocircuit $(S \setminus \{e\}, \{e\})$ form a contradiction to the orthogonality property (*). This then shows that indeed $|C \cap J'| = |C \cap (J' \setminus S)| = |C \cap (J \setminus S)| = |C \cap J|$ is odd, as required.

In the second case, the orthogonality property (**) of regular oriented matroids applied with the directed circuit $C \cup \{e\}$ and the signed cocircuit $(\{e\}, S \setminus \{e\})$ within \vec{M} yield that the equation $|(C \cup \{e\}) \cap (S \setminus \{e\})| = |(C \cup \{e\}) \cap \{e\}| = 1$ holds. So let $C \cap S = \{f\}$ for some element $f \in E(\vec{M}) \setminus \{e\}$. By definition of J' , if $e \notin J$, then $|C \cap J'| = |C \cap J| = |(C \cup \{e\}) \cap J|$, which is odd. If $e \in J$, then we have (modulo 2)

$$|C \cap J'| = |C \cap (J \Delta S)| = |(C \cap J) \Delta (C \cap S)| \equiv |C \cap J| + |\{f\}| = |(C \cup \{e\}) \cap J|,$$

which is odd. Hence, we have shown that $|C \cap J'|$ is odd in every case, which yields that \vec{M}/e is a non-even oriented matroid. This concludes the proof. \square

Lemma H.2.5 allows us to immediately prove the correctness of Proposition H.1.8.

Proof of Proposition H.1.8. We prove both directions of the equivalence. Suppose first that \vec{M} is non-even. Then by Lemma H.2.5 every oriented matroid isomorphic to a GB-minor of \vec{M} is non-even as well. Hence it suffices to observe that none of

the matroids $M(\overset{\leftrightarrow}{C}_k)$ for odd $k \geq 3$ is non-even. However, this follows directly since any element set J in $M(\overset{\leftrightarrow}{C}_k)$ intersecting every directed circuit an odd number of times corresponds to an edge set in $\overset{\leftrightarrow}{C}_k$ intersecting every directed cycle an odd number of times, which cannot exist since by Theorem H.1.3 none of the digraphs $\overset{\leftrightarrow}{C}_k$ is non-even for an odd $k \geq 3$.

Vice versa, suppose that no GB-minor of \vec{M} is isomorphic to $M(\overset{\leftrightarrow}{C}_k)$ for any odd $k \geq 3$. Let D be a digraph such that $\vec{M} \simeq M(D)$. We claim that D must be non-even. Suppose not, then by Theorem H.1.3 D admits a butterfly minor isomorphic to $\overset{\leftrightarrow}{C}_k$ for some odd $k \geq 3$. We now claim that $M(D)$ has a GB-minor isomorphic to $M(\overset{\leftrightarrow}{C}_k)$. For this, it evidently suffices to verify the following general statement:

If an edge e of a digraph F is butterfly-contractible in F , then within $M(F)$ the corresponding element e of $M(F)$ is butterfly-contractible.

Indeed, let $e = (u, v)$ for distinct vertices $u, v \in V(D)$. Then by definition either u has out-degree 1 or v has in-degree 1 in D . In the first case, e is the unique edge leaving u in the cut $D[\{u\}, V(D) \setminus \{u\}]$, while in the second case e is the only edge entering v in the cut $D[V(D) \setminus \{v\}, \{v\}]$. Since every cut is an edge-disjoint union of bonds, we can find in both cases a bond containing e where e is the only edge directed away resp. towards the side of the bond that contains u resp. v .

Since the oriented bonds in D yield the signed cocircuits of $M(D)$, this shows that there is a cocircuit S in $M(D)$ such that $(S \setminus \{e\}, \{e\})$ is a signed cocircuit. Hence, e is a butterfly-contractible element of $M(D)$. This shows that $M(\overset{\leftrightarrow}{C}_k)$ is isomorphic to a GB-minor of $M(D) \simeq \vec{M}$ which contradicts our initial assumption that no GB-minor of \vec{M} is isomorphic to $M(\overset{\leftrightarrow}{C}_k)$. Hence, D is non-even, and there exists $J \subseteq E(D)$ such that every directed cycle in D contains an odd number of edges from J . The same set J also certifies that $\vec{M} \simeq M(D)$ is non-even, and this concludes the proof of the equivalence. \square

H.3. On the Complexity of the Even Directed Circuit Problem

The formulation of Problem H.1.4 is rather vague, as it is not clear by which means the oriented matroid \vec{M} is given as an input to an algorithm designed for

solving the problem, and in which way we will measure its efficiency. For the latter, it is natural to aim for an algorithm which performs a polynomial number of elementary steps in terms of the number of elements of \vec{M} . This also resembles the even dicycle problem in digraphs, where we aim to find an algorithm running in polynomial time in $|E(D)|$.

For the former, it is not immediately clear how to encode the (oriented) matroid, and hence how to make information contained in the (oriented) matroid available to the algorithm. For instance, if the list of all circuits of a matroid is given as input to an algorithm, one can decide in linear time whether there exists an even (directed) circuit. This list, however, will usually have exponential size in the number of elements, and therefore disqualify as a good reference value for efficiency of the algorithm. For that reason, different computational models (and efficiency measures) for algorithmic problems in matroids (see [51]) and oriented matroids (see [7]) have been proposed in the literature. These models are based on the concept of *oracles*. For a family $\mathcal{F} \subseteq 2^{E(M)}$ of objects characterising the matroid M , an *oracle* is a function $f : 2^{E(M)} \rightarrow \{\mathbf{true}, \mathbf{false}\}$ assigning to every subset a truth value indicating whether or not the set is contained in \mathcal{F} . If \mathcal{F} for instance corresponds to the collection of circuits, cocircuits, independent sets, or bases of a matroid, we speak of a *circuit-*, *cocircuit-*, *independence-*, or *basis-oracle*. Similarly, for oriented matroids we can define several oracles [7]. Maybe the most natural choice for an oriented matroid-oracle for Problem H.1.4 is the *circuit oracle*, which given any subset of the element set together with a $\{+, -\}$ -signing of its elements, reveals whether or not this signed subset forms a signed circuit of the oriented matroid. This computational model applied to Problem H.1.4 yields the following question.

Question H.3.1. Does there exist an algorithm which, given an oriented matroid \vec{M} , decides whether there exists a directed circuit in \vec{M} of even size, by calling the circuit-oracle of \vec{M} only $\mathcal{O}(|E(\vec{M})|^c)$ times for some $c \in \mathbb{N}$?

However, as it turns out, the answer to the above problem is easily seen to be negative, even when the input oriented matroid \vec{M} is graphic.

Proposition H.3.2. *Any algorithm deciding whether a given oriented graphic matroid on n elements, for some $n \in \mathbb{N}$, contains an even directed circuit must use at least $2^{n-1} - 1$ calls to the circuit-oracle for some instances.*

Proof. Suppose towards a contradiction there was an algorithm which decides whether a given oriented graphic matroid contains an even directed circuit and uses at most $2^{n-1} - 2$ oracle calls for any input oriented graphic matroid on elements $E := \{1, \dots, n\}$. Now, playing the role of the oracle, we will answer all of the (at most $2^{n-1} - 2$) calls of the algorithm by **false**. Since there are exactly $2^{n-1} - 1$ non-empty sets $Y \in 2^E$ of even size, there must be an even non-empty subset Y of E such that the algorithm did not call the oracle with any input signed set whose support is Y . But this means the algorithm cannot distinguish between the oriented graphic matroids (E, \mathcal{C}_0) and (E, \mathcal{C}_Y) , where $\mathcal{C}_0 := \emptyset$ and $\mathcal{C}_Y := \{(Y, \emptyset), (\emptyset, Y)\}$, which result in the same oracle-answers to the calls by the algorithm, while (E, \mathcal{C}_0) contains no even directed circuit, but (E, \mathcal{C}_Y) does. This shows that the algorithm does not work correctly, and this contradiction proves the assertion. \square

The above result and its proof give a hint that maybe in general the use of oriented matroid-oracles to measure the efficiency of algorithms solving Problem H.1.4 is doomed to fail. One should therefore look for a different encoding of the input oriented matroids in order to obtain a sensible algorithmic problem. In this paper, we solve this issue by restricting the class of possible input oriented matroids to *oriented regular matroids*, which allow for a much simpler and compact encoding via their representation by totally unimodular matrices (cf. Theorems H.2.1 and H.2.3). The following finally is the actual algorithmic problem we are going to discuss in this paper.

Problem H.3.3. Is there an algorithm which decides, given as input a totally unimodular matrix $A \in \mathbb{R}^{m \times n}$ for some $m, n \in \mathbb{N}$, whether $\vec{M}[A]$ contains an even directed circuit, and runs in time polynomial in mn ?

The alert reader might be wondering what happens if in the above problem we aim to detect odd instead of even directed circuits. The reason why this problem is not a center of study in our paper is that it admits a simple polynomial time solution, which is given in the form of Proposition H.3.15 at the end of this section.

The next statement translates the main results from [90] and [80] to our setting to show that Problem H.3.3 has a positive answer if we restrict to oriented graphic matroids as inputs.

Lemma H.3.4. *There exists an algorithm which, given as input any totally unimodular matrix $A \in \mathbb{R}^{m \times n}$ for some $m, n \in \mathbb{N}$ such that $\vec{M}[A]$ is a oriented graphic matroid, decides whether $\vec{M}[A]$ contains a directed circuit of even size, and which runs in time polynomial in mn .*

Proof. The main results of Robertson et al. [90] and McCuaig [80] yield polynomial time algorithms which, given as input a digraph D (by its vertex- and edge-list) returns whether or not D contains an even directed cycle. Therefore, given a totally unimodular matrix $A \in \mathbb{R}^{m \times n}$ such that $\vec{M}[A]$ is graphic, if we can construct in time polynomial in mn a digraph D such that $\vec{M}[A] \simeq M(D)$, then we can decide whether $\vec{M}[A]$ contains a directed circuit of even size by testing whether D contains an even directed cycle using the algorithms from [80, 90]. Such a digraph can be found as follows:

First, we consider the *unoriented* matroid $M[A]$ defined by the matrix A , which is graphic. It follows from a result of Tutte [106] (see also more explicitly Bixby and Cunningham [10]) that there exists an algorithm which, given a binary matrix representation of size $m \times n$ of a graphic matroid M , computes a connected graph G with n edges such that $M(G) \simeq M$, and which runs in time polynomial in mn .

Since given the totally unimodular representation A of $M[A]$, we can derive a binary representation of the same size by simply changing -1 entries into 1 entries, we can apply one of the algorithms from [10, 106] to find a graph G satisfying $M(G) \simeq M[A]$ in polynomial time. Since $M(G) \simeq M[A]$, there must exist an orientation of $M(G)$ isomorphic to $\vec{M}[A]$, and this orientation in turn can be realized as $M(D)$ where D is an orientation of G^{**} . To find the desired orientation D of G in polynomial time, we first compute a decomposition of G into its blocks G_1, \dots, G_k (maximal connected subgraphs without cutvertices).

Next we (arbitrarily) select for every $i \in \{1, \dots, k\}$ a special ‘reference’-edge $e_i \in E(G_i)$. Note that two different orientations of G obtained from each other by reversing all edges in one block result in the same oriented matroid, as cycles in G are always entirely contained in one block. Hence for every $i \in \{1, \dots, k\}$ we can orient e_i arbitrarily and assume w.l.o.g. that this orientation coincides with the orientation in D . Note that every block of G which is not 2-connected must be a

**The fact that every orientation of $M(G)$ can be realised as $M(D)$ for an orientation D of G follows from a classical result by Bland and LasVergnas [12], who show that regular matroids (and particularly graphic ones) have a unique reorientation class.

K_2 forming a bridge in G . In this case, the only edge of the block is our chosen reference-edge and already correctly oriented. Now, for every $i \in \{1, \dots, k\}$ such that G_i is 2-connected and every edge $e \in E(G_i) \setminus \{e_i\}$ there is a cycle C in G_i containing both e_i and e . This cycle can be computed in polynomial time using a disjoint-paths algorithm between the endpoints of e and e_i . Now we consider the minimally linearly dependent set of columns in A corresponding to C , and compute the coefficients of a non-trivial linear combination resulting in 0. As we already know the orientation of $e_i \in E(C)$, this yields us the orientations of all edges on the cycle C in D and hence of the edge e . In this way, we can compute all orientations of edges in D in polynomial time in mn and find the digraph D such that $\vec{M}[A] \simeq M(D)$. As discussed above, this concludes the proof. \square

H.3.1. Proof of Theorem H.1.6

We prepare the proof by a set of useful definitions and lemmas dealing with circuit bases of regular matroids.

Definition H.3.5. Let M be a binary matroid. The *circuit space* of M is the \mathbb{F}_2 -linear vector space generated by the incidence vectors $\mathbf{1}_C \in \mathbb{F}_2^{E(M)}$ defined by $\mathbf{1}_C(e) := 1$ for $e \in C$ and $\mathbf{1}_C(e) := 0$ for $e \notin C$ and all circuits C of M . A *circuit basis* of M is a set of circuits of M whose incidence vectors form a basis of the circuit space. Equivalently, we can consider the circuit space as a \mathbb{F}_2 -linear subspace of the vector space whose elements are all the subsets of E and where the sum of two sets $X, Y \subseteq E(M)$ is defined as their symmetric difference $X \Delta Y$.

Definition H.3.6. Let \vec{M} be a regular oriented matroid and M be its underlying regular matroid. We call a circuit basis \mathcal{B} of M *directed* if all elements of \mathcal{B} are directed circuits of \vec{M} .

The next proposition is a well-known fact about the circuit space of a binary matroid.

Proposition H.3.7 (cf. Corollary 9.2.3, [84]). *Let M be a binary matroid. Then the dimension of the circuit space of M equals $|E(M)| - r(M)$.*

The following lemma is crucial for the proof of Theorem H.1.6 as well as for our work on digraphs in Section H.4. We will need the following matroid terminology:

Given a matroid M , a subset $A \subseteq E(M)$ is called *coindependent* if it is an independent set of the dual matroid M^* , or, formally, if there exists a basis B of M such that $A \cap B = \emptyset$, i.e., if and only if \overline{A} fully includes a basis (we also say that \overline{A} is *spanning* in this case).

Lemma H.3.8. *Let \vec{M} be an oriented regular matroid. If \vec{M} is totally cyclic, then the underlying matroid M admits a directed circuit basis. Furthermore, for every coindependent set A in M such that $\vec{M} \setminus A$ is totally cyclic, there exists a directed circuit basis of M such that every $a \in A$ is contained in exactly one circuit of the basis.*

Proof. We start by proving the first assertion concerning the existence of a directed circuit basis of M . We use induction on $|E(M)|$. If M consists of a single element, the claim holds trivially, since every circuit is a loop and thus directed. So assume now that $|E(M)| = k \geq 2$ and that the statement of the lemma holds for all oriented regular matroids on at most $k - 1$ elements. Choose some $e \in E(M)$ arbitrarily. Since \vec{M} is totally cyclic, there exists a directed circuit C_e containing e . Let us now consider the oriented regular matroid $\vec{M} \setminus e$. If $\vec{M} \setminus e$ is totally cyclic, then we can apply the induction hypothesis to $\vec{M} \setminus e$ and find a directed circuit basis \mathcal{B}^- of $M \setminus e$. Now consider the collection $\mathcal{B} = \mathcal{B}^- \cup \{C_e\}$ of directed circuits in \vec{M} . The incidence vectors of these circuits are linearly independent over \mathbb{F}_2 , as C_e is the only circuit yielding a non-zero entry at element e . Furthermore, we get by induction that $|\mathcal{B}| = |E(M)| - 1 - r(M \setminus e) + 1 = |E(M)| - r(M \setminus e) = |E(M)| - r(M)$. The last equality holds since e is contained in the circuit C_e and hence $\{e\}$ does not form a cocircuit.

As this matches the dimension of the circuit space of M , we have found a directed circuit basis of M , proving the inductive claim in this case.

It remains to prove the case where $\vec{M} \setminus e$ is not totally cyclic, i.e., there is an element not contained in a directed circuit. By Farkas' Lemma (Theorem H.2.4) applied to $\vec{M} \setminus e$ and this element there exists a directed cocircuit S in $\vec{M} \setminus e$. Then either (S, \emptyset) , $(S \cup \{e\}, \emptyset)$ or $(S, \{e\})$ form a signed cocircuit of \vec{M} . Since \vec{M} is totally cyclic, it contains no directed cocircuits, and hence only the latter case is possible, $(S, \{e\})$ must form a signed cocircuit and thus e is butterfly contractible.

Let us now consider the oriented regular matroid \vec{M}/e . Since \vec{M} is totally cyclic, so is \vec{M}/e . By the induction hypothesis there exists a directed circuit basis \mathcal{B}^-

of M/e . By definition, for every directed circuit $C \in \mathcal{B}^-$, either C is a directed circuit in \vec{M} not containing e , or $C \cup \{e\}$ is a directed circuit in \vec{M} , or $(C, \{e\})$ forms a signed circuit of \vec{M} . The latter is, however, impossible, as it would form a contradiction to the fact that e is a butterfly-contractible element of \vec{M} .

Hence, the set $\mathcal{B} := \{C \mid C \in \mathcal{B}^- \text{ circuit in } M\} \cup \{C \cup \{e\} \mid C \in \mathcal{B}^-, C \cup \{e\} \text{ circuit in } M\}$ consists of $|\mathcal{B}| = |\mathcal{B}^-| = |E(M)| - 1 - r(M/e) = |E(M)| - r(M)$ many circuits of M which are all directed ones in \vec{M} . Note that for the last equality we used that e is not a loop, as it is contained in the cocircuit $S \cup \{e\}$ of M . Finally, we claim that the binary incidence vectors of the elements of \mathcal{B} in $\mathbb{F}_2^{E(M)}$ are linearly independent. This follows since the restriction of these vectors to the coordinates $E(M) \setminus \{e\}$ equals the characteristic vectors of the elements of \mathcal{B}^- , which form a circuit basis of M/e . This shows that we have found a directed circuit basis of M , proving the inductive claim.

For the second assertion, let a coindependent set A in M be given and suppose that $\vec{M} \setminus A$ is totally cyclic. We claim that for every $a \in A$ there exists a directed circuit C_a in \vec{M} such that $C_a \cap A = \{a\}$. Equivalently, we may show that the oriented matroid $\vec{M} \setminus (A \setminus \{a\})$ has a directed circuit containing a . Towards a contradiction, suppose not, then by Farkas' Lemma (Theorem H.2.4) there exists a directed cocircuit S in $\vec{M} \setminus (A \setminus \{a\})$ containing a . Since A is coindependent, $\{a\}$ is not a cocircuit of $M \setminus (A \setminus \{a\})$ and hence $S \setminus \{a\} \neq \emptyset$. Every directed circuit in $\vec{M} \setminus (A \setminus \{a\})$ must be disjoint from S , and hence no $f \in S \setminus \{a\}$ is contained in a directed circuit of $\vec{M} \setminus A$, contradicting our assumption that $\vec{M} \setminus A$ is totally cyclic. It follows that for each $a \in A$ a directed circuit C_a with $C_a \cap A = \{a\}$ exists.

Next we apply the first assertion of this lemma to the totally cyclic oriented matroid $\vec{M} \setminus A$. We get that there is a directed circuit basis \mathcal{B}_A of $M \setminus A$. We claim that $\mathcal{B} := \mathcal{B}_A \cup \{C_a \mid a \in A\}$ forms a directed circuit basis of M satisfying the properties claimed in this lemma. Indeed, every circuit in \mathcal{B} is a directed circuit of \vec{M} , and for every $a \in A$ the circuit C_a is the only circuit in \mathcal{B} containing a . Since the characteristic vectors of the elements of \mathcal{B}_A are linearly independent as \mathcal{B}_A is a circuit basis of $M \setminus A$, we already get that the characteristic vectors of elements of \mathcal{B} are linearly independent using that the characteristic vector of C_a is the only basis-vector having a non-zero entry at the position corresponding to element a . To show that \mathcal{B} indeed is a circuit basis of M , it remains to verify that it has the required size. We have $|\mathcal{B}| = |A| + |\mathcal{B}_A| = |A| + |E(M \setminus A)| - r(M \setminus A) = |E(M)| - r(M)$, where

for the latter equality we used that $r(M \setminus A) = r(M)$ since A is coindependent. This concludes the proof of the second assertion. \square

In order to prove our next lemma, we need the following result, which was already used by Seymour and Thomassen.

Lemma H.3.9 ([96], Prop. 3.2). *Let E be a finite set and \mathcal{F} a family of subsets of E . Then precisely one of the following statements holds:*

- (i) *There is a subset $J \subseteq E$ such that $|F \cap J|$ is odd for every $F \in \mathcal{F}$.*
- (ii) *There are sets $F_1, \dots, F_k \in \mathcal{F}$, where $k \in \mathbb{N}$ is odd, such that $\Delta_{i=1}^k F_i = \emptyset$.*

Please note that (i) and (ii) cannot hold simultaneously because if k is odd and F_1, \dots, F_k all have odd intersection with J , then the symmetric difference $\Delta_{i=1}^k F_i$ has odd intersection with J .

We now derive the following corollary for totally cyclic oriented regular matroids by using Lemma H.3.8 and applying Lemma H.3.9 to a directed circuit basis.

Corollary H.3.10. *Let \vec{M} be a totally cyclic oriented regular matroid, and let \mathcal{B} be a directed circuit basis of M . Then there exists $J \subseteq E(\vec{M})$ such that $|C \cap J|$ is odd for every $C \in \mathcal{B}$.*

Proof. The claim is that (i) in Lemma H.3.9 with $E = E(\vec{M})$ and $\mathcal{F} := \mathcal{B}$ holds true, so it suffices to rule out (ii). However, the latter would contradict the linear independence of the basis \mathcal{B} . \square

Building on this corollary we derive equivalent properties for an oriented matroid to be non-even.

Proposition H.3.11. *Let \vec{M} be a totally cyclic oriented regular matroid and let \mathcal{B} be a directed circuit basis of M . Furthermore, let $J \subseteq E(M)$ be such that $|C \cap J|$ is odd for all $C \in \mathcal{B}$. Then the following statements are equivalent:*

- (i) *\vec{M} is non-even.*
- (ii) *If C_1, \dots, C_k are directed circuits of \vec{M} where $k \in \mathbb{N}$ is odd, then $\Delta_{i=1}^k C_i \neq \emptyset$.*
- (iii) *Every directed circuit of \vec{M} is the symmetric difference of an odd number of elements of \mathcal{B} .*

(iv) $|C \cap J|$ is odd for all directed circuits C of \vec{M} .

Proof.

“(i) \Rightarrow (ii)” This follows from Lemma H.3.9 applied to the set of all directed circuits of \vec{M} .

“(ii) \Rightarrow (iii)” Let C be a directed circuit of \vec{M} . Since \mathcal{B} is a circuit basis of M , we can write $C = \Delta_{i=1}^k C_i$ for some $k \in \mathbb{N}$ and $C_1, \dots, C_k \in \mathcal{B}$. If k were even, then the sum $C + \Delta_{i=1}^k C_i = \emptyset$ would yield a contradiction to (ii).

“(iii) \Rightarrow (iv)” Let C be a directed circuit of \vec{M} . By assumption, $C = \Delta_{i=1}^k C_i$ with $k \in \mathbb{N}$ being odd and $C_1, \dots, C_k \in \mathcal{B}$. Since J has odd intersection with all C_i , the set J has also odd intersection with C .

“(iv) \Rightarrow (i)” This implication follows directly from the definition of non-even.

□

Before we turn towards the proof of Theorem H.1.6 we need the following result, yielding a computational version of Theorem H.2.4 (Farkas’ lemma) for oriented regular matroids. Although we suspect the statement is well-known among experts, we include a proof for the sake of completeness.

Lemma H.3.12. *There exists an algorithm that, given as input a totally unimodular matrix $A \in \{-1, 0, 1\}^{m \times n}$ representing a regular oriented matroid $\vec{M} \simeq \vec{M}[A]$, and an element $e \in E(\vec{M})$, outputs either a directed circuit of \vec{M} containing e or a directed cocircuit of \vec{M} containing e , and which runs in polynomial time in mn .*

Proof. We first observe that we can decide in polynomial time in mn whether e is contained in a directed circuit or in a directed cocircuit of \vec{M} (by Farkas’ Lemma, we know that exactly one of these two options must be satisfied). Let us denote for every element $f \in E(\vec{M})$ by $x_f \in \{-1, 0, 1\}^m$ the corresponding column-vector of A . We need the following claim:

The element e is contained in a directed circuit of \vec{M} if and only if there exist non-negative scalars $\alpha_f \geq 0$ for $f \in E(\vec{M}) \setminus \{e\}$ such that $-x_e = \sum_{f \in E(\vec{M}) \setminus \{e\}} \alpha_f x_f$.

The necessity of this condition follows directly by definition of $\vec{M}[A]$: If e is contained in a directed circuit with elements e, f_1, \dots, f_k , then there are coefficients $\beta_e > 0$ and $\beta_i > 0$ for $1 \leq i \leq k$ such that $\beta_e x_e + \sum_{i=1}^k \beta_i x_{f_i} = 0$, i.e., $-x_e =$

$\sum_{i=1}^k \frac{\beta_i}{\beta_e} x_{f_i}$. On the other hand, if $-x_e$ is contained in the conic hull of $\{x_f | f \in E(\vec{M}) \setminus \{e\}\}$, then we can select an inclusion-wise minimal subset $F \subseteq E(\vec{M}) \setminus \{e\}$ such that $-x_e$ is contained in the conic hull of $\{x_f | f \in F\}$. We claim that $\{e\} \cup F$ forms a directed circuit of \vec{M} . By definition of F , it suffices to verify that the vectors x_e and x_f for $f \in F$ are minimally linearly dependent. However, this follows directly by Carathéodory's Theorem: The dimension of the subspace spanned by $\{x_f | f \in F\}$ equals $|F|$, for otherwise we could select a subset of at most $|F| - 1$ elements from $\{x_f | f \in F\}$ whose conic hull also contains $-x_e$, contradicting the minimality of F . This shows the equivalence claimed above.

We can now use a well-known linear programming algorithm for linear programs with integral constraints, compare [35,37,45,101] to decide in (strongly) polynomial time^{††} (and hence in polynomial time in mn) the feasibility of the linear inequality system

$$\sum_{f \in E(\vec{M} \setminus \{e\})} \alpha_f x_f = -x_e, \text{ with } \alpha_f \geq 0.$$

Therefore, we have shown that we can decide in polynomial time in mn whether or not e is contained in a directed circuit of \vec{M} . Next we give an algorithm which, given that e is contained in a directed circuit of \vec{M} , finds such a circuit in polynomial time:

During the procedure, we update a subset $Z \subseteq E(\vec{M})$, which maintains the property that it contains a directed circuit including e . At the end of the procedure Z will form such a directed circuit of \vec{M} . We initialise $Z := E(\vec{M})$. During each step of the procedure, we go through the elements $f \in Z \setminus \{e\}$ one by one and apply the above algorithm to test whether $\vec{M} \setminus ((E(\vec{M}) \setminus Z) \cup \{f\})$ contains a directed circuit including e . At the first moment such an element is found, we put $Z := Z \setminus \{f\}$ and repeat. If no such element is found, we stop and output Z .

Since we reduce the size of the set Z at each round of the procedure, the above algorithm runs in at most n rounds and calls the above decision algorithm for the existence of a directed circuit including e at most $n - 1$ times in every round. All in all, the algorithm runs in time polynomial in mn . It is obvious that the procedure maintains the property that Z contains a directed circuit including e and that at the end of the procedure all elements of Z must be contained in this circuit, i.e., Z forms a directed circuit with the desired properties.

^{††}Here we use the fact that all coefficients appearing in the linear system are $-1, 0$ or 1 .

To complete the proof we now give an algorithm which finds either a directed circuit or a directed cocircuit through a given element e of \vec{M} as follows: First we apply the first (decision) algorithm, which either tells us that e is contained in a directed circuit of \vec{M} , in which case we apply the second (detection) algorithm to find such a circuit. Otherwise we know that e is contained in a directed cocircuit of \vec{M} , in which case we compute in polynomial time a totally unimodular representing matrix A^* with at most n rows and n columns^{‡‡} of the dual regular oriented matroid \vec{M}^* . As we know that e is included in a directed circuit of \vec{M}^* , we can apply the second (detection) algorithm to A^* and \vec{M}^* instead of A and \vec{M} in order to find a directed cocircuit in \vec{M} containing e in polynomial time. \square

Given a regular oriented matroid \vec{M} we shall denote by $TC(\vec{M})$ the largest totally cyclic deletion minor of \vec{M} , i.e. the deletion minor of \vec{M} whose ground set is

$$E(TC(\vec{M})) := \bigcup \{C \mid C \text{ is a directed circuit of } \vec{M}\}.$$

From Lemma H.3.12 we directly have the following.

Corollary H.3.13. *There exists an algorithm that, given as input a totally unimodular matrix $A \in \{-1, 0, 1\}^{m \times n}$ representing a regular oriented matroid \vec{M} , where $m \in \mathbb{N}$ and $n = |E(\vec{M})|$, computes a submatrix $B \in \{-1, 0, 1\}^{m \times n'}$ of A representing $TC(\vec{M})$, where $n' = |E(TC(\vec{M}))|$, in time polynomial in mn . \square*

The last ingredient we shall need for the proof of Theorem H.1.6 is a computational version of the first statement of Lemma H.3.8 combined with Corollary H.3.10.

Lemma H.3.14. *Let \vec{M} be a totally cyclic regular oriented matroid represented by a totally unimodular matrix $A \in \{-1, 0, 1\}^{m \times n}$ for some $m \in \mathbb{N}$ and $n = |E(\vec{M})|$. Then we can compute a directed circuit basis \mathcal{B} of \vec{M} together with a set $J \subseteq E(\vec{M})$ such that $|J \cap B| \equiv 1 \pmod{2}$ for every $B \in \mathcal{B}$ in time polynomial in mn .*

^{‡‡}To find such a representing matrix, one can use Gaussian elimination to compute a basis \mathcal{B} of $\ker(A)$. Since A is totally unimodular, the vectors in \mathcal{B} can be taken to be $\{-1, 0, 1\}$ -vectors such that the matrix A^* consisting of the elements of \mathcal{B} written as row-vectors is totally unimodular as well. It then follows from the orthogonality property of regular oriented matroids that A^* indeed forms a representation of \vec{M}^* , using the fact that the row spaces of A and A^* are orthogonal complements.

Proof. We shall follow the inductive proof of Lemma H.3.8 to obtain a recursive algorithm for finding a desired directed circuit basis together with the desired set J . If $n = 1$, the unique element e of $E(\vec{M})$ is a directed loop, since \vec{M} is totally cyclic, and forms our desired directed circuit basis of \vec{M} . Furthermore, by setting $J := \{e\}$ we also get our desired set.

In the case $n \geq 2$, let us fix an arbitrary element e of $E(\vec{M})$ and compute a directed circuit C_e of \vec{M} containing e by applying Lemma H.3.12. Also using Lemma H.3.12, we can test in time polynomial in mn whether $\vec{M} \setminus e$ is totally cyclic. If so, we fix C_e as an element of our desired directed circuit base \mathcal{B} of \vec{M} and proceed as before with $\vec{M} \setminus e$ instead of \vec{M} . The set J is updated as follows: Suppose we have already computed a directed circuit base \mathcal{B}^- and a set J^- as in the statement of this lemma, but with respect to $\vec{M} \setminus e$. Then we set $\mathcal{B} := \mathcal{B}^- \cup \{C_e\}$. Now we check the parity of $|J^- \cap C_e|$ and set

$$J := \begin{cases} J^- & \text{if } |J^- \cap C_e| \equiv 1 \pmod{2} \\ J^- \cup \{e\} & \text{if } |J^- \cap C_e| \equiv 0 \pmod{2}. \end{cases}$$

As C_e is the only element of \mathcal{B} that contains e , the set J has odd intersection with every element of \mathcal{B} , as desired.

If $\vec{M} \setminus e$ is not totally cyclic, we compute a totally unimodular representative matrix $A' \in \{-1, 0, 1\}^{m \times (n-1)}$ of \vec{M}/e . This task can be executed in time polynomial in mn ^{§§}. Now \vec{M}/e is totally cyclic as \vec{M} is totally cyclic and we proceed as before with \vec{M}/e instead of \vec{M} . However, when our recursive algorithm already yields a directed circuit basis \mathcal{B}^- of \vec{M}/e as well as a set J^- for \vec{M}/e as in the statement of this lemma, we know as argued in the proof of Lemma H.3.8 that each element C of \mathcal{B}^- either is a directed circuit of \vec{M} or $C \cup \{e\}$ is a directed circuit of \vec{M} . Depending on this distinction we define our desired circuit basis \mathcal{B} of \vec{M} as in the proof of Lemma H.3.8 via

$$\mathcal{B} := \{C \mid C \in \mathcal{B}^- \text{ circuit in } M\} \cup \{C \cup \{e\} \mid C \in \mathcal{B}^-, C \cup \{e\} \text{ circuit in } M\}.$$

^{§§}To compute A' , select a non-zero entry in the column of A belonging to the element e . Pivoting on this element and exchanging rows transforms A in polynomial time in mn into a totally unimodular matrix $A'' \in \{-1, 0, 1\}^{m \times n}$ of \vec{M} in which the column corresponding to the element e of \vec{M} is $(1, 0, \dots, 0)^\top$. Then $\vec{M}[A] = \vec{M}[A'']$, and the matrix A' obtained from A'' by deleting the first row is a totally unimodular representation of \vec{M}/e .

To decide for each element $C \in \mathcal{B}^-$ whether C or $C \cup \{e\}$ is a directed circuit of \vec{M} we calculate $A\mathbf{1}_C$ where $\mathbf{1}_C$ denotes the incidence vector of C with respect to A . Then C forms a directed circuit of \vec{M} if and only if $A\mathbf{1}_C = 0$. As $|\mathcal{B}^-| = |\mathcal{B}| = |E(\vec{M})| - r(\vec{M})$ as argued in the proof of Lemma H.3.8 and by Proposition H.3.7, we have to do at most n of these computations to compute \mathcal{B} from \mathcal{B}^- . Regarding the set J we can simply set $J := J^-$. \square

We are now ready for the proof of Theorem H.1.6.

Proof of Theorem H.1.6. Assume first we have access to an oracle deciding whether an oriented regular matroid given by a representing totally unimodular matrix is non-even. Now suppose we are given a regular oriented matroid \vec{M} represented by a totally unimodular matrix $A \in \{-1, 0, 1\}^{m \times n}$ for some $m, n \in \mathbb{N}$ and we want to decide whether it contains a directed circuit of even size.

First we compute $TC(\vec{M})$, which can be done in time polynomial in mn by Corollary H.3.13. Now we use Lemma H.3.14 to compute a directed circuit basis of $TC(\vec{M})$ in time polynomial in mn . Then we go through the $|E(TC(\vec{M}))| - r(TC(\vec{M}))$ many elements of the basis and check whether one of these directed circuits has even size. If so, the algorithm terminates. Otherwise, every member of the basis has odd size. By Proposition H.3.11 with $J := E(TC(\vec{M}))$, we know that $TC(\vec{M})$ contains no directed circuit of even size if and only if $TC(\vec{M})$ is non-even. Since $TC(\vec{M})$ is the largest deletion minor of \vec{M} , which has the same directed circuits as \vec{M} , we know that $TC(\vec{M})$ is non-even if and only if \vec{M} is non-even. So we can decide the question using the oracle.

Conversely, assume we have access to an oracle which decides whether a given oriented regular matroid contains a directed circuit of even size. Again, our first step is to compute $TC(\vec{M})$ using Corollary H.3.13. By Lemma H.3.14 we then compute a directed circuit basis of $TC(\vec{M})$ and a set $J \subseteq E(TC(\vec{M}))$ such that every circuit in the basis has odd intersection with J .

Let \vec{M}' be the oriented matroid obtained from $TC(\vec{M})$ by duplicating every element $e \in E(TC(\vec{M})) \setminus J$ into two copies e_1 and e_2 that are in series[¶]. This way,

[¶]Concretely, this means that we transform every signed circuit of $TC(\vec{M})$ into a signed circuit of \vec{M}' by replacing every occurrence of an element $e \in E(TC(\vec{M})) \setminus J$ in a signed partition by the two elements e_1, e_2 in the same set of the signed partition. It is not hard to see that this indeed defines an oriented matroid, which is still regular.

every directed circuit in \vec{M}' intersects $E(\vec{M}') \setminus J$ in an even number of elements. Thus, for every directed circuit C in $TC(\vec{M})$, the size of the corresponding directed circuit in \vec{M}' is odd if and only if $|C \cap J|$ is odd. Hence, J intersects every directed circuit in $TC(\vec{M})$ an odd number of times if and only if \vec{M}' contains no even directed circuit. By Proposition H.3.11 this shows that $TC(\vec{M})$ is non-even if and only if \vec{M}' has no directed circuit of even size. Since $TC(\vec{M})$ is non-even if and only if \vec{M} is non-even, we can decide the non-evenness of \vec{M} by negating the output of the oracle with instance \vec{M}' . \square

With the tools developed in this section at hand we are ready for the proof of Proposition H.3.15.

Proposition H.3.15. *There is an algorithm which given as input a totally unimodular matrix $A \in \mathbb{R}^{m \times n}$ for some $m, n \in \mathbb{N}$, either returns an odd directed circuit of $\vec{M}[A]$ or concludes that no such circuit exists, and runs in time polynomial in mn .*

Proof. Let $A \in \mathbb{R}^{m \times n}$ be a totally unimodular matrix given as input and let $\vec{M} := \vec{M}[A]$. To decide whether \vec{M} contains a directed circuit of odd size, we first use Corollary H.3.13 to compute a totally unimodular representation of $TC(\vec{M})$ in polynomial time in mn . We now apply Lemma H.3.14 to compute in polynomial time a directed circuit basis \mathcal{B} of $TC(\vec{M})$. Going through the elements of \mathcal{B} one by one, we test whether one of the basis-circuits is odd, in which case the algorithm stops and returns this circuit. Otherwise, all circuits in \mathcal{B} are even. Since every circuit in the underlying matroid of $TC(\vec{M})$ can be written as a symmetric difference of elements of \mathcal{B} , every circuit in this matroid must be even. In particular, $TC(\vec{M})$ and hence \vec{M} do not contain any odd directed circuits, and the algorithm terminates with this conclusion. \square

H.4. Digraphs Admitting an Odd Dijoin

This section is dedicated to the proof of our main result, Theorem H.1.9. The overall strategy to achieve this goal is to work on digraphs and their families of bonds directly. The object that certifies that the bond matroid of a digraph is non-even is called an *odd dijoin*.

Definition H.4.1. Let D be a digraph. A subset $J \subseteq E(D)$ is called an *odd dijoin* if $|J \cap S|$ is odd for every directed bond S in D .

Let D be a digraph. The *contraction* D/A of an edge set $A \subseteq E(D)$ in D is understood as the digraph arising from D by deleting all edges of A and identifying each weak connected component of $D[A]$ into a corresponding vertex. Note that this might produce new loops arising from edges spanned between vertices incident with A but not included in A . Note that contracting a loop is equivalent to deleting the loop.

An edge $e = (x, y)$ of a digraph D , which is not a loop, is said to be *deletable* (or *transitively reducible*) if there is a directed path in D starting in x and ending in y which does not use e . Note that an edge $e \in E(D)$ is deletable if and only if e is a butterfly-contractible element of $M^*(D)$.

For two digraphs D_1, D_2 , we say that D_1 is a *cut minor* of D_2 if it can be obtained from D_2 by a finite series of edge contractions, deletions of deletable edges, and deletions of isolated vertices.

Our next lemma guarantees that the property of admitting an odd dijoin is closed under the cut minor relation.

Lemma H.4.2. *Let D_1, D_2 be digraphs such that D_1 is a cut minor of D_2 . If D_2 admits an odd dijoin, then so does D_1 .*

Proof. The statement follows by applying Lemma H.2.5 to $M^*(D_1)$ and $M^*(D_2)$, noting that deleting isolated vertices from a digraph does not change the induced oriented bond matroid. \square

Our goal will be to characterise the digraphs admitting an odd dijoin in terms of forbidden cut minors. In the following, we prepare this characterisation by providing a set of helpful statements. For an undirected graph G , we define the *cutspace* of G as the \mathbb{F}_2 -linear vector space generated by the bonds in G , whose addition operation is the symmetric difference and whose neutral element is the empty set. The following statements are all obtained in a straightforward way by applying the oriented matroid results Lemma H.3.8, Corollary H.3.10, and Proposition H.3.11 respectively to the oriented bond matroid $M^*(D)$ induced by D .

Corollary H.4.3. *Let D be a weakly connected and acyclic digraph with underlying multi-graph G . Then the cut space of G admits a basis \mathcal{B} whose elements are the edge sets of minimal directed cuts in D . Moreover, if $A \subseteq E(D)$ is a set of edges*

such that D/A is acyclic and $G[A]$ is a forest, then one can choose \mathcal{B} such that every edge $e \in A$ appears in exactly one cut of the basis. \square

Corollary H.4.4. *Let D be a digraph and let \mathcal{B} be a basis of the cut space consisting of minimal directed cuts. Then there is an edge set $J' \subseteq E(D)$ such that $|J' \cap B|$ is odd for all $B \in \mathcal{B}$.* \square

Proposition H.4.5. *Let D be a digraph, \mathcal{B} be a basis of the cut space consisting of directed bonds, and let $J' \subseteq E(D)$ be such that $|B \cap J'|$ is odd for all $B \in \mathcal{B}$. Then the following statements are equivalent:*

(i) D has an odd dijoin.

(ii) If B_1, \dots, B_k are directed bonds of D with k odd, then $\Delta_{i=1}^k B_i \neq \emptyset$.

(iii) Every directed bond of D can be written as the symmetric difference of an odd number of elements of \mathcal{B} .

(iv) J' is an odd dijoin of D . \square

H.4.1. Forbidden cut minors for digraphs with an odd dijoin

Next we characterise the digraphs admitting an odd dijoin in terms of forbidden cut minors. For this purpose, we identify the digraphs without an odd dijoin for which every proper cut minor has an odd dijoin. We call such a digraph a *minimal obstruction*. A digraph $D = (V, E)$ is said to be *oriented* if it has no loops, no parallel, and no anti-parallel edges. Furthermore, D is called *transitively reduced* if for every edge $e = (v, w) \in E$ the only directed path in D starting at v and ending in w consists of e itself, or equivalently, if no edge in D is deletable.

We start with the following crucial lemma, which will be used multiple times to successively find the structure minimal obstructions must have.

Lemma H.4.6. *Let D be a minimal obstruction. Then the underlying multi-graph G of D is 2-vertex-connected. Furthermore, D is oriented, acyclic, and transitively reduced.*

Proof. Assume that D has no odd dijoin, but every cut minor of D has one. Then it is easy to check that $|V(D)| \geq 4$.

To prove that G must be 2-vertex-connected, suppose towards a contradiction that G can be written as the union of two proper subgraphs G_1, G_2 with the property that $|V(G_1) \cap V(G_2)| \leq 1$. Then the orientations D_1, D_2 induced on G_1, G_2 by D are proper cut minors of D : Indeed, for $i \in \{1, 2\}$ we can obtain D_i from D by contracting all edges in D_{3-i} and then deleting all the resulting isolated vertices outside $V(D_i)$. Since D_1, D_2 are proper cut minors of D , they must admit odd dijoins J_1, J_2 , respectively. However, since D_1 and D_2 share at most a single vertex, the directed bonds of D are either directed bonds of D_1 or of D_2 . Hence, the disjoint union $J_1 \cup J_2$ defines an odd dijoin of D and yields the desired contradiction.

To prove acyclicity, assume towards a contradiction that there is a directed cycle C in D . Let us consider the digraph $D/E(C)$. This is a proper cut minor of D and therefore must have an odd dijoin J . However, the directed bonds in $D/E(C)$ are the same as the directed bonds in D edge-disjoint from C , and since C is directed, these are already all the directed bonds of D . Hence J is an odd dijoin also for D , which is a contradiction.

To prove that D is transitively reduced, assume towards a contradiction that there was an edge $e = (x, y) \in E(D)$ and a directed path P from x to y not containing e . Then e is a deletable edge and $D - e$ is a cut minor of D , which therefore must have an odd dijoin $J \subseteq E(D) \setminus \{e\}$. Note that a directed cut S in D either does not intersect $\{e\} \cup E(P)$ at all or contains e and exactly one edge from P . To see this, note first that S cannot contain more than one edge from the directed path P . Since the cut S and the cycle formed by the edges in $\{e\} \cup E(P)$ intersect an even number of times, the set $S \cap (\{e\} \cup E(P))$ is either empty or consists of exactly two arcs, namely e and one arc from P .

We now claim that for every directed bond $B = D[X, Y]$ in D , we get that $B \setminus \{e\}$ is a directed bond of $D - e$. This will then prove that J is also an odd dijoin of D , and yield the desired contradiction. To verify the above statement, consider first the case that $B \cap (\{e\} \cup E(P)) = \emptyset$. Since B is a dibond of D , both $D[X]$ and $D[Y]$ are weakly connected. Because e is contained in the cycle $P \cup e$, which is either fully included in $D[X]$ or in $D[Y]$, it follows that also $(D - e)[X], (D - e)[Y]$ remain weakly connected. In the other case, namely that B contains e and exactly one edge from P , trivially, $(D - e)[X] = D[X]$ and $(D - e)[Y] = D[Y]$ remain weakly connected, and hence $B \setminus \{e\}$ is also a dibond.

Clearly, the fact that D is oriented follows from D being simultaneously acyclic and transitively reduced. This concludes the proof of the lemma. \square

From this, we directly have the following useful observations.

Corollary H.4.7. *Let D be a minimal obstruction. Then for every edge $e \in E(D)$, the digraph D/e is acyclic. Similarly, for every vertex $v \in V(D)$ which is either a source or a sink, the digraph $D/E(v)$, with $E(v) := D[\{v\}, V(D) \setminus \{v\}]$, is acyclic.*

Proof. Let e be an edge of D . Since D is a minimal obstruction, we know by Lemma H.4.6 that e is no loop. Now assume towards a contradiction that there was a directed cycle in D/e . As D itself is acyclic according to Lemma H.4.6, this implies that there is a directed path P in D connecting the end vertices of e , which does not contain e itself. This path together with e now either contradicts the fact that D is acyclic or the fact that D is transitively reduced, both of which hold due to Lemma H.4.6.

For the second part assume w.l.o.g. (using the symmetry given by reversing all edges) that v is a source. Suppose for a contradiction there was a directed cycle in $D/E(v)$. This implies the existence of a directed path P in $D - v$ which connects two different vertices in the neighbourhood of v , say it starts in $w_1 \in N(v)$ and ends in $w_2 \in N(v)$. Now the directed path $(v, w_1) + P$ witnesses that the directed edge (v, w_2) is deletable contradicting that D is transitively reduced. This concludes the proof of the second statement. \square

Lemma H.4.8. *Let D be a minimal obstruction. If $A \subseteq E(D)$ is such that D/A is acyclic and such that $D[A]$ is a forest, then there is a directed bond in D which contains A .*

Proof. By Corollary H.4.3 there is a basis \mathcal{B} of the cut space consisting of directed bonds such that each $e \in A$ is contained in exactly one of the bonds in the basis. Moreover, by Corollary H.4.4 there is $J' \subseteq E(D)$ such that each $B \in \mathcal{B}$ has odd intersection with J' . Since D has no odd dijoin, there has to be a directed bond B_0 in D such that $|B_0 \cap J'|$ is even. Let $B_0 = \Delta_{i=1}^m B_i$ be the unique linear combination with pairwise distinct $B_1, \dots, B_m \in \mathcal{B}$. Clearly, m must be even. Let D' be the cut minor obtained from D by contracting the edges in $E(D) \setminus \bigcup_{i=1}^m B_i$. The bonds B_0, B_1, \dots, B_m are still directed bonds in D' and satisfy $\Delta_{i=0}^m B_i = \emptyset$, while $m + 1$ is odd. The equivalence of (i) and (ii) in Proposition H.4.5 now yields

that D' has no odd dijoin. By the minimality of D we thus must have $D = D'$ and $\bigcup_{i=1}^m B_i = E(D)$. It follows that every $e \in A$ is contained in exactly one of the bonds B_i and thus also in B_0 . Therefore, $B_0 \supseteq A$. \square

Corollary H.4.9. *Let $D = (V, E)$ be a minimal obstruction. For $i \in \{1, 2\}$ let $\emptyset \neq A_i \subseteq E$ be such that $D[A_i]$ is a forest and D/A_i is acyclic. Suppose there is a directed cut $\partial(X)$ in D separating A_1 from A_2 , i.e., such that $A_1 \subseteq E(D[X])$ and $A_2 \subseteq E(D[V \setminus X])$. Then there exists a directed bond in D containing $A_1 \cup A_2$.*

Proof. Let $A := A_1 \dot{\cup} A_2$. As A_1 and A_2 induce vertex-disjoint forests, $D[A]$ is a forest as well. Since no edge is directed from a vertex in $V \setminus X$ to a vertex in X , no directed circuit in D/A can contain a contracted vertex from A_1 and a contracted vertex from A_2 , so every directed circuit must already exist in D/A_1 or in D/A_2 . Because these two digraphs are acyclic, D/A is acyclic. Hence, by Lemma H.4.8, A is fully included in a directed bond of D . This proves the assertion. \square

With the next proposition we shall make the structure of minimal obstructions much more precise. To state the result, we shall make use of the following definition.

Definition H.4.10. Let $n_0, n_1, n_2 \in \mathbb{N}$. Then we denote by $\mathcal{D}(n_0, n_1, n_2)$ the digraph (V, E) , where $V = V_0 \dot{\cup} V_1 \dot{\cup} V_2$ with $V_i = [n_i]$ for $i \in \{1, 2, 3\}$, and $E = (V_0 \times V_1) \dot{\cup} (V_1 \times V_2)$.

Proposition H.4.11. *Let D be a minimal obstruction. Then D is isomorphic to $\mathcal{D}(n_1, n_2, n_3)$ for some integers $n_1, n_2, n_3 \geq 0$.*

Proof. First let us set $D = (V, E)$. We shall split the proof into several claims, starting with the following one.

Claim H.4.12. *D contains no directed path of length 3.*

Suppose towards a contradiction that $v_0, e_1, v_1, e_2, v_2, e_3, v_3$ is a directed path of length 3 in D with $e_1 = (v_0, v_1), e_2 = (v_1, v_2), e_3 = (v_2, v_3)$. By Corollary H.4.7, D/e_1 and D/e_3 are acyclic. Moreover, because D is acyclic by Lemma H.4.6, the edge e_2 is contained in a directed cut $\partial(X)$ in D , separating $\{e_1\}$ and $\{e_3\}$. By Corollary H.4.9 this means that there is a directed bond $\partial(Y)$ in D containing both e_1 and e_2 . This however means that $v_0, v_2 \in Y$ and $v_1, v_3 \notin Y$. Hence, e_2 is an edge in D starting in $V(D) \setminus Y$ and ending in Y , a contradiction since $\partial(Y)$ is a directed bond. This completes the proof of Claim H.4.12.

For $i \in \{0, 1, 2\}$ let V_i denote the set of vertices $v \in V$ such that the longest directed path ending in v has length i . By definition of the V_i and since D is acyclic, there is no edge from a vertex in V_i to a vertex in V_j for $i \geq j$, as otherwise this would give rise to a directed path of length $i + 1$ ending in a vertex of V_j .

By Claim H.4.12 we know that $V = V_0 \dot{\cup} V_1 \dot{\cup} V_2$ holds. We move on by proving the following claim.

Claim H.4.13. *Every vertex $v \in V_1$ is adjacent to every vertex $u \in V_0$.*

Let $v \in V_1$ and $u \in V_0$. Assume for a contradiction that u is not adjacent to v . By definition of V_1 there is an edge $f = (u', v)$ with $u' \in V_0$. By Corollary H.4.7, D/f and $D/E(u)$ are acyclic because u is a source. Let $X \supseteq \{u', v\}$ be the set of all vertices from which v can be reached via a directed path. Clearly $\partial(X)$ is a directed cut in D . As $u \in V_0 \setminus X$ is a source, we conclude that $\{u\} \cup N(u) \subseteq V \setminus X$. This however means that the directed cut $\partial(X)$ separates f from the edges in $E(u)$. By Corollary H.4.9, this means that there is a directed bond $\partial(Y)$ in D containing $E(u) \cup \{f\}$. Since $E(u) = \partial(\{u\})$ itself is a directed cut in D , this contradicts the fact that $\partial(Y)$ is an inclusion-wise minimal directed cut in D , and proves Claim H.4.13.

We proceed with another claim.

Claim H.4.14. *D does not contain any edge from V_0 to V_2 .*

Let $u \in V_0$ and $w \in V_2$. By definition of V_2 there is some $v \in V_1$ such that $(v, w) \in E$. By Claim H.4.13, $(u, v) \in E$. Because D is transitively reduced by Lemma H.4.6, we obtain $(u, w) \notin E$. So the proof of Claim H.4.14 is complete.

Now we come to the last claim we need for the proof of this proposition.

Claim H.4.15. *Every vertex $v \in V_1$ is adjacent to every vertex $w \in V_2$.*

Let $v \in V_1, w \in V_2$ and suppose for a contradiction that w is not adjacent to v . Let $f = (u, v)$ be an edge with $u \in V_0$. By Lemma H.4.8, D/f and $D/E(w)$ are acyclic because w is a sink. Let $X \supseteq \{u, v\}$ be the set of all vertices from which v can be reached via a directed path. Again, $\partial(X)$ forms a directed cut in D . Claim H.4.14 implies that $N(w) \subseteq V_1 \setminus \{v\} \subseteq V \setminus X$. This means $\partial(X)$ separates f from the edges in $E(w)$, contradicting Corollary H.4.9 again.

By combining all four claims we obtain $E = (V_0 \times V_1) \dot{\cup} (V_1 \times V_2)$, and the proof of this proposition is complete. \square

Now Proposition [H.4.11](#) puts us in the comfortable situation that the only possible minimal obstructions to having an odd dijoin are part of a 3-parameter class of simply structured digraphs. The rest of this section is devoted to determine the conditions on n_1, n_2, n_3 that need to be imposed such that $\mathcal{D}(n_1, n_2, n_3)$ is a minimal obstruction. It will be helpful to use the well-known concept of so-called T -joins.

Definition H.4.16. Let G be an undirected graph and $T \subseteq V(G)$ be some vertex set. A subset $J \subseteq E(G)$ of edges is called a T -join, if in the subgraph $H := G[J]$ of G , every vertex in T has odd, and every vertex in $V(G) \setminus T$ has even degree.

The following result is folklore.

Lemma H.4.17. *A graph G with some vertex set $T \subseteq V(G)$ admits a T -join if and only if T has an even number of vertices in each connected component of G .*

We continue with an observation about odd dijoins in digraphs of the form $\mathcal{D}(n_1, n_2, 0)$.

Observation H.4.18. Let $n_1, n_2 \geq 1$. Then the digraph $\mathcal{D}(n_1, n_2, 0) \simeq \mathcal{D}(0, n_1, n_2)$ has an odd dijoin if and only if $\min(n_1, n_2) \leq 1$ or $n_1, n_2 \geq 2$ and $n_1 \equiv n_2 \pmod{2}$.

Proof. If $\min(n_1, n_2) \leq 1$, then all directed bonds in $\mathcal{D}(n_1, n_2, 0)$ consist of single edges. Thus, $J := E(\mathcal{D}(n_1, n_2, 0))$ defines an odd dijoin. If $n_1, n_2 \geq 2$, the directed bonds in $\mathcal{D}(n_1, n_2, 0)$ are exactly those cuts with one vertex on one side of the cut and all other vertices on the other side. Hence, there is an odd dijoin if and only if the complete bipartite graph with partition classes of size n_1, n_2 has a T -join, where T contains all $n_1 + n_2$ vertices. The statement is now implied by Lemma [H.4.17](#). \square

Next we characterise when the digraphs $\mathcal{D}(n_1, n_2, n_3)$ admit an odd dijoin.

Proposition H.4.19. *Let $n_1, n_2, n_3 \geq 1$ be integers. Then $\mathcal{D}(n_1, n_2, n_3)$ has an odd dijoin if and only if one of the following holds:*

- (i) $n_2 = 1$.
- (ii) $n_2 = 2$ and $n_1 \equiv n_3 \pmod{2}$.
- (iii) $n_2 \geq 3$, and $n_1 \equiv n_3 \equiv 1 \pmod{2}$.

Proof. If $n_2 = 1$, then $\mathcal{D}(n_1, n_2, n_3)$ is an oriented star. Clearly, here, the directed bonds consist of single edges, and therefore, $J := E(\mathcal{D}(n_1, 1, n_3))$ defines an odd dijoin.

If $n_2 = 2$, it is easily seen that $\mathcal{D}(n_1, 2, n_3)$ is a planar digraph, which admits a directed planar dual isomorphic to a bicycle $\overset{\leftrightarrow}{C}_{n_1+n_3}$ of length $n_1 + n_3$. By planar duality, we know that $\mathcal{D}(n_1, 2, n_3)$ has an odd dijoin if and only if there is a subset of edges of $\overset{\leftrightarrow}{C}_k$ which intersects every directed cycle an odd number of times. By Theorem H.1.3 we know that such an edge set exists if and only if $n_1 + n_3$ is even, that is, $n_1 \equiv n_3 \pmod{2}$.

Therefore, we assume that $n_2 \geq 3$ for the rest of the proof. We now first show the necessity of (iii). So assume that $D := \mathcal{D}(n_1, n_2, n_3)$ has an odd dijoin J . We observe that the underlying multi-graph of D is 2-connected. Hence, for every vertex $x \in V_1 \cup V_2 \cup V_3$, the cut $E(x)$ of all edges incident with x is a minimal cut of the underlying multi-graph, and it is directed whenever $x \in V_1 \cup V_3$. Therefore, $U(D[J])$ must have odd degree at every vertex in $V_1 \cup V_3$. Moreover, we observe that for any proper non-empty subset $X \subsetneq V_2$, the cut in D induced by the partition $(V_1 \cup X, (V_2 \setminus X) \cup V_3)$ is minimal and directed. In the following, we denote this cut by $F(X)$. Now for every vertex $x \in V_2$, choose some $x' \in V_2 \setminus \{x\}$ and consider the minimal directed cuts $F(\{x'\}), F(\{x, x'\})$. Both are minimal directed cuts (here, we use that $n_2 \geq 3$) and thus must have odd intersection with J . Moreover, the symmetric difference $F(\{x'\}) \Delta F(\{x, x'\})$ contains exactly the set $E(x)$ of edges incident with x in D . We conclude the following:

$$\begin{aligned} |E(x) \cap J| &= |(F(\{x'\}) \Delta F(\{x, x'\})) \cap J| \equiv |F(\{x'\}) \cap J| + |F(\{x, x'\}) \cap J| \\ &\equiv 1 + 1 \equiv 0 \pmod{2} \end{aligned}$$

As $x \in V_2$ was chosen arbitrarily, we conclude that J must be a T -join of the underlying multi-graph of $\mathcal{D}(n_1, n_2, n_3)$ where $T = V_1 \cup V_3$. Now Lemma H.4.17 implies that $|T| = n_1 + n_3$ must be even and hence $n_1 \equiv n_3 \pmod{2}$.

We claim that (iii) must be satisfied, i.e., n_1 and n_3 are odd. Assume towards a contradiction that this is not the case. Hence, by our observation above both n_1 and n_3 are even. Let $x \in V_2$ be some vertex, and consider the directed bond $F(\{x\})$. We can rewrite this bond as the symmetric difference of the directed cut $\partial(V_1) = \{(v_1, v_2) \mid v_1 \in V_1, v_2 \in V_2\}$ and the cut $E(x)$ of all edges incident with x . Because $|E(u) \cap J|$ is odd for every $u \in V_1$, we obtain that

$|\partial(V_1) \cap J| = \sum_{u \in V_1} |E(u) \cap J|$ must be even. However, since also $|E(x) \cap J|$ is even, this means that $|F(\{x\}) \cap J| \equiv |\partial(V_1) \cap J| + |E(x) \cap J| \equiv 0 \pmod{2}$, which is the desired contradiction, as J is an odd dijoin. So (iii) must be satisfied.

To prove the reverse direction, assume that (iii) is fulfilled, i.e., $n_1 \equiv n_3 \equiv 1 \pmod{2}$. We shall construct an odd dijoin of $\mathcal{D}(n_1, n_2, n_3)$. For this purpose, we choose J to be a T -join of the underlying multi-graph where $T = V_1 \cup V_3$. We claim that this defines an odd dijoin of $\mathcal{D}(n_1, n_2, n_3)$. It is not hard to check that the directed bonds of $\mathcal{D}(n_1, n_2, n_3)$ are the cuts $E(v)$ for vertices $v \in V_1 \cup V_3$ and the cuts $F(X)$ as described above, where $\emptyset \neq X \subsetneq V_2$. By the definition of a T -join, all of the directed bonds of the first type have an odd intersection with J , so it suffices to consider the bonds of the second type. Consider again the directed cut $\partial(V_1)$ in $\mathcal{D}(n_1, n_2, n_3)$. For any $\emptyset \neq X \subsetneq V_2$, we can write $F(X)$ as the symmetric difference $F(X) = \partial(V_1) \Delta \Delta_{x \in X} E(x)$. We therefore conclude that

$$\begin{aligned} |F(X) \cap J| &\equiv |\partial(V_1) \cap J| + \sum_{x \in X} \underbrace{|E(x) \cap J|}_{\text{even}} \pmod{2} \\ &\equiv |\partial(V_1) \cap J| = \sum_{x \in V_1} \underbrace{|E(x) \cap J|}_{\text{odd}} \equiv n_1 \equiv 1 \pmod{2}. \end{aligned}$$

This verifies that J is an odd dijoin, and completes the proof of the proposition. \square

We shall now use these insights to characterise minimal obstructions. For this let us first introduce new notation.

Let D be a digraph consisting of a pair h_1, h_2 of ‘‘hub vertices’’ and other vertices x_1, \dots, x_n , where $n \geq 3$, such that for every $i \in [n]$, the vertex x_i has either precisely two outgoing or precisely two incoming edges to both h_1, h_2 , and these are all the edges of D . In this case, we refer to D as a *diamond*. Pause to note that independent of which vertices x_i are sinks and sources, the bond matroid induced by D is always isomorphic to $M^*(\vec{K}_{2,n})$.

Furthermore, we call any digraph isomorphic to \vec{K}_{n_1, n_2} for some $n_1, n_2 \geq 2$, a *one-direction*.

We shall call both, diamonds and one-directions, *odd* if the total number of vertices of these digraphs is odd.

Lemma H.4.20. *All odd diamonds and all odd one-directions are minimal obstructions.*

Proof. It is directly seen from Observation H.4.18 and Proposition H.4.19 that indeed, odd diamonds and odd one-directions do not possess an odd dijoin. Therefore it remains to show that all proper cut minors of these digraphs have odd dijoints. Because both odd diamonds and odd one-directions are weakly 2-connected, transitively reduced and acyclic, the only cut minor operation applicable to them in the first step is the contraction of a single edge. By Lemma H.4.2 it therefore suffices to show that for both types of digraphs, the contraction of any edge results in a digraph admitting an odd dijoin. We first consider odd diamonds. Let $D = \mathcal{D}(n_1, 2, n_3)$ with $n_1, n_2 \geq 1$ and $n_1 + n_2$ odd, and let $e \in E(D)$ be arbitrary. In the planar directed dual graph of D , an odd bicycle with $n_1 + n_2$ vertices, there is a directed dual edge corresponding to e . It is easily seen by duality that D/e has an odd dijoin if and only if the odd bicycle of order $n_1 + n_2 \geq 3$ with a single deleted edge has an edge set intersecting every directed cycle an odd number of times. However, this is the case, because such a digraph is non-even by Theorem H.1.3.

Now we consider odd one-directions. Let $D = \mathcal{D}(n_1, n_2, 0)$ with $n_1, n_2 \geq 2$ and $n_1 + n_2$ odd, and let $e = (x, y) \in E(D)$ be arbitrary. Then in the digraph D/e , define J to be the set of all edges incident with the contraction vertex. It is easily observed that J intersects every minimal directed cut exactly once and thus indeed, every proper cut minor has an odd dijoin. This completes the proof. \square

Now we are able to prove a dual version of Theorem H.1.3 and characterise the existence of odd dijoints in terms of forbidden cut minors.

Theorem H.4.21. *A digraph admits an odd dijoin if and only if it does neither have an odd diamond nor an odd one-direction as a cut minor.*

Proof. By Lemma H.4.2, a digraph has an odd dijoin if and only if it does not contain a minimal obstruction as a cut minor. Hence it suffices to show that a digraph D is a minimal obstruction if and only if it is isomorphic to an odd diamond or an odd one-direction. The fact that these digraphs indeed are minimal obstructions was proved in Lemma H.4.20. So it remains to show that these are the only minimal obstructions.

Let D be an arbitrary minimal obstruction. By Proposition H.4.11 $D \simeq \mathcal{D}(n_1, n_2, n_3)$ for some integers $n_1, n_2, n_3 \geq 0$. By the definition of a minimal obstruction, we know that D has no odd dijoin, while for every edge $e \in E(D)$,

the digraph D/e is a cut minor of D and therefore has one. We know due to Lemma H.4.6 that D is weakly 2-connected. Hence, we either have $\min(n_1, n_3) = 0$, so (by symmetry) w.l.o.g. $n_3 = 0$, or $n_1, n_3 \geq 1$ and therefore $n_2 \geq 2$.

In the first case, we know by Observation H.4.18 and using that D has no odd dijoin, that $n_1, n_2 \geq 2$ and $n_1 \not\equiv n_2 \pmod{2}$. So D is an odd one-direction, which verifies the claim in the case of $\min(n_1, n_3) = 0$.

Next assume that $n_1, n_3 \geq 1$ and $n_2 \geq 2$. Let $e = (x_1, x_2) \in E(D)$ with $x_i \in V_i$ for $i = 1, 2$ be an arbitrary edge going from the first layer V_1 to the second layer V_2 . Denote by c the vertex of D/e corresponding to the contracted edge e . Then in the digraph D/e , all edges $\{(c, v_3) \mid v_3 \in V_3\}$ as well as all the edges in $\{(v_1, v_2) \mid v_1 \in V_1 \setminus \{x_1\}, v_2 \in V_2 \setminus \{x_2\}\}$ admit parallel paths since $n_2 \geq 2$ and, therefore, are deletable. Successive deletion yields a cut minor D' of D/e , and thus of D , with vertex set

$$V(D') = (V_1 \setminus \{x_1\}) \cup \{c\} \cup (V_2 \setminus \{x_2\}) \cup V_3$$

and edge set

$$\begin{aligned} E(D') = & \{(v_1, c) \mid v_1 \in V_1 \setminus \{x_1\}\} \cup \{(c, v_2) \mid v_2 \in V_2 \setminus \{x_2\}\} \\ & \cup \{(v_2, v_3) \mid v_2 \in V_2 \setminus \{x_2\}, v_3 \in V_3\}. \end{aligned}$$

Now after contracting all edges of D' of the set $\{(v_1, c) \mid v_1 \in V_1 \setminus \{x_1\}\}$ we find that D' , and hence D , has a proper cut minor isomorphic to $\mathcal{D}(1, n_2 - 1, n_3)$ with corresponding layers $\{c\}, V_2 \setminus \{x_2\}$ and V_3 .

Applying a symmetric argument (starting by contracting an edge going from V_2 to V_3), we find that D also has a proper cut minor isomorphic to $\mathcal{D}(n_1, n_2 - 1, 1)$.

Using these insights, we now show that $n_2 = 2$ holds. Suppose for a contradiction that $n_2 \geq 3$ holds. Assume first that $n_2 \geq 4$, and therefore $n_2 - 1 \geq 3$. Using statement (iii) of Proposition H.4.19 and that $\mathcal{D}(1, n_2 - 1, n_3)$ and $\mathcal{D}(n_1, n_2 - 1, 1)$ both have odd dijoin, we must have $n_1 \equiv n_3 \equiv 1 \pmod{2}$. In the case that $n_2 = 3$, we similarly observe from statement (ii) of Proposition H.4.19 with the digraphs $\mathcal{D}(1, 2, n_3)$ and $\mathcal{D}(n_1, 2, 1)$ that both n_1 and n_3 must be odd. Now using statement (iii) of Proposition H.4.19 with the digraph $D \simeq \mathcal{D}(n_1, n_2, n_3)$ we can conclude that D must admit an odd dijoin as well, a contradiction.

Hence, we must have $n_2 = 2$. Using again statement (ii) of Proposition H.4.19 with $D \simeq \mathcal{D}(n_1, 2, n_3)$, we get that $n_1 + n_3$ must be odd. Therefore D is isomorphic

to an odd diamond with $2 + n_1 + n_3$ many vertices. This concludes the proof of the theorem. \square

We are now ready to give the proof of Theorem [H.1.9](#).

Proof of Theorem [H.1.9](#). Let \vec{M} be an oriented cographic matroid, and let D be a digraph such that $\vec{M} \simeq M^*(D)$. Let us first note that by definition, \vec{M} is non-even if and only if D has an odd dijoin. Hence, for the equivalence claimed in this theorem it suffices to show that D has an odd dijoin if and only if $M^*(D)$ does not have a GB-minor isomorphic to $\vec{K}_{m,n}$ for $m, n \geq 2$ such that $m+n$ is odd. Suppose first that D has an odd dijoin and $M^*(D)$ is non-even. Then by Lemma [H.2.5](#), every GB-minor of $M^*(D)$ is non-even as well, and hence, no such minor can equal $M^*(\vec{K}_{m,n})$ for any $m, n \geq 2$ with $m+n$ is odd, since $\vec{K}_{m,n}$ does not have an odd dijoin for any such m and n by Lemma [H.4.20](#). This proves the first implication of the equivalence.

Conversely, let us suppose that $M^*(D)$ does not have a GB-minor isomorphic to $M^*(\vec{K}_{m,n})$ for any $m, n \geq 2$ such that $m+n$ is odd. We shall show that D admits an odd dijoin. For this we use Theorem [H.4.21](#) and verify that D has neither an odd diamond nor an odd one-direction as a cut minor. This however follows directly from the fact that the bond matroid induced by any odd diamond of order n is isomorphic to $M^*(\vec{K}_{2,n-2})$ as well as the easy observation that if D' is a cut minor of D , then $M^*(D')$ is a GB-minor of $M^*(D)$. This finishes the proof of the claimed equivalence. \square

H.5. Concluding remarks

For every odd $k \geq 3$ it holds that $M(\vec{C}_k) \simeq M^*(\vec{K}_{k,2}) \simeq M^*(\vec{K}_{2,k})$, and hence, the list of smallest excluded GB-minors characterising non-evenness for oriented cographic matroids strictly extends the list for graphic ones. We find this quite surprising and did not expect it when we initiated our research on the subject.

Seymour [[95](#)] has proved a theorem about generating the class of regular matroids, showing that every regular matroid can be built up from graphic matroids, cographic matroids and a certain 10-element matroid R_{10} by certain sum operations. The matroid R_{10} is regular, but neither graphic nor cographic. It is given by the

following totally unimodular representing matrix:

$$R_{10} = M \left[\begin{array}{c} \left(\begin{array}{ccccccccc} 1 & 0 & 0 & 0 & 0 & -1 & 1 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 & 0 & 1 & -1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 1 & -1 & 1 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 1 & -1 & 1 \\ 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 & 1 & -1 \end{array} \right) \end{array} \right].$$

Seymour introduced three different kinds of sum operation which join two regular matroids M_1 and M_2 whose element sets are either disjoint (1-sum), intersect in a single non-loop element (2-sum) or in a common 3-circuit (3-sum) into a bigger regular matroid $M_1 \Delta M_2$ (for a precise definition of these operations we refer to the introduction of [95]).

Theorem H.5.1 ([95]). *Every regular matroid can be built up from graphic matroids, cographic matroids and R_{10} by repeatedly applying 1-sums, 2-sums and 3-sums.*

This theorem shows that graphic matroids, cographic matroids and R_{10} constitute the most important building blocks of regular matroids. Using a brute force implementation, we checked by computer that every orientation of R_{10} containing no $M^*(\vec{K}_{m,n})$ as a GB-minor for any $m, n \geq 2$ such that $m + n$ is odd, is already non-even. We therefore expect the total list of forbidden minors for all non-even oriented matroids to not be larger than the union of the forbidden minors for graphic (Proposition H.1.8) and cographic (Theorem H.1.9) non-even oriented matroids. In other words, we conjecture the following.

Conjecture H.5.2. *A regular oriented matroid M is non-even if and only if none of its GB-minors is isomorphic to $M^*(\vec{K}_{m,n})$ for some $m, n \geq 2$ such that $m + n$ is odd.*

The natural way of working on this conjecture would be to try and show that a smallest counterexample is not decomposable as the 1-, 2- or 3-sum of two smaller oriented regular matroids. Apart from the obvious open problem of resolving the computational complexity of the even circuit problem (Problem H.1.4) for regular oriented matroids in general, already resolving the case of cographic matroids would be interesting.

Problem H.5.3. Is there a polynomially bounded algorithm that, given as input a digraph D , decides whether or not D contains a directed bond of even size? Equivalently, is there a polynomially bounded recognition algorithm for digraphs admitting an odd dijoin?

Conclusively, given our characterisation of digraphs admitting an odd dijoin in terms of forbidden cut minors, the following question naturally comes up.

Problem H.5.4. Let F be a fixed digraph. Is there a polynomially bounded algorithm that, given as input a digraph D , decides whether or not D contains a cut minor isomorphic to F ?

Bibliography

- [1] R. Aharoni and E. Berger, *Menger's theorem for infinite graphs*, *Invent. Math.* **176** (2009), no. 1, 1–62, DOI [10.1007/s00222-008-0157-3](https://doi.org/10.1007/s00222-008-0157-3). MR2485879 ↑248
- [2] T. Andreae, *Classes of locally finite ubiquitous graphs*, *J. Combin. Theory Ser. B* **103** (2013), no. 2, 274–290, DOI [10.1016/j.jctb.2012.11.003](https://doi.org/10.1016/j.jctb.2012.11.003). MR3018070 ↑17, 29, 30, 47, 92
- [3] ———, *On disjoint configurations in infinite graphs*, *J. Graph Theory* **39** (2002), no. 4, 222–229, DOI [10.1002/jgt.10016](https://doi.org/10.1002/jgt.10016). MR1894467 ↑8, 17, 47, 91
- [4] ———, *On a problem of R. Halin concerning infinite graphs. II*, *Discrete Math.* **29** (1980), no. 3, 219–227, DOI [10.1016/0012-365X\(80\)90149-1](https://doi.org/10.1016/0012-365X(80)90149-1). MR560764 ↑91
- [5] ———, *Über eine Eigenschaft lokalfiniten, unendlicher Bäume*, *J. Combin. Theory Ser. B* **27** (1979), no. 2, 202–215, DOI [10.1016/0095-8956\(79\)90082-0](https://doi.org/10.1016/0095-8956(79)90082-0) (German, with English summary). MR546863 ↑18, 19, 23, 29, 30, 33, 91, 107
- [6] ———, *Bemerkung zu einem Problem aus der Theorie der unendlichen Graphen*, *Abh. Math. Sem. Univ. Hamburg* **46** (1977), 91–95, DOI [10.1007/BF02993015](https://doi.org/10.1007/BF02993015) (German). MR505832 ↑8, 17, 91
- [7] A. Bachem and A. Reinhold, *On the complexity of the Farkas-Property of oriented matroids* (1989). <http://e-archive.informatik.uni-koeln.de/65> Preprint. ↑272
- [8] J. Bang-Jensen and G. Gutin, *Digraphs*, 2nd ed., Springer Monographs in Mathematics, Springer-Verlag London, Ltd., London, 2009. Theory, algorithms and applications. MR2472389 ↑236, 264
- [9] C. Berge, *Hypergraphs*, North-Holland Mathematical Library, vol. 45, North-Holland Publishing Co., Amsterdam, 1989. Combinatorics of finite sets; Translated from the French. MR1013569 ↑235, 236, 242, 244
- [10] R. E. Bixby and W. H. Cunningham, *Converting linear programs to network problems*, *Math. Oper. Res.* **5** (1980), no. 3, 321–357, DOI [10.1287/moor.5.3.321](https://doi.org/10.1287/moor.5.3.321). MR594849 ↑274
- [11] A. Björner, M. Las Vergnas, B. Sturmfels, N. White, and G. M. Ziegler, *Oriented matroids*, 2nd ed., *Encyclopedia of Mathematics and its Applications*, vol. 46, Cambridge University Press, Cambridge, 1999. MR1744046 ↑266, 269
- [12] R. G. Bland and M. Las Vergnas, *Orientability of matroids*, *J. Combinatorial Theory Ser. B* **24** (1978), no. 1, 94–123, DOI [10.1016/0095-8956\(78\)90080-1](https://doi.org/10.1016/0095-8956(78)90080-1). MR485461 ↑261, 268, 269, 274

- [13] N. Bowler, C. Elbracht, J. Erde, J. P. Gollin, K. Heuer, M. Pitz, and M. Teegen, *Topological ubiquity of trees*, J. Combinatorial Theory Ser. B **157** (2022), 70–95, DOI [10.1016/j.jctb.2022.05.011](https://doi.org/10.1016/j.jctb.2022.05.011). MR4438889 ↑7, 8, 51, 56, 80, 81, 91, 94, 96, 101, 105, 128, 129
- [14] ———, *Ubiquity in graphs II: Ubiquity of graphs with nowhere-linear end structure*, available at [arXiv:1809.00602](https://arxiv.org/abs/1809.00602). Preprint. ↑7, 8, 18, 19, 28, 97, 104, 106, 118, 119, 126, 128, 130, 132
- [15] ———, *Ubiquity in graphs III: Ubiquity of locally finite graphs with extensive tree-decompositions*, available at [arXiv:2012.13070](https://arxiv.org/abs/2012.13070). Preprint. ↑7, 8, 18, 28, 48, 50, 65, 85
- [16] H. Broersma and H. J. Veldman, *Restrictions on induced subgraphs ensuring Hamiltonicity or pancyclicity of $K_{1,3}$ -free graphs*, Contemporary methods in graph theory, Bibliographisches Inst., Mannheim, 1990, pp. 181–194. MR1126227 ↑192
- [17] H. Bruhn and M. Stein, *On end degrees and infinite cycles in locally finite graphs*, Combinatorica **27** (2007), no. 3, 269–291, DOI [10.1007/s00493-007-2149-0](https://doi.org/10.1007/s00493-007-2149-0). MR2345811 ↑166, 168
- [18] H. Bruhn and X. Yu, *Hamilton Cycles in Planar Locally Finite Graphs*, SIAM J. Discrete Math. **22** (2008), no. 4, 1381–1392, DOI [10.1137/050631458](https://doi.org/10.1137/050631458). MR2345811 ↑160
- [19] G. Călinescu, A. Dumitrescu, and J. Pach, *Reconfigurations in graphs and grids*, SIAM J. Discrete Math. **22** (2008), no. 1, 124–138, DOI [10.1137/060652063](https://doi.org/10.1137/060652063). MR2383232 ↑58
- [20] T. L. Chan, *Contractible edges in 2-connected locally finite graphs*, Electron. J. Combin. **22** (2015), no. 2, Paper 2.47, 13, DOI [10.37236/4414](https://doi.org/10.37236/4414). MR3367290 ↑160
- [21] M. Chudnovsky, K. Edwards, R. Kim, A. Scott, and P. Seymour, *Disjoint dijoins*, J. Combin. Theory Ser. B **120** (2016), 18–35, DOI [10.1016/j.jctb.2016.04.002](https://doi.org/10.1016/j.jctb.2016.04.002). MR3504077 ↑234
- [22] G. Cornuéjols, *Combinatorial optimization: Packing and covering*, CBMS-NSF Regional Conference Series in Applied Mathematics, vol. 74, Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 2001. MR1828452 ↑233
- [23] G. Cornuéjols and B. Guenin, *On dijoins*, Discrete Math. **243** (2002), no. 1-3, 213–216, DOI [10.1016/S0012-365X\(01\)00209-6](https://doi.org/10.1016/S0012-365X(01)00209-6). MR1874739 ↑234
- [24] R. Diestel, *Graph theory*, 5th ed., Graduate Texts in Mathematics, vol. 173, Springer, Berlin, 2017. MR3644391 ↑19, 20, 24, 26, 40, 51, 52, 83, 87, 95, 101, 102, 107, 108, 159, 163, 165, 167, 168, 191, 193, 195, 213, 215, 217, 218, 236, 243, 264
- [25] ———, *Locally finite graphs with ends: a topological approach* (2012), available at [arXiv:0912.4213v3](https://arxiv.org/abs/0912.4213v3). Post-publication manuscript. ↑159, 163, 168, 191, 193, 213, 215
- [26] R. Diestel and D. Kühn, *Topological paths, cycles and spanning trees in infinite graphs*, European J. Combin. **25** (2004), no. 6, 835–862, DOI [10.1016/j.ejc.2003.01.002](https://doi.org/10.1016/j.ejc.2003.01.002). MR2079902 ↑168, 218

- [27] ———, *On infinite cycles I*, *Combinatorica* **24** (2004), no. 1, 69–89, DOI [10.1007/s00493-004-0005-z](https://doi.org/10.1007/s00493-004-0005-z). MR2057684 ↑9, 159, 191, 213
- [28] ———, *On infinite cycles II*, *Combinatorica* **24** (2004), no. 1, 91–116, DOI [10.1007/s00493-004-0006-y](https://doi.org/10.1007/s00493-004-0006-y). MR2057685 ↑9, 159, 191, 213
- [29] ———, *Graph-theoretical versus topological ends of graphs*, *J. Combin. Theory Ser. B* **87** (2003), no. 1, 197–206, DOI [10.1016/S0095-8956\(02\)00034-5](https://doi.org/10.1016/S0095-8956(02)00034-5). Dedicated to Crispin St. J. A. Nash-Williams. MR1967888 ↑18, 32, 159, 165, 195, 218
- [30] R. Diestel and R. Thomas, *Excluding a countable clique*, *J. Combin. Theory Ser. B* **76** (1999), no. 1, 41–67, DOI [10.1006/jctb.1998.1886](https://doi.org/10.1006/jctb.1998.1886). MR1687330 ↑65
- [31] D. Duffus, M. S. Jacobson, and R. J. Gould, *Forbidden subgraphs and the Hamiltonian theme*, *The theory and applications of graphs* (Kalamazoo, Mich., 1980), Wiley, New York, 1981, pp. 297–316. MR634535 ↑161
- [32] J. Edmonds and R. Giles, *A min-max relation for submodular functions on graphs*, *Studies in integer programming* (Proc. Workshop, Bonn, 1975), North-Holland, Amsterdam, 1977, pp. 185–204. *Ann. of Discrete Math.*, Vol. 1. MR0460169 ↑234, 239
- [33] P. Feofiloff and D. H. Younger, *Directed cut transversal packing for source-sink connected graphs*, *Combinatorica* **7** (1987), no. 3, 255–263, DOI [10.1007/BF02579302](https://doi.org/10.1007/BF02579302). MR918396 ↑234, 235, 236, 253
- [34] H. Fleischner, *The square of every two-connected graph is Hamiltonian*, *J. Combinatorial Theory Ser. B* **16** (1974), 29–34, DOI [10.1016/0095-8956\(74\)90091-4](https://doi.org/10.1016/0095-8956(74)90091-4). MR332573 ↑213
- [35] A. Frank and É. Tardos, *An application of simultaneous Diophantine approximation in combinatorial optimization*, *Combinatorica* **7** (1987), no. 1, 49–65, DOI [10.1007/BF02579200](https://doi.org/10.1007/BF02579200). MR905151 ↑280
- [36] H. Freudenthal, *Über die Enden topologischer Räume und Gruppen*, *Math. Z.* **33** (1931), no. 1, 692–713, DOI [10.1007/BF01174375](https://doi.org/10.1007/BF01174375) (German). MR1545233 ↑165, 195, 218
- [37] P. Gács and L. Lovász, *Khachiyan’s algorithm for linear programming*, *Math. Programming Stud.* **14** (1981), 61–68, DOI [10.1007/bfb0120921](https://doi.org/10.1007/bfb0120921). MR600122 ↑280
- [38] A. Georgakopoulos, *Infinite Hamilton cycles in squares of locally finite graphs*, *Adv. Math.* **220** (2009), no. 3, 670–705, DOI [10.1016/j.aim.2008.09.014](https://doi.org/10.1016/j.aim.2008.09.014). MR2483226 ↑160, 213
- [39] S. Goodman and S. Hedetniemi, *Sufficient conditions for a graph to be Hamiltonian*, *J. Combinatorial Theory Ser. B* **16** (1974), 175–180, DOI [10.1016/0095-8956\(74\)90061-6](https://doi.org/10.1016/0095-8956(74)90061-6). MR357222 ↑192
- [40] J. P. Gollin and K. Heuer, *On the infinite Lucchesi-Younger conjecture I*, *J. Graph Theory* **98** (2021), no. 1, 27–48, DOI [10.1002/jgt.22680](https://doi.org/10.1002/jgt.22680). MR4313226 ↑234, 249, 250
- [41] J. P. Gollin, K. Heuer, and K. Stavropoulos, *Disjoint dijoins for classes of dicuts in finite and infinite digraphs*, available at [arXiv:2109.03518](https://arxiv.org/abs/2109.03518). Preprint. ↑7, 11, 12

- [42] J. P. Gollin and K. Heuer, *Characterising k -connected sets in infinite graphs*, available at [arXiv:1811.06411](https://arxiv.org/abs/1811.06411). Preprint. ↑82
- [43] ———, *On the infinite Lucchesi-Younger Conjecture II*. In preparation. ↑234
- [44] B. Guenin and R. Thomas, *Packing directed circuits exactly*, *Combinatorica* **31** (2011), no. 4, 397–421, DOI [10.1007/s00493-011-1687-5](https://doi.org/10.1007/s00493-011-1687-5). MR2861237 ↑259
- [45] L. G. Hačijan, *A polynomial algorithm in linear programming*, *Dokl. Akad. Nauk SSSR* **244** (1979), no. 5, 1093–1096 (Russian). MR522052 ↑280
- [46] R. Halin, *A problem in infinite graph-theory*, *Abh. Math. Sem. Univ. Hamburg* **43** (1975), 79–84, DOI [10.1007/BF02995937](https://doi.org/10.1007/BF02995937). MR427153 ↑18, 50, 84
- [47] ———, *A note on Menger’s theorem for infinite locally finite graphs*, *Abh. Math. Sem. Univ. Hamburg* **40** (1974), 111–114, DOI [10.1007/BF02993589](https://doi.org/10.1007/BF02993589). MR335355 ↑
- [48] ———, *Die Maximalzahl fremder zweiseitig unendlicher Wege in Graphen*, *Math. Nachr.* **44** (1970), 119–127, DOI [10.1002/mana.19700440109](https://doi.org/10.1002/mana.19700440109) (German). MR270953 ↑17, 91
- [49] ———, *Über die Maximalzahl fremder unendlicher Wege in Graphen*, *Math. Nachr.* **30** (1965), 63–85, DOI [10.1002/mana.19650300106](https://doi.org/10.1002/mana.19650300106) (German). MR190031 ↑7, 17, 47, 52, 91, 102, 112
- [50] M. Hamann, F. Lehner, and J. Pott, *Extending cycles locally to Hamilton cycles*, *Electron. J. Combin.* **23** (2016), no. 1, Paper 1.49, 17, DOI [10.37236/3960](https://doi.org/10.37236/3960). MR3484754 ↑160
- [51] D. Hausmann and B. Korte, *Algorithmic versus axiomatic definitions of matroids*, *Math. Programming Stud.* **14** (1981), 98–111, DOI [10.1007/bfb0120924](https://doi.org/10.1007/bfb0120924). MR600125 ↑272
- [52] K. Heuer, *Hamilton-laceable bi-powers of locally finite bipartite graphs*, available at [arXiv:2104.03199v2](https://arxiv.org/abs/2104.03199v2). Extended preprint. ↑7, 9, 11
- [53] ———, *Hamilton-laceable bi-powers of locally finite bipartite graphs*, *Discrete Math.* **345** (2022), no. 7, Paper No. 112777, DOI [10.1016/j.disc.2021.112777](https://doi.org/10.1016/j.disc.2021.112777). MR4390911 ↑7, 9, 10, 11, 160
- [54] ———, *Hamiltonicity in locally finite graphs: two extensions and a counterexample*, *Electron. J. Combin.* **25** (2018), no. 3, Paper 3.13, 29, DOI [10.37236/6773](https://doi.org/10.37236/6773). MR3853865 ↑160
- [55] ———, *A sufficient local degree condition for Hamiltonicity in locally finite claw-free graphs*, *European J. Combin.* **55** (2016), 82–99, DOI [10.1016/j.ejc.2016.01.003](https://doi.org/10.1016/j.ejc.2016.01.003). MR3474794 ↑160, 200
- [56] ———, *A sufficient condition for Hamiltonicity in locally finite graphs*, *European J. Combin.* **45** (2015), 97–114, DOI [10.1016/j.ejc.2014.08.025](https://doi.org/10.1016/j.ejc.2014.08.025). MR3286624 ↑160, 169, 195, 196, 200
- [57] K. Heuer and D. Sarikaya, *Forcing Hamiltonicity in locally finite graphs via forbidden induced subgraphs I: nets and bulls* (2020), available at [arXiv:2006.09160](https://arxiv.org/abs/2006.09160). Preprint. ↑7, 9, 10, 193, 195, 197

- [58] ———, *Forcing Hamiltonicity in locally finite graphs via forbidden induced subgraphs II: paws* (2020), available at [arXiv:2006.09166](https://arxiv.org/abs/2006.09166). Preprint. ↑7, 9, 10, 161
- [59] K. Heuer, R. Steiner, and S. Wiederrecht, *Even circuits in oriented matroids*, *Comb. Theory* **2** (2022), no. 1, Paper No. 3, 32, DOI [10.5070/c62156875](https://doi.org/10.5070/c62156875). MR4405992 ↑7, 11, 12, 13
- [60] T. Jech, *Set theory*, Springer Monographs in Mathematics, Springer-Verlag, Berlin, 2003. The third millennium edition, revised and expanded. MR1940513 ↑236
- [61] T. Johnson, N. Robertson, P. D. Seymour, and R. Thomas, *Directed tree-width*, *J. Combin. Theory Ser. B* **82** (2001), no. 1, 138–154, DOI [10.1006/jctb.2000.2031](https://doi.org/10.1006/jctb.2000.2031). MR1828440 ↑259
- [62] H. A. Jung, *Wurzelbäume und unendliche Wege in Graphen*, *Math. Nachr.* **41** (1969), 1–22, DOI [10.1002/mana.19690410102](https://doi.org/10.1002/mana.19690410102) (German). MR266807 ↑216, 217, 224
- [63] K. Kawarabayashi and S. Kreutzer, *The Directed Grid Theorem*, *Proceedings of the Forty-Seventh Annual ACM Symposium on Theory of Computing*, 2015, pp. 655–664, DOI [10.1145/2746539.2746586](https://doi.org/10.1145/2746539.2746586), available at [arXiv:1411.5681](https://arxiv.org/abs/1411.5681). ↑259
- [64] V. Klee, R. Ladner, and R. Manber, *Signsolvability revisited*, *Linear Algebra Appl.* **59** (1984), 131–157, DOI [10.1016/0024-3795\(84\)90164-2](https://doi.org/10.1016/0024-3795(84)90164-2). MR743051 ↑259
- [65] P. Komjáth, *A note on minors of uncountable graphs*, *Math. Proc. Cambridge Philos. Soc.* **117** (1995), no. 1, 7–9, DOI [10.1017/S0305004100072881](https://doi.org/10.1017/S0305004100072881). MR1297892 ↑91
- [66] D. Kornhauser, G. Miller, and P. Spirakis, *Coordinating Pebble Motion On Graphs, The Diameter Of Permutation Groups, And Applications*, *Proceedings of the 25th Annual Symposium on Foundations of Computer Science*, 1984, 1984, pp. 241–250, DOI [10.1109/SFCS.1984.715921](https://doi.org/10.1109/SFCS.1984.715921). ↑58
- [67] I. Kříž and R. Thomas, *The Menger-like property of the tree-width of infinite graphs*, *J. Combin. Theory Ser. B* **52** (1991), no. 1, 86–91, DOI [10.1016/0095-8956\(91\)90093-Y](https://doi.org/10.1016/0095-8956(91)90093-Y). MR1109422 ↑114
- [68] A. Kündgen, B. Li, and C. Thomassen, *Cycles through all finite vertex sets in infinite graphs*, *European J. Combin.* **65** (2017), 259–275, DOI [10.1016/j.ejc.2017.06.006](https://doi.org/10.1016/j.ejc.2017.06.006). MR3679848 ↑221
- [69] J. Lake, *A problem concerning infinite graphs*, *Discrete Math.* **14** (1976), no. 4, 343–345, DOI [10.1016/0012-365X\(76\)90066-2](https://doi.org/10.1016/0012-365X(76)90066-2). MR419297 ↑17, 91
- [70] R. Laver, *Better-quasi-orderings and a class of trees*, *Studies in foundations and combinatorics*, *Adv. in Math. Suppl. Stud.*, vol. 1, Academic Press, New York-London, 1978, pp. 31–48. MR520553 ↑18, 22
- [71] ———, *On Fraïssé’s order type conjecture*, *Ann. of Math. (2)* **93** (1971), 89–111, DOI [10.2307/1970754](https://doi.org/10.2307/1970754). MR279005 ↑48
- [72] F. Laviolette, *Decompositions of infinite graphs. I. Bond-faithful decompositions*, *J. Combin. Theory Ser. B* **94** (2005), no. 2, 259–277, DOI [10.1016/j.jctb.2005.01.003](https://doi.org/10.1016/j.jctb.2005.01.003). MR2145516 ↑235, 256

- [73] O. Lee and Y. Wakabayashi, *Note on a min-max conjecture of Woodall*, J. Graph Theory **38** (2001), no. 1, 36–41, DOI [10.1002/jgt.1022](https://doi.org/10.1002/jgt.1022). MR1849557 ↑233, 234
- [74] O. Lee and A. Williams, *Packing dicycle covers in planar graphs with no $K^5 - e$ minor*, LATIN 2006: Theoretical informatics, Lecture Notes in Comput. Sci., vol. 3887, Springer, Berlin, 2006, pp. 677–688, DOI [10.1007/11682462_62](https://doi.org/10.1007/11682462_62). MR2256372 ↑234, 235, 236, 251
- [75] F. Lehner, *On spanning tree packings of highly edge connected graphs*, J. Combin. Theory Ser. B **105** (2014), 93–126, DOI [10.1016/j.jctb.2013.12.004](https://doi.org/10.1016/j.jctb.2013.12.004). MR3171783 ↑160
- [76] B. Li, *Faithful subgraphs and Hamiltonian circles of infinite graphs* (2019), available at [arXiv:1902.06402](https://arxiv.org/abs/1902.06402). Preprint. ↑160
- [77] ———, *Hamiltonicity of bi-power of bipartite graphs, for finite and infinite cases* (2019), available at [arXiv:1902.06403](https://arxiv.org/abs/1902.06403). Preprint. ↑160, 213, 214, 215, 219, 221, 222
- [78] C. L. Lucchesi and D. H. Younger, *A minimax theorem for directed graphs*, J. London Math. Soc. (2) **17** (1978), no. 3, 369–374, DOI [10.1112/jlms/s2-17.3.369](https://doi.org/10.1112/jlms/s2-17.3.369). MR500618 ↑233
- [79] R. Manber and J. Y. Shao, *On digraphs with the odd cycle property*, J. Graph Theory **10** (1986), no. 2, 155–165, DOI [10.1002/jgt.3190100203](https://doi.org/10.1002/jgt.3190100203). MR890220 ↑259
- [80] W. McCuaig, *Pólya’s permanent problem*, Electron. J. Combin. **11** (2004), no. 1, Research Paper 79, 83, DOI [10.37236/1832](https://doi.org/10.37236/1832). MR2114183 ↑13, 259, 273, 274
- [81] A. Mészáros, *Note: a note on disjoint dijoins*, Combinatorica **38** (2018), no. 6, 1485–1488, DOI [10.1007/s00493-018-3862-6](https://doi.org/10.1007/s00493-018-3862-6). MR3910884 ↑234
- [82] G. J. Minty, *On the axiomatic foundations of the theories of directed linear graphs, electrical networks, and network programming*, J. Math. Mech. **15** (1966), 485–520. ↑269
- [83] C. St. J. A. Nash-Williams, *On well-quasi-ordering infinite trees*, Proc. Cambridge Philos. Soc. **61** (1965), 697–720, DOI [10.1017/s0305004100039062](https://doi.org/10.1017/s0305004100039062). MR175814 ↑18, 21, 154
- [84] J. Oxley, *Matroid theory*, 2nd ed., Oxford Graduate Texts in Mathematics, vol. 21, Oxford University Press, Oxford, 2011. MR2849819 ↑264, 265, 275
- [85] M. Pitz, *A note on minor antichains of uncountable graphs* (2020), available at [arXiv:2005.05816](https://arxiv.org/abs/2005.05816). Preprint. ↑91
- [86] M. F. Pitz, *Hamilton cycles in infinite cubic graphs*, Electron. J. Combin. **25** (2018), no. 3, Paper 3.3, 11, DOI [10.37236/7033](https://doi.org/10.37236/7033). MR3829289 ↑160
- [87] H. E. Robbins, *Questions, Discussions, and Notes: A Theorem on Graphs, with an Application to a Problem of Traffic Control*, Amer. Math. Monthly **46** (1939), no. 5, 281–283, DOI [10.2307/2303897](https://doi.org/10.2307/2303897). MR1524589 ↑233
- [88] N. Robertson and P. Seymour, *Graph minors XXIII. Nash-Williams’ immersion conjecture*, J. Combin. Theory Ser. B **100** (2010), no. 2, 181–205, DOI [10.1016/j.jctb.2009.07.003](https://doi.org/10.1016/j.jctb.2009.07.003). MR2595703 ↑112

- [89] N. Robertson and P. D. Seymour, *Graph minors. IV. Tree-width and well-quasi-ordering*, J. Combin. Theory Ser. B **48** (1990), no. 2, 227–254, DOI [10.1016/0095-8956\(90\)90120-O](https://doi.org/10.1016/0095-8956(90)90120-O). MR1046757 ↑111
- [90] N. Robertson, P. D. Seymour, and R. Thomas, *Permanents, Pfaffian orientations, and even directed circuits*, Ann. of Math. (2) **150** (1999), no. 3, 929–975, DOI [10.2307/121059](https://doi.org/10.2307/121059). MR1740989 ↑13, 259, 273, 274
- [91] Z. Ryjáček, *Hamiltonicity in claw-free graphs through induced bulls*, Discrete Math. **140** (1995), no. 1-3, 141–147, DOI [10.1016/0012-365X\(94\)00292-Q](https://doi.org/10.1016/0012-365X(94)00292-Q). MR1333716 ↑162
- [92] A. Schrijver, *Min-max relations for directed graphs*, Bonn Workshop on Combinatorial Optimization (Bonn, 1980), Ann. Discrete Math., vol. 16, North-Holland, Amsterdam-New York, 1982, pp. 261–280. MR686312 ↑234
- [93] ———, *A counterexample to a conjecture of Edmonds and Giles*, Discrete Math. **32** (1980), no. 2, 213–215, DOI [10.1016/0012-365X\(80\)90057-6](https://doi.org/10.1016/0012-365X(80)90057-6). MR592858 ↑234, 239
- [94] M. Sekanina, *On an ordering of the set of vertices of a connected graph*, Spisy Přírod. Fak. Univ. Brno **1960** (1960), 137–141 (English, with Russian summary). MR0140095 ↑213
- [95] P. D. Seymour, *Decomposition of regular matroids*, J. Combin. Theory Ser. B **28** (1980), no. 3, 305–359, DOI [10.1016/0095-8956\(80\)90075-1](https://doi.org/10.1016/0095-8956(80)90075-1). MR579077 ↑296, 297
- [96] P. Seymour and C. Thomassen, *Characterization of even directed graphs*, J. Combin. Theory Ser. B **42** (1987), no. 1, 36–45, DOI [10.1016/0095-8956\(87\)90061-X](https://doi.org/10.1016/0095-8956(87)90061-X). MR872406 ↑12, 13, 259, 260, 278
- [97] S. Shelah, *A compactness theorem for singular cardinals, free algebras, Whitehead problem and transversals*, Israel J. Math. **21** (1975), no. 4, 319–349, DOI [10.1007/BF02757993](https://doi.org/10.1007/BF02757993). MR389579 ↑18, 43
- [98] F. B. Shepherd, *Hamiltonicity in claw-free graphs*, J. Combin. Theory Ser. B **53** (1991), no. 2, 173–194, DOI [10.1016/0095-8956\(91\)90074-T](https://doi.org/10.1016/0095-8956(91)90074-T). MR1129551 ↑10, 160, 162, 169, 179
- [99] F. B. Shepherd and A. Vetta, *Visualizing, finding and packing dijoins*, Graph theory and combinatorial optimization, GERAD 25th Anniv. Ser., vol. 8, Springer, New York, 2005, pp. 219–254, DOI [10.1007/0-387-25592-3_8](https://doi.org/10.1007/0-387-25592-3_8). MR2180135 ↑234
- [100] L. Soukup, *Elementary submodels in infinite combinatorics*, Discrete Math. **311** (2011), no. 15, 1585–1598, DOI [10.1016/j.disc.2011.01.025](https://doi.org/10.1016/j.disc.2011.01.025). MR2800978 ↑256
- [101] É. Tardos, *A strongly polynomial algorithm to solve combinatorial linear programs*, Oper. Res. **34** (1986), no. 2, 250–256, DOI [10.1287/opre.34.2.250](https://doi.org/10.1287/opre.34.2.250). MR861043 ↑280
- [102] R. Thomas, *Well-quasi-ordering infinite graphs with forbidden finite planar minor*, Trans. Amer. Math. Soc. **312** (1989), no. 1, 279–313, DOI [10.2307/2001217](https://doi.org/10.2307/2001217). MR932450 ↑93, 111, 116, 154

- [103] ———, *A counterexample to “Wagner’s conjecture” for infinite graphs*, Math. Proc. Cambridge Philos. Soc. **103** (1988), no. 1, 55–57, DOI [10.1017/S0305004100064616](https://doi.org/10.1017/S0305004100064616). MR913450 ↑18, 47, 91
- [104] C. Thomassen, *Sign-nonsingular matrices and even cycles in directed graphs*, Linear Algebra Appl. **75** (1986), 27–41, DOI [10.1016/0024-3795\(86\)90179-5](https://doi.org/10.1016/0024-3795(86)90179-5). MR825397 ↑259
- [105] ———. Personal communication. ↑234
- [106] W. T. Tutte, *An algorithm for determining whether a given binary matroid is graphic*, Proc. Amer. Math. Soc. **11** (1960), 905–917, DOI [10.2307/2034435](https://doi.org/10.2307/2034435). MR117173 ↑274
- [107] ———, *A homotopy theorem for matroids. I, II*, Trans. Amer. Math. Soc. **88** (1958), 144–174, DOI [10.2307/1993243](https://doi.org/10.2307/1993243). MR101526 ↑265
- [108] M. E. Watkins, *Infinite paths that contain only shortest paths*, J. Combin. Theory Ser. B **41** (1986), no. 3, 341–355, DOI [10.1016/0095-8956\(86\)90055-9](https://doi.org/10.1016/0095-8956(86)90055-9). MR864581 ↑167
- [109] D. J. A. Welsh, *Matroid theory*, Academic Press [Harcourt Brace Jovanovich, Publishers], London-New York, 1976. L. M. S. Monographs, No. 8. MR0427112 ↑264
- [110] A. M. Williams and B. Guenin, *Advances in packing directed joins*, Proceedings of GRACO2005, Electron. Notes Discrete Math., vol. 19, Elsevier Sci. B. V., Amsterdam, 2005, pp. 249–255, DOI [10.1016/j.endm.2005.05.034](https://doi.org/10.1016/j.endm.2005.05.034). MR2173794 ↑234
- [111] D. R. Woodall, *Menger and König systems*, Theory and applications of graphs (Proc. Internat. Conf., Western Mich. Univ., Kalamazoo, Mich., 1976), Lecture Notes in Math., vol. 642, Springer, Berlin, 1978, pp. 620–635. MR499529 ↑11, 233
- [112] ———, *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** (1976), no. 2, 132–134, DOI [10.1016/0095-8956\(76\)90052-6](https://doi.org/10.1016/0095-8956(76)90052-6). MR427154 ↑17, 91
- [113] http://lemon.cs.elte.hu/egres/open/Woodall's_conjecture. ↑233
- [114] http://www.openproblemgarden.org/op/woodalls_conjecture. ↑233