

**The Eulerian problem
and further results
in the theory of infinite graphs**

Habilitationschrift

vorgelegt im
Fachbereich Mathematik
Universität Hamburg

von

Max F. Pitz

Hamburg, im Mai 2019

Contents

Chapter 1. Overview	7
Part 1. Structural infinite graph theory	11
Chapter 2. A counterexample to the reconstruction conjecture for locally finite trees	13
2.1. Introduction	13
2.2. Sketch of the construction	15
2.3. Closure with respect to promises	18
2.4. The construction	21
2.5. The trees are also edge-hypomorphic	33
Chapter 3. Non-reconstructible locally finite graphs	35
3.1. Introduction	35
3.2. Sketch of the construction	37
3.3. Closure with respect to promises	41
3.4. Thickening the graph	44
3.5. The construction	46
3.6. A non-reconstructible graph with countably many ends	62
Chapter 4. Topological ubiquity of trees	65
4.1. Introduction	65
4.2. Preliminaries	67
4.3. Well-quasi-orders and κ -embeddability	68
4.4. Linkages between rays	70
4.5. G -tribes and concentration of G -tribes towards an end	74
4.6. Countable subtrees	77
4.7. The induction argument	84
Chapter 5. Ubiquity of graphs with non-linear end structure	89
5.1. Introduction	89
5.2. Preliminaries	93
5.3. The Ray Graph	95
5.4. A pebble-pushing game	96
5.5. Pebbly ends	100

5.6.	<i>G</i> -tribes and concentration of <i>G</i> -tribes towards an end	103
5.7.	Ubiquity of minors of the half grid	106
5.8.	Proof of main results	107
Chapter 6. Ubiquity of locally finite graphs with extensive tree decompositions . .		111
6.1.	Introduction	111
6.2.	Proof sketch	112
6.3.	Preliminaries	119
6.4.	Extensive tree-decompositions and self minors	123
6.5.	Existence of extensive tree-decompositions	126
6.6.	The structure of non-pebbly ends	132
6.7.	Grid-like and half-grid-like ends	140
6.8.	<i>G</i> -tribes and concentration of <i>G</i> -tribes towards an end	148
6.9.	The inductive argument	156
6.10.	Outlook: connections with well-quasi-ordering and better-quasi-ordering .	168
Chapter 7. Minimal obstructions for normal spanning trees		171
7.1.	Introduction	171
7.2.	Collections of infinite subsets of \mathbb{N} , and (\aleph_0, \aleph_1) -graphs	173
7.3.	Finding almost disjoint (\aleph_0, \aleph_1) -subgraphs	176
7.4.	The situation under Martin's Axiom	179
7.5.	A third type of (\aleph_0, \aleph_1) -graph	183
7.6.	More on indivisible graphs	186
Part 2. Topological infinite graph theory		189
Chapter 8. Hamilton decompositions of one-ended Cayley graphs		191
8.1.	Introduction	191
8.2.	Notation and preliminaries	193
8.3.	The covering lemma and a high-level proof of the main theorem	195
8.4.	Proof of the Covering Lemma	196
8.5.	Hamiltonian decompositions of products	208
8.6.	Open Problems	209
Chapter 9. Hamilton cycles in infinite cubic graphs		211
9.1.	Introduction	211
9.2.	Two facts about end degrees	213
9.3.	Affirmative results for second Hamilton cycles	213
9.4.	Examples witnessing optimality	217
Chapter 10. Circuits through prescribed edges		221
10.1.	Introduction	221

10.2. Preliminaries	223
10.3. A reduction to the bridge case	224
10.4. Proving the bridge case	226
10.5. Concluding remarks and an open question	231
Chapter 11. n -Arc connected graphs	235
11.1. Introduction	235
11.2. Preliminaries	237
11.3. Characterizing n -ac Graphs	242
Chapter 12. n -Arc and n -circle connected graph-like spaces	261
12.1. Introduction	261
12.2. Locally finite graphs, and their Freudenthal compactification	264
12.3. General graph-like continua	270
12.4. The graph-like examples	276
Chapter 13. Graph-like compacta: characterizations and Eulerian loops	285
13.1. Introduction	285
13.2. Properties and characterizations of graph-like continua	293
13.3. Eulerian graph-like continua	303
13.4. Bruhn & Stein Parity	307
Chapter 14. Eulerian spaces	315
14.1. Introduction	315
14.2. Eulerian maps and Peano graphs	328
14.3. Approximating by Eulerian decompositions	342
14.4. Product-structured ground spaces	350
14.5. One-dimensional spaces	370
Bibliography	393

CHAPTER 1

Overview

This habilitation thesis represents the outcome of six extensive research projects I have been involved in since taking up my postdoc position at the University of Hamburg in the fall of 2015. The first three of these projects fall into the area of structural infinite graph theory displaying connections to infinite combinatorics and logic. The other three fall into the area of infinite topological graph theory with various connections to topology and geometry. The results presented here are based on thirteen papers [31, 32, 33, 35, 36, 37, 66, 70, 76, 77, 78, 101, 134] conceived and written together with varying groups of collaborators.

The first part of this thesis on structural infinite graph theory contains results and solutions to the following three problems:

The reconstruction problem for locally finite graphs. (Chapters 2 & 3). The reconstruction problem asks whether local information about a graph determines its global isomorphism type. In these two chapters we construct examples showing that, perhaps surprisingly, this question for infinite, locally finite graphs is to be answered in the negative. These results summarily solve a group of seven problems in the literature, which were open for more than 50 years. This construction is joint work with N. Bowler, J. Erde, P. Heinig and F. Lehner [36, 35], published in the *Bulletin of the London Mathematical Society* and *Journal of Combinatorial Theory Series B*.

The main idea of our construction relies on the same principle that lies behind the proof of the completeness theorem for first-order logic: In each step, we can arrange for one additional requirement – while paying the price of simultaneously introducing countably many new future tasks still to be solved. However, by a suitable book-keeping procedure, one can satisfy all requirements in an ω -length recursion.

The ubiquity problem. (Chapters 4, 5 & 6). A graph G is called *ubiquitous* if whenever some graph Γ contains arbitrarily many disjoint copies of G , then Γ must contain a family of infinitely many disjoint copies of G . The *Ubiquity Conjecture*, due to Andreae, suggests that every locally finite connected graph is ubiquitous. While proven for certain well-behaved classes of graphs, the general case remains a challenging open problem with connections to the theory of well- and better-quasi orderings of graphs.

Our results include the solution of the ubiquity problem for trees in Chapter 4, solving a 40-year-old problem raised by Halin and Andreae, as well as the solution of the ubiquity

problem for locally finite graphs of finite tree-width in Chapter 6. This is joint work with N. Bowler, C. Elbracht, J. Erde, P. Gollin, K. Heuer and M. Teegen [31, 32, 33], submitted for publication.

Forbidden minors for normal spanning trees. (Chapter 7). This chapter contains the solution of two 15-year-old problems due to Diestel and Leader about the existence of *normal spanning trees* in infinite graphs. This is joint work with N. Bowler and S. Geschke [37], published in *Fundamenta Mathematicae*.

Normal spanning trees are amongst the most useful objects governing the structure of finite and infinite connected graphs. While every countable connected graph has a normal spanning tree, not all uncountable graphs do, and a challenging problem is to characterise the existence of normal spanning trees by a small list of forbidden minors. Our main result is that two natural questions by Diestel and Leader about these forbidden minors are independent of the usual ZFC-axioms of set theory: For one direction we prove a strong structural result under Martin's axiom, and for the other we present a construction under CH that takes advantage of an Aronszajn tree.

The second part of this thesis on infinite topological graph theory contains the following Hamiltonicity and Eulerianity results:

Hamiltonicity results for infinite graphs with ends. (Chapters 8 & 9). We present an affirmative solution to Alspach's problem about decompositions of infinite Cayley graphs into Hamiltonian double rays, for a large class of abelian groups: Every Cayley graph of a one-ended abelian group generated by a finite set of non-torsion elements has such a Hamilton decomposition. This is joint work with J. Erde und F. Lehner [66], to appear in the *Journal of Combinatorial Theory Series B*.

We then investigate a problem by Mohar whether there exist infinite cubic graphs that are uniquely Hamiltonian (by a theorem of Thomason, every finite Hamiltonian cubic graph contains at least three distinct Hamilton cycles). We show that in the one-ended case, there always exists a second Hamilton cycle (and construct an example showing that there might not be a third), while constructing examples of uniquely Hamiltonian cubic graphs as soon as there are at least two ends. This paper is single authored [134] and published in the *Electronic Journal of Combinatorics*.

Circuits, paths and cycles containing prescribed edges and points. (Chapters 10, 11 & 12). First for a given $n \in \mathbb{N}$, we characterise the finite graphs in which any n edges lie on a common circuit (a closed walk that repeats no edges). This is joint work with P. Knappe, [101], submitted for publication. Next, we characterise for which graphs any n topological points (i.e. vertices or interior points of edges) lie on a common topological path or cycle, respectively. Finally, we extend this characterisation to locally finite graphs

with ends, and describe the jump in complexity for this problem when considering graph-like continua. The finite case is joint work with P. Gartside and A. Mamatelashvili, [76], submitted for publication, and the infinite case is joint work with P. Gartside [78], published in *Topology and its Applications*.

The Eulerian problem for topological spaces. (Chapters 13 & 14). We first generalise Diestel and Kühn’s theory of topological Euler tours from locally finite graphs with ends to graph-like continua. This is joint work with B. Espinoza and P. Gartside [70], submitted for publication. The abstract viewpoint of graph-like spaces allows us to simplify also other results by Georgakopoulos [79], Bruhn & Stein [39], and Berger & Bruhn [20] about locally finite graphs with ends.

Finally, we generalise the theory of topological Euler tours further to geometric structures such as hyperbolic graphs with Gromov boundary, and finally to arbitrary topological spaces. This is joint work with P. Gartside [77], submitted for publication. In the process we uncover a connection to a 90-year-old problem in topology: characterising the irreducible images of the circle. Our first main result is that these irreducible images of the circle are precisely the Eulerian spaces, the proof of which relies on a Baire Category function space argument. This new viewpoint makes it possible to apply combinatorial tools to study such topological spaces. In particular, we define a natural notion of *edge-cuts* in topological spaces and conjecture that a space is Eulerian if and only if it is a Peano continuum where all edge-cuts have even size. As our second main result, we confirm this conjecture for a variety of topological spaces, in particular for all one-dimensional ones. This subsumes and extends all known results about the Eulerianity of infinite graphs and continua to date.

This paper marks a new point in the combinatorial theory of topological graphs with ends, in the sense that its cycle space theory, topological spanning trees, and fundamental cycles and -cuts are here for the first time applied to help solving a natural, longstanding open problem arising outside of combinatorics.

Acknowledgements. I would like to thank my collaborators in the above projects, especially Nathan Bowler, Joshua Erde, Paul Gartside and Florian Lehner: Working together on these problems has been inspiring, exciting and rewarding. It has also been a joy to work in Hamburg during the past four years, and I would like to thank Reinhard Diestel for his guidance and inspiration about all things related to infinite graphs.

Part 1

Structural infinite graph theory

CHAPTER 2

A counterexample to the reconstruction conjecture for locally finite trees

Two graphs G and H are *hypomorphic* if there exists a bijection $\varphi: V(G) \rightarrow V(H)$ such that $G - v \cong H - \varphi(v)$ for each $v \in V(G)$. A graph G is *reconstructible* if $H \cong G$ for all H hypomorphic to G .

It is well known that not all infinite graphs are reconstructible. However, the Harary-Schwenk-Scott Conjecture from 1972 suggests that all locally finite trees are reconstructible.

In this paper, we construct a counterexample to the Harary-Schwenk-Scott Conjecture. Our example also answers four other questions of Nash-Williams, Halin and Andreae on the reconstruction of infinite graphs.

2.1. Introduction

We say that two graphs G and H are (*vertex-*)hypomorphic if there exists a bijection φ between the vertices of G and H such that the induced subgraphs $G - v$ and $H - \varphi(v)$ are isomorphic for each vertex v of G . Any such bijection is called a *hypomorphism*. We say that a graph G is *reconstructible* if $H \cong G$ for every H hypomorphic to G . The following conjecture, attributed to Kelly and Ulam, is perhaps one of the most famous unsolved problems in the theory of graphs.

CONJECTURE 2.1.1 (The Reconstruction Conjecture). *Every finite graph with at least three vertices is reconstructible.*

For an overview of results towards the Reconstruction Conjecture for finite graphs see the survey of Bondy and Hemminger [29]. Harary [90] proposed the Reconstruction Conjecture for infinite graphs, however Fisher [73] found a counterexample, which was simplified to the following counterexample by Fisher, Graham and Harary [74]: consider the infinite tree G in which every vertex has countably infinite degree, and the graph H formed by taking two disjoint copies of G , which we will write as $G \sqcup G$. For each vertex v of G , the induced subgraph $G - v$ is isomorphic to $G \sqcup G \sqcup \dots$, a disjoint union of countably many copies of G , and similarly for each vertex w of H , the induced subgraph $H - w$ is isomorphic to $G \sqcup G \sqcup \dots$ as well. Therefore, any bijection from $V(G)$ to $V(H)$ is a hypomorphism, but G and H are clearly not isomorphic. Hence, the tree G is not reconstructible.

These examples, however, contain vertices of infinite degree. Regarding locally finite graphs, Harary, Schwenk and Scott [91] showed that there exists a non-reconstructible locally finite forest. However, they conjectured that the Reconstruction Conjecture should hold for locally finite trees.

CONJECTURE 2.1.2 (The Harary-Schwenk-Scott Conjecture). *Every locally finite tree is reconstructible.*

This conjecture has been verified in a number of special cases. Kelly [100] showed that finite trees on at least three vertices are reconstructible. Bondy and Hemminger [28] showed that every tree with at least two but a finite number of ends is reconstructible, and Thomassen [151] showed that this also holds for one-ended trees. Andreae [12] proved that also every tree with countably many ends is reconstructible.

A survey of Nash-Williams [127] on the subject of reconstruction problems in infinite graphs gave the following three main open problems in this area, which have remained open until now.

PROBLEM 2.1.3 (Nash-Williams). Is every locally finite connected infinite graph reconstructible?

PROBLEM 2.1.4 (Nash-Williams). If two infinite trees are hypomorphic, are they also isomorphic?

PROBLEM 2.1.5 (Halin). If G and H are hypomorphic, do there exist embeddings $G \hookrightarrow H$ and $H \hookrightarrow G$?

Problem 2.1.4 has been emphasized in Andreae's [14], which contains partial affirmative results on Problem 2.1.4. A positive answer to Problem 2.1.3 or 2.1.4 would verify the Harary-Schwenk-Scott Conjecture. In this paper we construct a pair of trees which are not only a counterexample to the Harary-Schwenk-Scott Conjecture, but also answer the three questions of Nash-Williams and Halin in the negative. Our counterexample will in fact have bounded degree.

THEOREM 2.1.6. *There are two (vertex)-hypomorphic infinite trees T and S with maximum degree three such that there is no embedding $T \hookrightarrow S$ or $S \hookrightarrow T$.*

Our example also provides a strong answer to a question by Andreae [13] about edge-reconstructibility. Two graphs G and H are *edge-hypomorphic* if there exists a bijection $\varphi: E(G) \rightarrow E(H)$ such that $G - e \cong H - \varphi(e)$ for each $e \in E(G)$. A graph G is *edge-reconstructible* if $H \cong G$ for all H edge-hypomorphic to G . In [13] Andreae constructed countable forests which are not edge-reconstructible, but conjectured that no locally finite such examples can exist.

PROBLEM 2.1.7 (Andreae). Is every locally finite graph with infinitely many edges edge-reconstructible?

Our example answers Problem 2.1.7 in the negative: the trees T and S we construct for Theorem 2.1.6 will also be edge-hypomorphic. Besides answering Problem 2.1.7, this appears to be the first known example of two non-isomorphic graphs that are *simultaneously* vertex- and edge-hypomorphic.

The Reconstruction Conjecture has also been considered for general locally finite graphs. Nash-Williams [126] showed that any locally finite graph with at least three, but a finite number of ends is reconstructible, and in [128], he established the same result for two-ended graphs. The following problems, also from [127], remain open:

PROBLEM 2.1.8 (Nash-Williams). Is every locally finite graph with exactly one end reconstructible?

PROBLEM 2.1.9 (Nash-Williams). Is every locally finite graph with countably many ends reconstructible?

In a paper in preparation [36], we will extend the methods developed in the present paper to also construct counterexamples to Problems 2.1.8 and 2.1.9.

This paper is organised as follows. In the next section we will give a short, high-level overview of our counterexample to the Harary-Schwenk-Scott Conjecture. In Section 2.3, we will develop the technical tools necessary for our construction, and in Section 2.4, we will prove Theorem 2.1.6.

For standard graph theoretical concepts we follow the notation in [54].

2.2. Sketch of the construction

In this section we sketch the main ideas of the construction. For the sake of simplicity we only indicate how to ensure that the trees T and S are vertex-hypomorphic and non-isomorphic, but not that they are edge-hypomorphic as well, nor that neither embeds into the other.

Our plan is to build the trees T and S recursively, where at each step of the construction we ensure for some vertex v already chosen for T that there is a corresponding vertex w of S with $T - v \cong S - w$, or vice versa. This will ensure that by the end of the construction, the trees we have built are hypomorphic.

More precisely, at step n we will construct subtrees T_n and S_n of our eventual trees, where some of the leaves of these subtrees have been coloured in two colours, say red and blue. We will only further extend the trees from these coloured leaves, and we will extend from leaves of the same colour in the same way.

That is, the plan is that there should be two further rooted trees R and B such that T can be obtained from T_n by attaching copies of R at all red leaves and copies of B at all blue leaves, and S can be obtained from S_n in the same way. At step n , however, we do not yet know what these trees R and B will eventually be.

Nevertheless, we can ensure that the induced subgraphs, $T - v$ and $S - w$, of the vertices we have dealt with so far really will match up. More precisely, by step n we have vertices x_1, \dots, x_n of T_n and y_1, \dots, y_n of S_n for which we intend that $T - x_j$ should be isomorphic to $S - y_j$ for each j . We ensure this by arranging that for each j there is an isomorphism from $T_n - x_j$ to $S_n - y_j$ which preserves the colours of the leaves.

The T_n will be nested, and we will take T to be the union of all of them; similarly the S_n will be nested and we take S to be the union of all of them.

There is a trick to ensure that T and S do not end up being isomorphic. First we ensure, for each n , that there is no isomorphism from T_n to S_n . We also ensure that the part of T or S beyond any coloured leaf of T_n or S_n begins with a long non-branching path (called a *bare* path), longer than any such path appearing in T_n or S_n . Call the length of these long paths k_{n+1} .

Suppose now for a contradiction that there is an isomorphism from T to S . Then there must exist some large n such that the isomorphism sends some vertex t of T_n to a vertex s of S_n . However, T_n is the component of T containing t after all bare paths of length k_{n+1} have been removed¹, and so it must map isomorphically onto the component of S containing s after all bare paths of length k_{n+1} have been removed, namely onto S_n . However, there is no isomorphism from T_n onto S_n , so we have the desired contradiction.



FIGURE 2.1. A first approximation of T_{n+1} on the left, and S_{n+1} on the right. All dotted lines are non-branching paths of length k_{n+1} .

Suppose now that we have already constructed T_n and S_n and wish to construct T_{n+1} and S_{n+1} . Suppose further that we are given a vertex v of T_n for which we wish to find a partner w in S_{n+1} so that $T - v$ and $S - w$ are isomorphic. We begin by building a tree $\hat{T}_n \not\cong T_n$ which has some vertex w such that $T_n - v \cong \hat{T}_n - w$. This can be done by taking the components of $T_n - v$ and arranging them suitably around the new vertex w .

We will take S_{n+1} to include S_n and \hat{T}_n , with the copies of red and blue leaves in \hat{T}_n also coloured red and blue respectively. As indicated on the right in Figure 2.1, we add paths of length k_{n+1} to some blue leaf b of S_n and to some red leaf \hat{r} of \hat{T}_n and join these paths at their other endpoints by some edge e_n . We also join two new leaves y and g to the endvertices of e_n . We colour the leaf y yellow and the leaf g green (to avoid confusion

¹Here and throughout this section we will omit minor technical details for brevity.

with the red and blue leaves from step n , we take the two colours applied to the leaves in step $n + 1$ to be yellow and green).

To ensure that $T_{n+1} - v \cong S_{n+1} - w$, we take T_{n+1} to include T_n together with a copy \hat{S}_n of S_n , coloured appropriately and joined up in the same way, as indicated on the left in Figure 2.1.

The only problem up to this point is that we have not been faithful to our intention of extending in the same way at each red or blue leaf of T_n and S_n . Thus, we now copy the same subgraph appearing beyond r in Fig. 2.1, including its coloured leaves, onto all the other red leaves of S_n and T_n . Similarly we copy the subgraph appearing beyond the blue leaf b of S_n onto all other blue leaves of S_n and T_n .

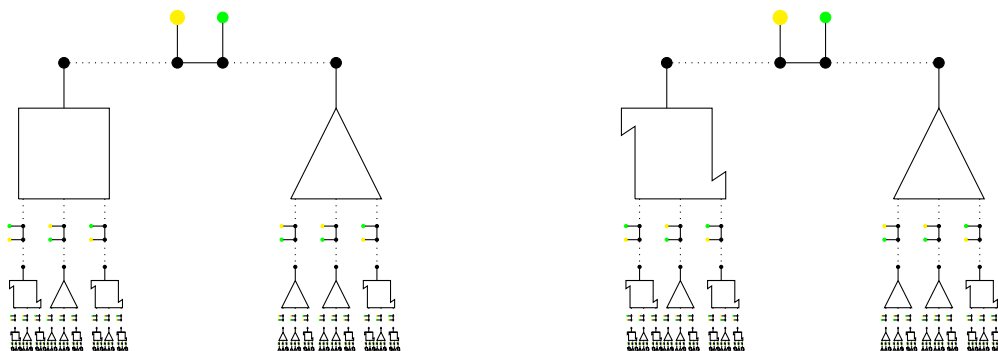


FIGURE 2.2. A sketch of T_{n+1} and S_{n+1} after countably many steps.

At this point, we would have kept our promise of adding the same thing behind every red and blue leaf of T_n and S_n , and hence would have achieved $T_{n+1} - x_j \cong S_{n+1} - y_j$ for all $j \leq n$. However, by gluing the additional copies to blue and red leaves of T_n and S_n , we now have ruined the isomorphism between $T_{n+1} - v$ and $S_{n+1} - w$. In order to repair this, we also have to copy the graphs appearing beyond r and b in Fig. 2.1 respectively onto all red and blue leaves of \hat{S}_n and \hat{T}_n . This repairs $T_{n+1} - v \cong S_{n+1} - w$, but again violates our initial promises. In this way, we keep adding, step by step, further copies of the graphs appearing beyond r and b in Fig. 2.1 respectively onto all red and blue leaves of everything we have constructed so far.

At every step we preserved the colours of leaves in all newly added copies, so we get new red leaves and blue leaves, and we continue the process of copying onto those new leaves as well. After countably many steps we have dealt with all red or blue leaves. We take these new trees to be S_{n+1} and T_{n+1} . They are non-isomorphic, since after removing all long bare paths, T_{n+1} contains T_n as a component, whereas S_{n+1} does not.

Figure 2.2 shows how T_{n+1} and S_{n+1} might appear. We have now fulfilled our intention of sticking the same thing onto all red leaves and the same thing onto all blue leaves, but we have also ensured that $T_{n+1} - v \cong S_{n+1} - w$, as desired.

2.3. Closure with respect to promises

In this section, we formalise the ideas set forth in the proof sketch of how to extend a graph so that it looks the same beyond certain sets of leaves.

Given a directed edge $\vec{e} = x\vec{y}$ in some forest $G = (V, E)$, we denote by $G(\vec{e})$ the unique component of $G - e$ containing the vertex y . We think of $G(\vec{e})$ as a rooted tree with root y . As indicated in the previous section, in order to make T and S hypomorphic at the end, we will often have to guarantee $S(\vec{e}) \cong T(\vec{f})$ for certain pairs of edges \vec{e} and \vec{f} .

DEFINITION 2.3.1 (Promise structure). A *promise structure* $\mathcal{P} = (G, \vec{P}, \mathcal{L})$ consists of:

- a forest G ,
- $\vec{P} = \{\vec{p}_i : i \in I\}$ a set of directed edges $\vec{P} \subseteq \vec{E}(G)$, and
- $\mathcal{L} = \{L_i : i \in I\}$ a set of pairwise disjoint sets of leaves of G .

Often, when the context is clear, we will not make a distinction between \mathcal{L} and the set $\bigcup_i L_i$, for notational convenience.

We will call an edge $\vec{p}_i \in \vec{P}$ a *promise edge*, and leaves $\ell \in L_i$ *promise leaves*. A promise edge $\vec{p}_i \in \vec{P}$ is called a *placeholder-promise* if the component $G(\vec{p}_i)$ consists of a single leaf $c_i \in L_i$, then called a *placeholder-leaf*. We write

$$\mathcal{L}_p = \{L_i : i \in I, \vec{p}_i \text{ a placeholder-promise}\} \text{ and } \mathcal{L}_q = \mathcal{L} \setminus \mathcal{L}_p.$$

Given a leaf ℓ in G , there is a unique edge $q_\ell \in E(G)$ incident with ℓ , and this edge has a natural orientation \vec{q}_ℓ towards ℓ . Informally, we think of the ‘promise’ $\ell \in L_i$ as saying that if we extend G to a graph $H \supset G$, we will do so in such a way that $H(\vec{q}_\ell) \cong H(\vec{p}_i)$. Given a promise structure $\mathcal{P} = (G, \vec{P}, \mathcal{L})$, we would like to construct a graph $H \supset G$ which satisfies all the promises in \mathcal{P} . This will be done by the following kind of extension.

DEFINITION 2.3.2 (Leaf extension). Given an inclusion $H \supseteq G$ of forests and a set L of leaves of G , H is called a *leaf extension*, or more specifically an *L -extension*, of G , if:

- every component of H contains precisely one component of G , and
- for every vertex $h \in H \setminus G$ and every vertex $g \in G$ in the same component as h , the unique $g - h$ path in H meets L .

In the remainder of this section we describe a construction of a forest $\text{cl}(G)$ which has the following properties.

PROPOSITION 2.3.3. *Let G be a forest and let $(G, \vec{P}, \mathcal{L})$ be a promise structure. Then there is a forest $\text{cl}(G)$ such that:*

- (cl.1) $\text{cl}(G)$ is an \mathcal{L}_q -extension of G , and
- (cl.2) for every $\vec{p}_i \in \vec{P}$ and all $\ell \in L_i$,

$$\text{cl}(G)(\vec{p}_i) \cong \text{cl}(G)(\vec{q}_\ell)$$

are isomorphic as rooted trees.

We first describe the construction of $\text{cl}(G)$, and then verify the properties asserted in Proposition 2.3.3. Let us define a sequence of promise structures $(H^{(i)}, \vec{P}, \mathcal{L}^{(i)})$ as follows. We set $(H^{(0)}, \vec{P}, \mathcal{L}^{(0)}) = (G, \vec{P}, \mathcal{L})$. We construct a sequence of graphs

$$G = H^{(0)} \subseteq H^{(1)} \subseteq H^{(2)} \subseteq \dots,$$

and each $H^{(n)}$ will get a promise structure whose set of promise edges is equal to \vec{P} again, yet whose set of promise leaves depends on n as follows: given $(H^{(n)}, \vec{P}, \mathcal{L}^{(n)})$, we construct $H^{(n+1)}$ by gluing, for each i , at every promise leaf $\ell \in L_i^{(n)}$ a rooted copy of $G(\vec{p}_i)$. As promise leaves for $H^{(n+1)}$ we take all promise leaves from the newly added copies of $G(\vec{p}_i)$. That is, if a leaf $\ell \in G(\vec{p}_i)$ was such that $\ell \in L_j$, then every copy of that leaf will be in $L_j^{(n+1)}$.

Formally, suppose that $(G, \vec{P}, \mathcal{L})$ is a promise structure. For each $\vec{p}_i \in \vec{P}$ let $C_i = G(\vec{p}_i)$ and let c_i be the root of this tree. If U is a set and H is a graph, then we denote by $U \times H$ the graph whose vertices are pairs (u, v) with $u \in U$ and v a vertex of H , and with an edge from (u, v) to (u, w) whenever vw is an edge of H . Let $(H^{(0)}, \vec{P}, \mathcal{L}^{(0)}) = (G, \vec{P}, \mathcal{L})$ and given $(H^{(n)}, \vec{P}, \mathcal{L}^{(n)})$ let us define:

- $H^{(n+1)}$ to be the quotient of $H^{(n)} \sqcup \bigsqcup_{i \in I} (L_i^{(n)} \times C_i)$ w.r.t. the relation

$$l \sim (l, c_i) \text{ for } l \in L_i^{(n)} \in \mathcal{L}^{(n)}.$$

- $\mathcal{L}^{(n+1)} = \{L_i^{(n+1)} : i \in I\}$ with $L_i^{(n+1)} = \bigcup_{j \in I} L_j^{(n)} \times (C_j \cap L_i)$.

There is a sequence of natural inclusions $G = H^{(0)} \subseteq H^{(1)} \subseteq \dots$ and we define $\text{cl}(G)$ to be the direct limit of this sequence.

DEFINITION 2.3.4 (Promise-respecting map). Let G be a forest, $F^{(1)}$ and $F^{(2)}$ be leaf extensions of G , and $\mathcal{P}^{(1)} = (F^{(1)}, \vec{P}, \mathcal{L}^{(1)})$ and $\mathcal{P}^{(2)} = (F^{(2)}, \vec{P}, \mathcal{L}^{(2)})$ be promise structures with $\vec{P} \subseteq \vec{E}(G)$. Suppose $X^{(1)} \subseteq V(F^{(1)})$ and $X^{(2)} \subseteq V(F^{(2)})$.

A bijection $\varphi: X^{(1)} \rightarrow X^{(2)}$ is \vec{P} -respecting (with respect to $\mathcal{P}^{(1)}$ and $\mathcal{P}^{(2)}$) if the image of $L_i^{(1)} \cap X^{(1)}$ under φ is $L_i^{(2)} \cap X^{(2)}$ for all i .

Since both promise structures $\mathcal{P}^{(1)}$ and $\mathcal{P}^{(2)}$ refer to the same edge set \vec{P} , we can think of them as defining a $|\vec{P}|$ -colouring on some sets of leaves. Then a mapping is \vec{P} -respecting if it preserves leaf colours.

LEMMA 2.3.5. Let $(G, \vec{P}, \mathcal{L})$ be a promise structure and let $G = H^{(0)} \subseteq H^{(1)} \subseteq \dots$ be as defined above. Then the following statements hold:

- $H^{(n)}$ is an \mathcal{L}_q -extension of G for all n ,
- $\Delta(H^{(n+1)}) = \Delta(H^{(n)})$ for all n , and

- For each $\ell \in L_i \in \mathcal{L}$ there exists a sequence of \vec{P} -respecting rooted isomorphisms $\varphi_{\ell,n}: H^{(n)}(\vec{p}_i) \rightarrow H^{(n+1)}(\vec{q}_\ell)$ such that $\varphi_{\ell,n+1}$ extends $\varphi_{\ell,n}$ for all $n \in \mathbb{N}$.

PROOF. The first two statements are clear. We will prove the third by induction on n . To construct $H^{(1)}$ from G , we glued a rooted copy of $G(\vec{p}_i)$ to each $\ell \in L_i$, keeping all copies of promise leaves. Hence, for any given $\ell \in L_i$, the natural isomorphism $\varphi_{\ell,0}: G(\vec{p}_i) \rightarrow H^{(1)}(\vec{q}_\ell)$ is \vec{P} -respecting as desired.

Now suppose that $\varphi_{\ell,n}$ exists for all $\ell \in \mathcal{L}$. To form $H^{(n+1)}(\vec{p}_i)$, we glued on a copy of $G(\vec{p}_i)$ to each $\ell \in L_i^{(n)} \cap H^{(n)}(\vec{p}_i)$, and to construct $H^{(n+2)}(\vec{q}_\ell)$, we glued on a copy of $G(\vec{p}_i)$ to each $\ell \in L_i^{(n+1)} \cap H^{(n+1)}(\vec{q}_\ell)$, in both cases keeping all copies of promise leaves.

Therefore, since $\varphi_{\ell,n}$ was a \vec{P} -respecting rooted isomorphism from $H^{(n)}(\vec{p}_i)$ to $H^{(n+1)}(\vec{q}_\ell)$, we can combine the individual isomorphisms between the newly added copies of $G(\vec{p}_i)$ with $\varphi_{\ell,n}$ to form $\varphi_{\ell,n+1}$. \square

We can now complete the proof of Proposition 2.3.3.

PROOF OF PROPOSITION 2.3.3. First, we note that $G \subseteq \text{cl}(G)$, and since each $H^{(n)}$ is an \mathcal{L}_q -extension of G for all n , so is $\text{cl}(G)$. Also, since each $H^{(n)}$ is a forest it follows that $\text{cl}(G)$ is a forest.

Let us show that $\text{cl}(G)$ satisfies property (cl.2). Since we have the sequence of inclusions $G = H^{(0)} \subseteq H^{(1)} \subseteq \dots$, it follows that $\text{cl}(G)(\vec{q}_\ell)$ is the direct limit of the sequence $H^{(0)}(\vec{q}_\ell) \subseteq H^{(1)}(\vec{q}_\ell) \subseteq \dots$ and also $\text{cl}(G)(\vec{p}_i)$ is the direct limit of the sequence $H^{(0)}(\vec{p}_i) \subseteq H^{(1)}(\vec{p}_i) \subseteq \dots$. By Lemma 2.3.5 there is a sequence of rooted isomorphisms $\varphi_{\ell,n}: H^{(n)}(\vec{p}_i) \rightarrow H^{(n+1)}(\vec{q}_\ell)$ such that $\varphi_{\ell,n+1}$ extends $\varphi_{\ell,n}$, so $\varphi_\ell = \bigcup_n \varphi_{\ell,n}$ is the required isomorphism. \square

We remark that it is possible to show that $\text{cl}(G)$ is in fact determined, uniquely up to isomorphism, by the properties (cl.1) and (cl.2). Also we note that since each $H^{(n)}$ has the same maximum degree as G , it follows that $\Delta(\text{cl}(G)) = \Delta(G)$.

There is a natural promise structure on $\text{cl}(G)$ given by the placeholder promises in \vec{P} and their corresponding promise leaves. In the construction sketch from Section 2.2, these leaves corresponded to the yellow and green leaves. We now show how to keep track of the placeholder promises when taking the closure of a promise structure.

Note that if \vec{p}_i is a placeholder promise, then for each $(H^{(n)}, \mathcal{P}, \mathcal{L}^{(n)})$ we have $L_i^{(n)} \supseteq L_i^{(n-1)}$. Indeed, for each leaf in $L_i^{(n-1)}$ we glue a copy of the component c_i together with the associated promises on the leaves in this component. However, c_i is just a single vertex, with a promise corresponding to \vec{p}_i , and hence $L_i^{(n)} \supseteq L_i^{(n-1)}$. For every placeholder promise $\vec{p}_i \in \vec{P}$ we define $\text{cl}(L_i) = \bigcup_n L_i^{(n)}$.

DEFINITION 2.3.6 (Closure of a promise structure). The *closure* of the promise structure $(G, \mathcal{P}, \mathcal{L})$ is the promise structure $\text{cl}(\mathcal{P}) = \left(\text{cl}(G), \text{cl}(\vec{P}), \text{cl}(\mathcal{L}) \right)$, where:

- $\text{cl}(\vec{P}) = \left\{ \vec{p}_i : \vec{p}_i \in \vec{P} \text{ is a placeholder-promise} \right\}$, and
- $\text{cl}(\mathcal{L}) = \{ \text{cl}(L_i) : \vec{p}_i \in \vec{P} \text{ is a placeholder-promise} \}$.

We note that, since each isomorphism $\varphi_{\ell,n}$ from Lemma 2.3.5 was \vec{P} -respecting, it is possible to strengthen Proposition 2.3.3 in the following way.

PROPOSITION 2.3.7. *Let G be a forest and let $(G, \vec{P}, \mathcal{L})$ be a promise structure. Then the forest $\text{cl}(G)$ satisfies:*

(cl.3) *for every $\vec{p}_i \in \vec{P}$ and every $\ell \in L_i$,*

$$\text{cl}(G)(\vec{p}_i) \cong \text{cl}(G)(\vec{q}_\ell)$$

are isomorphic as rooted trees, and this isomorphism is $\text{cl}(\vec{P})$ -respecting with respect to $\text{cl}(\mathcal{P})$.

PROOF. Since each isomorphism $\varphi_{\ell,n} : H^{(n)}(\vec{p}_i) \rightarrow H^{(n+1)}(\vec{q}_\ell)$ in Proposition 2.3.5 is \vec{P} -respecting, we have

$$\varphi_{\ell,n} \left(L_i^{(n)} \cap H^{(n)}(\vec{p}_i) \right) = L_i^{(n+1)} \cap H^{(n+1)}(\vec{q}_\ell).$$

For each placeholder promise we have that $\text{cl}(L_i) = \bigcup_n L_i^{(n)}$, and so it follows that

$$\text{cl}(L_i) \cap \text{cl}(G)(\vec{q}_\ell) = \bigcup_n \left(L_i^{(n)} \cap H^{(n)}(\vec{q}_\ell) \right)$$

and

$$\text{cl}(L_i) \cap \text{cl}(G)(\vec{p}_i) = \bigcup_n \left(L_i^{(n)} \cap H^{(n)}(\vec{p}_i) \right).$$

From this it follows that $\varphi_\ell = \bigcup_n \varphi_{\ell,n}$ is a $\text{cl}(\vec{P})$ -respecting isomorphism between $\text{cl}(G)(\vec{p}_i)$ and $\text{cl}(G)(\vec{q}_\ell)$ as rooted trees. \square

It is precisely this property (cl.3) of the promise closure that will allow us, in Claim 8 below, to maintain partial hypomorphisms during our recursive construction.

2.4. The construction

In this section we construct two hypomorphic locally finite trees neither of which embed into the other, establishing our main theorem announced in the introduction.


2.4.1. Preliminary definitions.

DEFINITION 2.4.1 (Bare path). A path $P = v_0, v_1, \dots, v_n$ in a graph G is called a *bare path* if $\deg_G(v_i) = 2$ for all internal vertices v_i for $0 < i < n$. The path P is a *maximal bare path* (or *maximally bare*) if in addition $\deg_G(v_0) \neq 2 \neq \deg_G(v_n)$. An infinite path $P = v_0, v_1, v_2, \dots$ is *maximally bare* if $\deg_G(v_0) \neq 2$ and $\deg_G(v_i) = 2$ for all $i \geq 1$.


LEMMA 2.4.2. *Let T be a tree and $e \in E(T)$. If every maximal bare path in T has length at most $k \in \mathbb{N}$, then every maximal bare path in $T - e$ has length at most $2k$.*

PROOF. We first note that every maximal bare path in $T - e$ has finite length, since any infinite bare path in $T_n - e$ would contain a subpath which is an infinite bare path in T . If $P = \{x_0, x_1, \dots, x_n\}$ is a maximal bare path in $T - e$ which is not a subpath of any maximal bare path in T , then there is at least one $1 \leq i \leq n - 1$ such that e is adjacent to x_i , and since T was a tree, x_i is unique. Therefore, both $\{x_0, x_1, \dots, x_i\}$ and $\{x_i, x_{i+1}, \dots, x_n\}$ are maximal bare paths in T . By assumption both i and $n - i$ are at most k , and so the length of P is at most $2k$, as claimed. \square

DEFINITION 2.4.3 (Bare extension). Given a forest G , a subset B of leaves of G , and a component T of G , we say that a tree $\hat{T} \supset T$ is a *bare extension* of T at B to length k if \hat{T} can be obtained from T by adjoining, at each vertex $l \in B \cap V(T)$, a new path of length k starting at l and a new leaf whose only neighbour is l .



A tree T with designated leaf set B .



A bare extension of T at B .

FIGURE 2.3. Building a bare extension of a tree T at B to length k . All dotted lines are maximal bare paths of length k .

Note that the new leaves attached to each $l \in B$ ensure that the paths of length k are indeed maximal bare paths.

DEFINITION 2.4.4 (k -ball). For G a subgraph of H , the k -ball $\text{Ball}_H(G, k)$ is the induced subgraph of H on the set of vertices within distance k of some vertex of G .

DEFINITION 2.4.5 (Binary tree). For $k \geq 1$, the *binary tree of height k* is the unique rooted tree on $2^k - 1 = 1 + 2 + \dots + 2^{k-1}$ vertices such that the root has degree 2, there are 2^{k-1} leaves, and all other vertices have degree 3. By a *binary tree* we mean a binary tree of height k for some $k \in \mathbb{N}$.

2.4.2. The back-and-forth construction. We prove the following theorem.

THEOREM 2.4.6. *There are two (vertex-)hypomorphic infinite trees T and S with maximum degree 3 such that there is no embedding $T \hookrightarrow S$ or $S \hookrightarrow T$.*




FIGURE 2.4. The binary tree of height 3.

To do this we shall recursively construct, for each $n \in \mathbb{N}$,

- disjoint (possibly infinite) rooted trees T_n and S_n ,
- disjoint (possibly infinite) sets R_n and B_n of leaves of the forest $T_n \sqcup S_n$,
- finite sets $X_n \subseteq V(T_n)$ and $Y_n \subseteq V(S_n)$, and bijections $\varphi_n: X_n \rightarrow Y_n$,
- a family of isomorphisms $\mathcal{H}_n = \{h_{n,x}: T_n - x \rightarrow S_n - \varphi_n(x): x \in X_n\}$,
- strictly increasing sequences of integers $k_n \geq 2$ and $b_n \geq 3$,

such that (letting all objects indexed by -1 be the empty set) for all $n \in \mathbb{N}$:

- (†1) $T_{n-1} \subseteq T_n$ and $S_{n-1} \subseteq S_n$ as induced subgraphs,
- (†2) the vertices of T_n and S_n all have degree at most 3,
- (†3) the root of T_n is in R_n and the root of S_n is in B_n ,
- (†4) all binary trees appearing as subgraphs of $T_n \sqcup S_n$ are finite and have height at most b_n ,
- (†5) all bare paths in $T_n \sqcup S_n$ are finite and have length at most k_n ,
- (†6) $\text{Ball}_{T_n}(T_{n-1}, k_{n-1} + 1)$ is a bare extension of T_{n-1} at $R_{n-1} \cup B_{n-1}$ to length $k_{n-1} + 1$ and does not meet $R_n \cup B_n$,
- (†7) $\text{Ball}_{S_n}(S_{n-1}, k_{n-1} + 1)$ is a bare extension of S_{n-1} at $R_{n-1} \cup B_{n-1}$ to length $k_{n-1} + 1$ and does not meet $R_n \cup B_n$,
- (†8) there is no embedding from T_n into any bare extension of S_n at $R_n \cup B_n$ to any length, nor from S_n into any bare extension of T_n at $R_n \cup B_n$ to any length,
- (†9) any embedding of T_n into a bare extension of T_n at $R_n \cup B_n$ to any length fixes the root of T_n and has image T_n ,
- (†10) any embedding of S_n into a bare extension of S_n at $R_n \cup B_n$ to any length fixes the root of S_n and has image S_n ,
- (†11) there are enumerations $V(T_n) = \{t_j: j \in J_n\}$ and $V(S_n) = \{s_j: j \in J_n\}$ such that
 - $J_{n-1} \subseteq J_n \subseteq \mathbb{N}$,
 - $\{t_j: j \in J_n\}$ extends the enumeration $\{t_j: j \in J_{n-1}\}$ of $V(T_{n-1})$, and similarly for $\{s_j: j \in J_n\}$,
 - $|\mathbb{N} \setminus J_n| = \infty$,
 - $\{0, 1, \dots, n\} \subseteq J_n$,
- (†12) $\{t_j, s_j: j \leq n\} \cap (R_n \cup B_n) = \emptyset$,
- (†13) the finite sets of vertices X_n and Y_n satisfy $|X_n| = n = |Y_n|$, and
 - $X_{n-1} \subseteq X_n$ and $Y_{n-1} \subseteq Y_n$,
 - $\varphi_n \upharpoonright X_{n-1} = \varphi_{n-1}$,
 - $\{t_j: j \leq n\} \subseteq X_{2n+1}$ and $\{s_j: j \leq n\} \subseteq Y_{2(n+1)}$,

- $(X_n \cup Y_n) \cap (R_n \cup B_n) = \emptyset$,
- (†14) the families of isomorphisms \mathcal{H}_n satisfy
- $h_{n,x} \upharpoonright (T_{n-1} - x) = h_{n-1,x}$ for all $x \in X_{n-1}$,
 - the image of $R_n \cap V(T_n)$ under $h_{n,x}$ is $R_n \cap V(S_n)$, and
 - the image of $B_n \cap V(T_n)$ under $h_{n,x}$ is $B_n \cap V(S_n)$ for all $x \in X_n$.

2.4.3. The construction yields the desired non-reconstructible trees. By property (†1), we have $T_0 \subseteq T_1 \subseteq T_2 \subseteq \dots$ and $S_0 \subseteq S_1 \subseteq S_2 \subseteq \dots$. Let T and S be the union of the respective chains. It is clear that T and S are trees, and that as a consequence of (†2), both trees have maximum degree 3.

We claim that the map $\varphi = \bigcup_n \varphi_n$ is a hypomorphism between T and S . Indeed, it follows from (†11) and (†13) that φ is a well-defined bijection from $V(T)$ to $V(S)$. To see that φ is a hypomorphism, consider any vertex x of T . This vertex appears as some t_j in our enumeration of $V(T)$, so by (†14) the map

$$h_x := \bigcup_{n>2j} h_{n,x}: T - x \rightarrow S - \varphi(x)$$

is an isomorphism between $T - x$ and $S - \varphi(x)$.

Now suppose for a contradiction that $f: T \hookrightarrow S$ is an embedding of T into S . Then $f(t_0)$ is mapped into S_n for some $n \in \mathbb{N}$. Properties (†5) and (†6) imply that after deleting all maximal bare paths in T of length $> k_n$, the connected component of t_0 is a bare extension of T_n to length 0. Further, by (†7), $\text{Ball}_S(S_n, k_n + 1)$ is a bare extension of S_n at $R_n \cup B_n$ to length $k_n + 1$. But combining the fact that $f(T_n) \cap S_n \neq \emptyset$ and the fact that T_n does not contain long maximal bare paths, it is easily seen that $f(T_n) \subseteq \text{Ball}_S(S_n, k_n + 1)$, contradicting (†8).²

The case $S \hookrightarrow T$ yields a contradiction in a symmetric fashion, completing the proof.

2.4.4. The base case: there are finite rooted trees T_0 and S_0 satisfying requirements (†1)–(†14). Choose a pair of non-isomorphic, equally sized trees T_0 and S_0 of maximum degree 3, and pick a leaf each as roots $r(T_0)$ and $r(S_0)$ for T_0 and S_0 , subject to conditions (†8)–(†10) with $R_0 = \{r(T_0)\}$ and $B_0 = \{r(S_0)\}$. A possible choice is given in Fig. 2.5. Here, (†8) is satisfied, because any embedding of T_0 into a bare extension of S_0 has to map the binary tree of height 3 in T_0 to the binary tree in S_0 , making it impossible to embed the middle leaf. Properties (†9) and (†10) are similar.

Let $J_0 = \{0, 1, \dots, |T_0| - 1\}$ and choose enumerations $V(T_0) = \{t_j: j \in J_0\}$ and $V(S_0) = \{s_j: j \in J_0\}$ with $t_0 \neq r(T_0)$ and $s_0 \neq r(S_0)$. This takes care of (†11) and (†12). Finally, (†13) and (†14) are satisfied for $X_0 = Y_0 = \mathcal{H}_0 = \varphi_0 = \emptyset$. Set $k_0 = 2$ and $b_0 = 3$.

²To get the non-embedding property, we have used (†5)–(†8) at every step n . While at the first glance, properties (†4), (†9) and (†10) do not seem to be needed at this point, they are crucial during the construction to establish (†8) at step $n + 1$. See Claim 5 below for details.

FIGURE 2.5. A possible choice for finite rooted trees T_0 and S_0 .

2.4.5. The inductive step: set-up. Now, assume that we have constructed trees T_k and S_k for all $k \leq n$ such that (†1)–(†14) are satisfied up to n . If $n = 2m$ is even, then we have $\{t_j : j \leq m - 1\} \subseteq X_n$, so in order to satisfy (†13) we have to construct T_{n+1} and S_{n+1} such that the vertex t_m is taken care of in our partial hypomorphism. Similarly, if $n = 2m + 1$ is odd, then we have $\{s_j : j \leq m - 1\} \subseteq Y_n$ and we have to construct T_{n+1} and S_{n+1} such that the vertex s_m is taken care of in our partial hypomorphism. Both cases are symmetric, so let us assume in the following that $n = 2m$ is even.

Now let v be the vertex with the least index in the set $\{t_j : j \in J_n\} \setminus X_n$, i.e.

$$(1) \quad v = t_i \text{ for } i = \min \{\ell : t_\ell \in V(T_n) \setminus X_n\}.$$

Then by assumption (†13), v will be t_m , unless t_m was already in X_n anyway. In any case, since $|X_n| = |Y_n| = n$, it follows from (†11) that $i \leq n$, so by (†12), v does not lie in our leaf sets $R_n \cup B_n$, i.e.

$$(2) \quad v \notin R_n \cup B_n.$$

In the next sections, we will demonstrate how to obtain trees $T_{n+1} \supset T_n$ and $S_{n+1} \supset S_n$ with $X_{n+1} = X_n \cup \{v\}$ and $Y_{n+1} = Y_n \cup \{\varphi_{n+1}(v)\}$ satisfying (†1)–(†10) and (†13)–(†14).

After we have completed this step, since $|\mathbb{N} \setminus J_n| = \infty$, it is clear that we can extend our enumerations of T_n and S_n to enumerations of T_{n+1} and S_{n+1} as required, making sure to first list some new elements that do not lie in $R_{n+1} \cup B_{n+1}$. This takes care of (†11) and (†12) and completes the recursion step $n \mapsto n + 1$.

2.4.6. The inductive step: construction. Given the two trees T_n and S_n , we extend each of them through their roots as indicated in Figure 2.6 to trees \tilde{T}_n and \tilde{S}_n respectively. The trees T_{n+1} and S_{n+1} will be obtained as components of the promise closure of the forest $G_n = \tilde{T}_n \sqcup \tilde{S}_n$ with respect to the coloured promise edges.

Since v is not the root of T_n , there is a first edge e on the unique path in T_n from v to the root.

$$(3) \quad \text{This edge we also call } e(v).$$

Then $T_n - e$ has two connected components: one that contains the root of T_n which we name $T_n(r)$, and one that contains v which we name $T_n(v)$.

Since every maximal bare path in T_n has length at most k_n by (†5), it follows from Lemma 2.4.2 that all maximal bare paths in $T_n - e$, and so all bare paths in $T_n(r)$ and $T_n(v)$, have bounded length. Let $k = \tilde{k}_n$ be twice the maximum of the length of bare paths in T_n , S_n , $T_n(r)$ and $T_n(v)$, which exists by (†5).




FIGURE 2.6. All dotted lines are maximal bare paths of length at least $k = \tilde{k}_n$. The trees D_n are binary trees of height $b_n + 3$, hence $D_n \not\rightarrow T_n$ and $D_n \not\rightarrow S_n$ by ((†4)).

To obtain \tilde{T}_n , we extend T_n through its root $r(T_n) \in R_n$ by a path

$$r(T_n) = u_0, u_1, \dots, u_{p-1}, u_p = r(\hat{S}_n)$$

of length $p = 4(\tilde{k}_n + 1) + 3$, where at its last vertex u_p we glue a rooted copy \hat{S}_n of S_n (via an isomorphism $\hat{w} \leftrightarrow w$), identifying u_p with the root of \hat{S}_n .

Next, we add two additional leaves at u_0 and u_p , so that $\deg(r(T_n)) = 3 = \deg(r(\hat{S}_n))$. Further, we add a leaf $r(T_{n+1})$ at u_{2k+2} , which will be our new root for the next tree T_{n+1} ; and another leaf g at u_{2k+5} . Finally, we take a copy D_n of a rooted binary tree of height $b_n + 3$ and connect its root via an edge to u_{2k+3} . This completes the construction of \tilde{T}_n .

The construction of \tilde{S}_n is similar, but with a twist. For its construction, we extend S_n through its root $r(S_n) \in B_n$ by a path

$$r(S_n) = v_p, v_{p-1}, \dots, v_1, v_0 = r(\hat{T}_n(r))$$

of length p , where at its last vertex v_0 we glue a copy $\hat{T}_n(r)$ of $T_n(r)$, identifying v_0 with the root of $\hat{T}_n(r)$. Then, we take a copy $\hat{T}_n(\hat{v})$ of $T_n(v)$ and connect \hat{v} via an edge to v_{k+1} .

(4) This edge we call $e(\hat{v})$.

Finally, as before, we add two leaves at v_0 and v_p so that $\deg(r(\hat{T}_n(r))) = 3 = \deg(r(S_n))$. Next, we add a leaf $r(S_{n+1})$ to v_{2k+5} , which will be our new root for the next tree S_{n+1} ; and another leaf y to v_{2k+2} . Finally, we take another copy \hat{D}_n of a rooted

binary tree of height $b_n + 3$ and connect its root via an edge to v_{2k+3} . This completes the construction of \tilde{S}_n .

By the induction hypothesis, certain leaves of T_n have been coloured with one of the two colours $R_n \cup B_n$, and also some leaves of S_n have been coloured with one of the two colours $R_n \cup B_n$. In the above construction, we colour leaves of \hat{S}_n , $\hat{T}_n(r)$ and $\hat{T}_n(\hat{v})$ accordingly:

$$(5) \quad \begin{aligned} \tilde{R}_n &= \left(R_n \cup \left\{ \hat{w} \in \hat{S}_n \cup \hat{T}_n(r) \cup \hat{T}_n(\hat{v}) : w \in R_n \right\} \right) \setminus \left\{ r(T_n), r(\hat{T}_n(r)) \right\}, \\ \tilde{B}_n &= \left(B_n \cup \left\{ \hat{w} \in \hat{S}_n \cup \hat{T}_n(r) \cup \hat{T}_n(\hat{v}) : w \in B_n \right\} \right) \setminus \left\{ r(S_n), r(\hat{S}_n) \right\}. \end{aligned}$$

Now put $G_n := \tilde{T}_n \sqcup \tilde{S}_n$ and consider the following promise structure $\mathcal{P} = (G_n, \vec{P}, \mathcal{L})$ on G_n , consisting of four promise edges $\vec{P} = \{\vec{p}_1, \vec{p}_2, \vec{p}_3, \vec{p}_4\}$ and corresponding leaf sets $\mathcal{L} = \{L_1, L_2, L_3, L_4\}$, as follows:

$$(6) \quad \begin{aligned} &\bullet \vec{p}_1 \text{ pointing in } T_n \text{ towards the root } r(T_n), \text{ with } L_1 = \tilde{R}_n, \\ &\bullet \vec{p}_2 \text{ pointing in } S_n \text{ towards the root } r(S_n), \text{ with } L_2 = \tilde{B}_n, \\ &\bullet \vec{p}_3 \text{ pointing in } \tilde{T}_n \text{ towards the root } r(T_{n+1}), \text{ with } L_3 = \{r(T_{n+1}), y\}, \\ &\bullet \vec{p}_4 \text{ pointing in } \tilde{S}_n \text{ towards the root } r(S_{n+1}), \text{ with } L_4 = \{r(S_{n+1}), g\}. \end{aligned}$$

Note that our construction so far has been tailored to provide us with a \vec{P} -respecting isomorphism

$$(7) \quad h: \tilde{T}_n - v \rightarrow \tilde{S}_n - \hat{v}.$$

Consider the closure $\text{cl}(G_n)$ with respect to the promise structure \mathcal{P} defined above. Since $\text{cl}(G_n)$ is a leaf-extension of G_n , it has two connected components, just as G_n . We now define

$$(8) \quad \begin{aligned} T_{n+1} &= \text{the component containing } T_n \text{ in } \text{cl}(G_n), \text{ and} \\ S_{n+1} &= \text{the component containing } S_n \text{ in } \text{cl}(G_n). \end{aligned}$$

It follows that $\text{cl}(G_n) = T_{n+1} \sqcup S_{n+1}$ and $\hat{v} \in V(S_{n+1})$. Further, since \vec{p}_3 and \vec{p}_4 are placeholder promises, $\text{cl}(G)$ carries a corresponding promise structure, see Definition 2.3.6. We define

$$(9) \quad R_{n+1} = \text{cl}(L_3) \quad \text{and} \quad B_{n+1} = \text{cl}(L_4).$$

Lastly, we set

$$(10) \quad \begin{aligned} X_{n+1} &= X_n \cup \{v\}, \\ Y_{n+1} &= Y_n \cup \{\hat{v}\}, \text{ and} \\ \varphi_{n+1} &= \varphi_n \cup \{(v, \hat{v})\}, \end{aligned}$$

and put

$$(11) \quad k_{n+1} = 2\tilde{k}_n + 3 \quad \text{and} \quad b_{n+1} = b_n + 3$$

The construction of trees T_{n+1} and S_{n+1} , coloured leaf sets R_{n+1} and B_{n+1} , the bijection $\varphi_{n+1}: X_{n+1} \rightarrow Y_{n+1}$, and integers k_{n+1} and b_{n+1} is now complete. In the following, we verify that (†1)–(†14) are indeed satisfied for the $(n+1)$ th instance.

2.4.7. The inductive step: verification.

CLAIM 1. T_{n+1} and S_{n+1} extend T_n and S_n . Moreover, they are rooted trees of maximum degree 3 such that their respective roots are contained in R_{n+1} and B_{n+1} . Hence, (†1)–(†3) are satisfied.

PROOF. Property (†1) follows from (cl.1), i.e. that $\text{cl}(G_n)$ is a leaf-extension of G_n . Thus, T_{n+1} is a leaf extension of \tilde{T}_n , which in turn is a leaf extension of T_n , and similar for S_n . This shows (†1).

As noted after the proof of Proposition 2.3.3, taking the closure does not affect the maximum degree, i.e. $\Delta(\text{cl}(G_n)) = \Delta(G_n) = 3$. This shows (†2).

Finally, (9) implies (†3), as $r(T_{n+1}) \in R_{n+1}$ and $r(S_{n+1}) \in B_{n+1}$. \square

CLAIM 2. All binary trees appearing as subgraphs of $T_{n+1} \sqcup S_{n+1}$ have height at most b_{n+1} , and every such tree of height b_{n+1} is some copy D_n or \hat{D}_n . Hence, T_{n+1} and S_{n+1} satisfy (†4).

PROOF. We first claim that all binary trees appearing as subgraphs of $\tilde{T}_n \sqcup \tilde{S}_n$ which are not contained in D_n or \hat{D}_n have height at most $b_n + 1$. Indeed, note that any binary tree appearing as a subgraph of T_n , $\hat{T}_n(r)$, $\hat{T}_n(v)$, \hat{S}_n or S_n has height at most b_n by the inductive hypothesis. Since the paths we added to the roots of T_n and \hat{S}_n to form \tilde{T}_n were sufficiently long, any binary tree appearing as a subgraph of \tilde{T}_n can only meet one of T_n , \hat{S}_n or D_n . Since the roots of T_n and \hat{S}_n are adjacent to two new vertices in \tilde{T}_n , one of degree 1, any such tree meeting T_n or \hat{S}_n must have height at most $b_n + 1$. By Figure 2.6 we see that any binary tree in \tilde{T}_n which meets D_n but whose root lies outside of D_n has height at most $3 \leq b_n + 1$. Consider then a binary tree whose root lies inside D_n , but that is not contained in D_n . Again, by Figure 2.6 we see that the root of D_n must lie in one of the bottom three layers of this binary tree. Hence, if the root of this tree lies on the k th level of D_n , then the tree can have height at most $\min\{b_n + 3 - k, k + 2\}$, and hence the tree has height at most $b_n/2 + 2 \leq b_n + 1$. Any other binary tree meeting D_n is then contained in D_n . It follows that the only binary tree of height $b_n + 3$ appearing as a subgraph of \tilde{T}_n is D_n , and a similar argument holds for \tilde{S}_n and \hat{D}_n .

Recall that T_{n+1} and S_{n+1} are the components of $\text{cl}(\tilde{T}_n \sqcup \tilde{S}_n)$ containing \tilde{T}_n and \tilde{S}_n respectively. If we refer back to Section 2.3 we see that T_{n+1} can be formed from \tilde{T}_n by repeatedly gluing components isomorphic to $\tilde{T}_n(\vec{p}_1)$ or $\tilde{S}_n(\vec{p}_2)$ to leaves. Consider a binary

tree appearing as a subgraph of T_{n+1} which is contained in \tilde{T}_n or one of the copies of $\tilde{T}_n(\vec{p}_1)$ or $\tilde{S}_n(\vec{p}_2)$. By the previous paragraph, this tree has height at most $b_n + 3$, and if it has height $b_n + 3$ it is a copy D_n or \hat{D}_n . Suppose then that there is a binary tree, of height b , whose root is in \tilde{T}_n , but is not contained in \tilde{T}_n . Such a tree must contain some vertex $\ell \in \tilde{T}_n$ which is adjacent to a vertex not in \tilde{T}_n . Hence, ℓ must have been a leaf in \tilde{T}_n at which a copy of $\tilde{T}_n(\vec{p}_1)$ or $\tilde{S}_n(\vec{p}_2)$ was glued on. However, the roots of each of these components are adjacent to just two vertices, one of degree 1, and hence this leaf ℓ must either be in the bottom, or second to bottom layer of the binary tree. Therefore, $b \leq b_n + 2$. A similar argument holds when the root lies in some copy of $\tilde{T}_n(\vec{p}_1)$ or $\tilde{S}_n(\vec{p}_2)$, and also for S_{n+1} .

Therefore, all binary trees appearing as subgraphs of $T_{n+1} \sqcup S_{n+1}$ have height at most $b_n + 3$, and every such tree is some copy D_n or \hat{D}_n . Hence, since $b_{n+1} = b_n + 3$, it follows that $b_{n+1} \geq b_n$ and T_{n+1} and S_{n+1} satisfy $(\dagger 4)$. \square

CLAIM 3. Every maximal bare path in $T_{n+1} \sqcup S_{n+1}$ has length at most k_{n+1} . Hence, T_{n+1} and S_{n+1} satisfy $(\dagger 5)$.

PROOF. We first claim that all maximal bare paths in $\tilde{T}_n \sqcup \tilde{S}_n$ have length at most $2\tilde{k}_n + 3$. Firstly, we note that any maximal bare path which is contained in T_n or \hat{S}_n has length at most $k_n \leq \tilde{k}_n$ by the induction hypothesis. Also, since the roots of T_n and \hat{S}_n have degree 3 in \tilde{T}_n , any maximal bare path is either contained in T_n or \hat{S}_n , or does not contain any interior vertices from T_n or \hat{S}_n . However, it is clear from the construction that any maximal bare path in \tilde{T}_n that does not contain any interior vertices from T_n or \hat{S}_n has length at most $2\tilde{k}_n + 3$. Similarly, any maximal bare path which is contained in $\hat{T}_n(r)$, $\hat{T}_n(v)$, or S_n has length at most \tilde{k}_n by definition. By the same reasoning as above, any maximal bare path in \tilde{S}_n not contained in $\hat{T}_n(r)$, $\hat{T}_n(v)$, or S_n has length at most $2\tilde{k}_n + 3$.

Again, recall that T_{n+1} can be formed from \tilde{T}_n by repeatedly gluing components isomorphic to $\tilde{T}_n(\vec{p}_1)$ or $\tilde{S}_n(\vec{p}_2)$ to leaves. Any maximal bare path in T_{n+1} which is contained in \tilde{T}_n or one of the copies of $\tilde{T}_n(\vec{p}_1)$ or $\tilde{S}_n(\vec{p}_2)$ has length at most $2\tilde{k}_n + 3$ by the previous paragraph. However, since every interior vertex in a maximal bare path has degree two, and the vertices in T_{n+1} at which we, at some point in the construction, stuck on copies of $\tilde{T}_n(\vec{p}_1)$ or $\tilde{S}_n(\vec{p}_2)$ have degree 3, any maximal bare path in T_{n+1} must be contained in \tilde{T}_n or one of the copies of $\tilde{T}_n(\vec{p}_1)$ or $\tilde{S}_n(\vec{p}_2)$. Again, a similar argument holds for S_{n+1} . Hence, all maximal bare paths in $T_{n+1} \sqcup S_{n+1}$ have length at most $2\tilde{k}_n + 3$. Therefore, since $k_{n+1} = 2\tilde{k}_n + 3$, it follows that $k_{n+1} \geq k_n$ and T_{n+1} and S_{n+1} satisfy $(\dagger 5)$. \square

CLAIM 4. $\text{Ball}_{T_{n+1}}(T_n, k_n + 1)$ is a bare extension of T_n at $R_n \cup B_n$ to length $k_n + 1$ and does not meet $R_{n+1} \cup B_{n+1}$ and similarly for S_{n+1} . Hence, T_{n+1} and S_{n+1} satisfy $(\dagger 6)$ and $(\dagger 7)$ respectively.

PROOF. We will show that T_{n+1} satisfies †6, the proof that S_{n+1} satisfies †7 is analogous. By Proposition 2.3.3, the tree T_{n+1} is an $\left((\tilde{R}_n \cup \tilde{B}_n) \cap V(\tilde{T}_n)\right)$ -extension of \tilde{T}_n . Hence T_{n+1} is an

$$(12) \quad \left(\left((\tilde{R}_n \cup \tilde{B}_n) \cap V(T_n)\right) \cup r(T_n)\right) = \left((R_n \cup B_n) \cap V(T_n)\right)\text{-extension of } T_n.$$

By looking at the construction of $\text{cl}(G)$ from Section 2.3, we see that T_{n+1} is also an L' -extension of the supertree $T' \supseteq T_n$ formed by gluing a copy of $\tilde{T}_n(\vec{p}_1)$ to every leaf in $R_n \cap V(T_n)$ and a copy of $\tilde{S}_n(\vec{p}_2)$ to every leaf in $B_n \cap V(T_n)$, where the leaves in L' are the inherited promise leaves from the copies of $\tilde{T}_n(\vec{p}_1)$ and $\tilde{S}_n(\vec{p}_2)$.

However, we note that every promise leaf in $\tilde{T}_n(\vec{p}_1)$ and $\tilde{S}_n(\vec{p}_2)$ is at distance at least $\tilde{k}_n + 1$ from the respective root, and so $\text{Ball}_{T_{n+1}}(T_n, \tilde{k}_n) = \text{Ball}_{T'}(T_n, \tilde{k}_n)$. However, $\text{Ball}_{T'}(T_n, \tilde{k}_n)$ can be seen immediately to be a bare extension of T_n at $R_n \cup B_n$ to length \tilde{k}_n , and since $\tilde{k}_n \geq k_n + 1$ it follows that $\text{Ball}_{T_{n+1}}(T_n, k_n + 1)$ is a bare extension of T_n at $R_n \cup B_n$ to length $k_n + 1$ as claimed.

Finally, we note that $R_{n+1} \cup B_{n+1}$ is the set of promise leaves $\text{cl}(\mathcal{L}_n)$. By the same reasoning as before, $\text{Ball}_{T_{n+1}}(T_n, k_n + 1)$ contains no promise leaf in $\text{cl}(\mathcal{L}_n)$, and so does not meet $R_{n+1} \cup B_{n+1}$ as claimed. \square

CLAIM 5. Let U_{n+1} be a bare extension of $\text{cl}(G_n) = T_{n+1} \sqcup S_{n+1}$ at $R_{n+1} \cup B_{n+1}$ to any length. Then any embedding of T_{n+1} or S_{n+1} into U_{n+1} fixes the respective root. Hence, T_{n+1} and S_{n+1} satisfy †8.

PROOF. Recall that the promise closure was constructed by recursively adding copies of rooted trees C_i and identifying their roots with promise leaves. For the promise structure $\mathcal{P} = (G_n, \vec{P}, \mathcal{L})$ on G_n we have $C_1 = \tilde{T}_n(\vec{p}_1)$ and $C_2 = \tilde{S}_n(\vec{p}_2)$.

Note that by †5, the image of any embedding $T_n \hookrightarrow U_{n+1}$ cannot contain a bare path of length $k_n + 1$. Also, by construction, every copy of T_n , S_n , $\hat{T}_n(r)$, or $\hat{T}_n(\hat{v})$ in T_{n+1} has the property that its $(k_n + 1)$ -ball in T_{n+1} is a bare extension to length $k_n + 1$ of this copy. Hence, if the root of T_n embeds into some copy of T_n , S_n , $\hat{T}_n(r)$, or $\hat{T}_n(\hat{v})$, then the whole tree embeds into a bare extension of this copy. The same is true for S_n .

By †8, there are no embeddings of T_n into a bare extension of S_n , or of S_n into a bare extension of T_n . Moreover, since both $\hat{T}_n(r)$ and $\hat{T}_n(\hat{v})$ are subtrees of T_n , there is no embedding of T_n or S_n into bare extensions of them by †8 and †9.

Thus, only the following embeddings are possible:

- T_n embeds into a bare extension of a copy of T_n , or S_n embeds into a bare extension of a copy of S_n . In both cases, the root must be preserved, as otherwise we contradict †9 or †10.

Let $f: T_{n+1} \hookrightarrow U_{n+1}$ be an embedding. By Claim 2, U_{n+1} contains no binary trees of height $b_n + 3$ apart from D_n , \hat{D}_n , and the copies of those two trees that were created by adding copies of C_1 and C_2 . Consequently f maps D_n to one of these copies, mapping

the root to the root. The neighbours of $r(T_{n+1})$ and g must map to vertices of degree 3 at distance two and three from the image of the root of D_n respectively, which forces $f(r(T_{n+1})) \in R_{n+1}$. If $f(r(T_{n+1})) = r(T_{n+1})$ then we are done.

Otherwise there are two possibilities for $f(r(T_{n+1}))$. If $f(r(T_{n+1}))$ is contained in a copy of C_1 , then $r(T_n)$ maps to a promise leaf other than the root in a copy of T_n , S_n , $\hat{T}_n(r)$, or $\hat{T}_n(\hat{v})$. If $f(r(T_{n+1})) = y$ or $f(r(T_{n+1}))$ is contained in a copy of C_2 , then $r(T_n)$ maps to a copy of $r(\hat{T}_n(r))$ or some vertex of $\hat{T}_n(\hat{v})$. In both cases the root of T_n does not map to the root of a copy of T_n , which is impossible by the first bullet point.

Finally, let $f: S_{n+1} \hookrightarrow U_{n+1}$ be an embedding. By the same arguments as above $f(r(S_{n+1})) \in B_{n+1}$. If f fixes $r(S_{n+1})$, we are done.

Otherwise we have again two cases. If $f(r(S_{n+1})) = g$, or $f(r(S_{n+1}))$ is contained in a copy of C_1 , then v_{k+1} (the neighbour of \hat{v} on the long path) would have to map to a vertex of degree 2, giving an immediate contradiction. If $f(r(S_{n+1}))$ is contained in a copy of C_2 , then $r(S_n)$ maps to a promise leaf other than the root in a copy of T_n , S_n , $\hat{T}_n(r)$, or $\hat{T}_n(\hat{v})$ which is also impossible by the observations in the bullet points. \square

CLAIM 6. Let U_{n+1} be as in Claim 5. Then there is no embedding of T_{n+1} or S_{n+1} into U_{n+1} whose image contains vertices outside of $\text{cl}(G_n)$, i.e. vertices that have been added to form the bare extension.

Since a root-preserving embedding of a locally finite tree into itself must be an automorphism, this together with the previous claim implies $(\dagger 9)$ and $(\dagger 10)$.

PROOF. We prove this claim for T_{n+1} , the proof for S_{n+1} is similar. Assume for a contradiction that there is a vertex w of T_{n+1} and an embedding $f: T_{n+1} \hookrightarrow U_{n+1}$ such that $f(w) \notin \text{cl}(G_n)$. By definition of bare extension, removing $f(w)$ from U_{n+1} splits the component of $f(w)$ into at most two components, one of which is a path.

Note first that w does not lie in a copy of D_n or \hat{D}_n , because these must map to binary trees of the same height by Claim 2. Furthermore, all vertices in $R_{n+1} \cup B_{n+1}$ have a neighbour of degree 3 whose neighbours all have degree ≥ 2 , thus $w \notin R_{n+1} \cup B_{n+1}$. Finally, only one component of $T_{n+1} - w$ can contain vertices of degree 3. Consequently, w must lie in a copy C of T_n , S_n , $\hat{T}_n(r)$, or $\hat{T}_n(\hat{v})$.

All maximal bare paths in the image $f(C)$ have length at most $k = \tilde{k}_n$, so $f(C)$ cannot intersect any copies of T_n , S_n , $\hat{T}_n(r)$, or $(\hat{T}_n(\hat{v}) + v_{k+1})$. Let r be the root of C (where $r = \hat{v}$ in the last case). Now $f(r)$ must have the following properties: it is a vertex of degree 3, and the root of a nearest binary tree of height b_{n+1} not containing $f(r)$ lies at distance d from $f(r)$, where $5 \leq d \leq 2k + 4$.

But the only vertices with these properties are contained in copies of T_n , \hat{S}_n , $\hat{T}_n(r)$, or $(\hat{T}_n(\hat{v}) + v_{k+1})$. This contradicts the fact that $f(C)$ does not intersect any of these copies. \square

CLAIM 7. The function φ_{n+1} is a well-defined bijection extending φ_n , such that its domain and range do not intersect $R_{n+1} \cup B_{n+1}$. Hence, property (†13) holds for $\varphi_{n+1}: X_{n+1} \rightarrow Y_{n+1}$.

PROOF. By the choice of x in (1) and the definition of $\varphi_{n+1}: X_{n+1} \rightarrow Y_{n+1}$ in (10), the first three items of property (†13) hold.

Since v does not lie in $R_n \cup B_n$ by (2), it follows by our construction of the promise structure $\mathcal{P} = (G_n, \vec{P}, \mathcal{L})$ in (5) and (6) that neither v nor $\hat{v} = \varphi_{n+1}(v)$ appear as promise leaves in \mathcal{L} . Furthermore, by the induction hypothesis, $(X_n \cup Y_n) \cap (R_n \cup B_n) = \emptyset$, so no vertex in $(X_n \cup Y_n)$ appears as a promise leaf in \mathcal{L} either. Thus, in formulas,

$$(13) \quad (X_{n+1} \cup Y_{n+1}) \cap \bigcup_{L \in \mathcal{L}} L = \emptyset.$$

In particular, since

$$(R_{n+1} \cup B_{n+1}) \cap G_n = (\text{cl}(L_3) \cup \text{cl}(L_4)) \cap G_n = L_3 \cup L_4,$$

and $X_{n+1} \cup Y_{n+1} \subseteq G_n$, we get $(X_{n+1} \cup Y_{n+1}) \cap (R_{n+1} \cup B_{n+1}) = \emptyset$. Thus, also the last item of (†13) is verified. \square

CLAIM 8. There is a family of isomorphisms $\mathcal{H}_{n+1} = \{h_{n+1,x}: x \in X_{n+1}\}$ witnessing that $T_{n+1} - x$ and $S_{n+1} - \varphi_{n+1}(x)$ are isomorphic for all $x \in X_{n+1}$, such that $h_{n+1,x}$ extends $h_{n,x}$ for all $x \in X_n$. Hence, property (†14) holds.

PROOF. There are four things to be verified for this claim. Firstly, we need an isomorphism $h_{n+1,v}$ witnessing that $T_{n+1} - v$ and $S_{n+1} - \hat{v}$ are isomorphic. Secondly, we need to *extend* all previous isomorphisms $h_{n,x}$ between $T_n - x$ and $S_n - \varphi_n(x)$ to $T_{n+1} - x$ and $S_{n+1} - \varphi_n(x)$. This will take care of the first item of (†14). To also comply with the remaining two items, we need to make sure that each isomorphism in

$$\mathcal{H}_{n+1} = \{h_{n+1,x}: x \in X_{n+1}\}$$

maps leaves in $R_{n+1} \cap V(T_{n+1})$ bijectively to leaves in $R_{n+1} \cap V(S_{n+1})$, and similarly for B_{n+1} .

To find the first isomorphism, note that by construction of the promise structure $\mathcal{P} = (G_n, \vec{P}, \mathcal{L})$ on G_n in (5), and properties (cl.1) and (cl.3) of the promise closure, the trees T_{n+1} and S_{n+1} are obtained from \tilde{T}_n and \tilde{S}_n by attaching at every leaf $r \in \tilde{R}_n$ a copy of the rooted tree $\text{cl}(G_n)(\vec{p}_1)$, and by attaching at every leaf $b \in \tilde{B}_n$ a copy of the rooted tree $\text{cl}(G_n)(\vec{p}_2)$.

By (13), neither v nor $\varphi_{n+1}(v)$ are mentioned in \mathcal{L} . As observed in (7), there is a \vec{P} -respecting isomorphism

$$h: \tilde{T}_n - v \rightarrow \tilde{S}_n - \varphi_{n+1}(v).$$

In other words, h maps promise leaves in $L_i \cap V(\tilde{T}_n)$ bijectively to the promise leaves in $L_i \cap V(\tilde{S}_n)$ for all $i = 1, 2, 3, 4$. Our plan is to extend h to an isomorphism between

$T_{n+1}-v$ and $S_{n+1}-\varphi_n(v)$ by mapping the corresponding copies of $\text{cl}(G_n)(\vec{p}_1)$ and $\text{cl}(G_n)(\vec{p}_2)$ attached to the various red and blue leaves to each other.

Formally, by (cl.3) there is for each $\ell \in (\tilde{R}_n \cup \tilde{B}_n) \cap V(T)$ a $\text{cl}(\vec{P})$ -respecting isomorphism of rooted trees

$$\text{cl}(G_n)(\vec{q}_\ell) \cong \text{cl}(G_n)(\vec{q}_{h(\ell)}).$$

Therefore, by combining the isomorphism h between $\tilde{T}_n - v$ and $\tilde{S}_n - \varphi_{n+1}(v)$ with these isomorphisms between each $\text{cl}(G_n)(\vec{q}_\ell)$ and $\text{cl}(G_n)(\vec{q}_{h(\ell)})$ we get a $\text{cl}(\vec{P})$ -respecting isomorphism

$$h_{n+1,v}: T_{n+1} - v \rightarrow S_{n+1} - \varphi_{n+1}(v).$$

And since R_{n+1} and B_{n+1} have been defined in (9) to be the promise leaf sets of $\text{cl}(\mathcal{P})$, by definition of $\text{cl}(\vec{P})$ -respecting (Def. 2.3.4), the image of $R_{n+1} \cap V(T_{n+1})$ under $h_{n+1,v}$ is $R_{n+1} \cap V(S_{n+1})$, and similarly for B_{n+1} .

It remains to extend the old isomorphisms in \mathcal{H}_n . As argued in (12), both trees T_{n+1} and S_{n+1} are leaf extensions of T_n and S_n at $R_n \cup B_n$ respectively. By property (cl.3), these leaf extensions are obtained by attaching at every leaf $r \in R_n$ a copy of the rooted tree $\text{cl}(G_n)(\vec{p}_1)$, and similarly by attaching at every leaf $b \in B_n$ a copy of the rooted tree $\text{cl}(G_n)(\vec{p}_2)$.

By induction assumption (†14), for each $x \in X_n$ the isomorphism

$$h_{n,x}: T_n - x \rightarrow S_n - \varphi_n(x)$$

maps the red leaves of T_n bijectively to the red leaves of S_n , and the blue leaves of T_n bijectively to the blue leaves of S_n . Thus, by property (cl.3), there are $\text{cl}(\vec{P})$ -respecting isomorphisms of rooted trees

$$\text{cl}(G_n)(\vec{q}_\ell) \cong \text{cl}(G_n)(\vec{q}_{h_{n,x}(\ell)})$$

for all $\ell \in (R_n \cup B_n) \cap V(T_n)$. By combining the isomorphism $h_{n,x}$ between $T_n - x$ and $S_n - \varphi_n(x)$ with these isomorphisms between each $\text{cl}(G_n)(\vec{q}_\ell)$ and $\text{cl}(G_n)(\vec{q}_{h_{n,x}(\ell)})$, we obtain a $\text{cl}(\vec{P})$ -respecting extension

$$h_{n+1,x}: T_{n+1} - x \rightarrow S_{n+1} - \varphi_n(x).$$

As before, by definition of $\text{cl}(\vec{P})$ -respecting, the image of $R_{n+1} \cap V(T_{n+1})$ under $h_{n+1,x}$ is $R_{n+1} \cap V(S_{n+1})$, and similarly for B_{n+1} .

Finally, by construction we have $h_{n+1,x} \upharpoonright (T_n - x) = h_{n,x}$ for all $x \in X_n$ as desired. The proof is complete. \square

2.5. The trees are also edge-hypomorphic

In this final section, we briefly indicate why the trees T and S yielded by our strategy above are automatically edge-hypomorphic: we claim the correspondence

$$\psi: e(x) \mapsto e(\varphi(x))$$

as introduced in (3) and (4) is an edge-hypomorphism between T and S . For this, we need to verify that

- (a) ψ is a bijection between $E(T)$ and $E(S)$, and that
- (b) the maps $h_x \cup \{(x, \varphi(x))\}: G - e(x) \rightarrow H - e(\varphi(x))$ are isomorphisms.

Regarding (b), observe that the map h as defined in (7) yields, by construction, also a \vec{P} -respecting isomorphism

$$h \cup \{(v, \hat{v})\}: \tilde{T}_n - e(v) \rightarrow \tilde{S}_n - e(\hat{v}),$$

and from there, the arguments are entirely the same as in the previous section.

For (a), we use the canonical bijection between the edge set of a rooted tree, and its vertices other than the root; namely the bijection mapping every such vertex to the first edge on its unique path to the root. Thus, given the enumeration of $V(T_n)$ and $V(S_n)$ in (†11), we obtain corresponding enumerations of $E(T_n)$ and $E(S_n)$, and since the rooted trees T_n and S_n are order-preserving subtrees of the rooted trees T_{n+1} and S_{n+1} (cf. Figure 2.6), it follows that also our enumerations of $E(T_n)$ and $E(S_n)$ extend the enumerations of $E(T_{n-1})$ and $E(S_{n-1})$ respectively. But now it follows from (†13) and the definition of ψ that by step $2(n+1)$ we have dealt with the first n edges in our enumerations of $E(T)$ and $E(S)$ respectively.

CHAPTER 3

Non-reconstructible locally finite graphs

Two graphs G and H are *hypomorphic* if there exists a bijection $\varphi: V(G) \rightarrow V(H)$ such that $G - v \cong H - \varphi(v)$ for each $v \in V(G)$. A graph G is *reconstructible* if $H \cong G$ for all H hypomorphic to G .

Nash-Williams proved that all locally finite graphs with a finite number ≥ 2 of ends are reconstructible, and asked whether locally finite graphs with one end or countably many ends are also reconstructible.

In this paper we construct non-reconstructible graphs of bounded maximum degree with one and countably many ends respectively, answering the two questions of Nash-Williams about the reconstruction of locally finite graphs in the negative.

3.1. Introduction

Two graphs G and H are *hypomorphic* if there exists a bijection φ between their vertex sets such that the induced subgraphs $G - v$ and $H - \varphi(v)$ are isomorphic for each vertex v of G . We say that a graph G is *reconstructible* if $H \cong G$ for every H hypomorphic to G . The *Reconstruction Conjecture*, a famous unsolved problem attributed to Kelly and Ulam, suggests that every finite graph with at least three vertices is reconstructible.

For an overview of results towards the Reconstruction Conjecture for finite graphs see the survey of Bondy and Hemminger [29]. The corresponding reconstruction problem for infinite graphs is false: the countable regular tree T_∞ , and two disjoint copies of it (written as $T_\infty \cup T_\infty$) are easily seen to be non-homeomorphic reconstructions of each other. This example, however, contains vertices of infinite degree. Regarding locally finite graphs, Harary, Schwenk and Scott [91] showed that there exists a non-reconstructible locally finite forest. However, they conjectured that the Reconstruction Conjecture should hold for locally finite trees. This conjecture has been verified for locally finite trees with at most countably many ends in a series of paper [12, 28, 151]. However, very recently, the present authors have constructed a counterexample to the conjecture of Harary, Schwenk and Scott.

THEOREM 3.1.1 (Bowler, Erde, Heinig, Lehner, Pitz [36]). *There exists a non-reconstructible tree of maximum degree three.*

The Reconstruction Conjecture has also been considered for general locally finite graphs. Nash-Williams [126] showed that if $p \geq 3$ is an integer, then any locally finite graph with exactly p ends is reconstructible; and in [128] he showed the same is true for $p = 2$. The case $p = 2$ is significantly more difficult. Broadly speaking this is because every graph with $p \geq 3$ ends has some identifiable finite ‘centre’, from which the ends can be thought of as branching out. A two-ended graph however can be structured like a double ray, without an identifiable ‘centre’.

The case of 1-ended graphs is even harder, and the following problems from a survey of Nash-Williams [127], which would generalise the corresponding results established for trees, have remained open.

PROBLEM 3.1.2 (Nash-Williams). Is every locally finite graph with exactly one end reconstructible?

PROBLEM 3.1.3 (Nash-Williams). Is every locally finite graph with countably many ends reconstructible?

In this paper, we extend our methods from [36] to construct examples showing that both of Nash-Williams’ questions have negative answers. Our examples will not only be locally finite, but in fact have bounded degree.

THEOREM 3.1.4. *There is a connected one-ended non-reconstructible graph with bounded maximum degree.*

THEOREM 3.1.5. *There is a connected countably-ended non-reconstructible graph with bounded maximum degree.*

Since every locally finite graph has either finitely many, countably many or continuum many ends, Theorems 3.1.1, 3.1.4 and 3.1.5 together with the results of Nash-Williams provide a complete picture about what can be said about number of ends versus reconstruction:

- A locally finite tree with at most countably many ends is reconstructible; but there are non-reconstructible locally finite trees with continuum many ends.
- A locally finite graph with at least two, but a finite number of ends is reconstructible; but there are non-reconstructible locally finite graphs with one, countably many, and continuum many ends respectively.

This paper is organised as follows: In the next section we give a short, high-level overview of our constructions which answer Nash-Williams’ problems. In Sections 3.3 and 3.4, we develop the technical tools necessary for our construction, and in Sections 3.5 and 3.6, we prove Theorems 3.1.4 and 3.1.5.

For standard graph theoretical concepts we follow the notation in [54].

3.2. Sketch of the construction

In this section we sketch the main ideas of the construction in three steps. First, we quickly recall our construction of two hypomorphic, non-isomorphic locally finite trees from [36]. We will then outline how to adapt the construction to obtain a one-ended-, and a countably-ended counterexample respectively.

3.2.1. The tree case. This section contains a very brief summary of the much more detailed sketch from [36]. The strategy is to build trees T and S recursively, where at each step of the construction we ensure for some new vertex v already chosen for T that there is a corresponding vertex w of S with $T - v \cong S - w$, or vice versa. This will ensure that by the end of the construction, the trees we have built are hypomorphic.

More precisely, at step n we will construct subtrees T_n and S_n of our eventual trees, where some of the leaves of these subtrees have been coloured in two colours, say red and blue. We will only further extend the trees from these coloured leaves, and we will extend from leaves of the same colour in the same way. We also make sure that earlier partial isomorphisms between $T_n - v_i \cong S_n - w_i$ preserve leaf colours. Together, these requirements guarantee that earlier partial isomorphisms always extend to the next step.

The T_n will be nested, and we will take T to be the union of all of them; similarly the S_n will be nested and we take S to be the union of all of them. To ensure that T and S do not end up being isomorphic, we first ensure, for each n , that there is no isomorphism from T_n to S_n . Our second requirement is that T or S beyond any coloured leaf of T_n or S_n begins with a long non-branching path, longer than any such path appearing in T_n or S_n . Together, this implies that T and S are not isomorphic.



FIGURE 3.1. A first approximation of T_{n+1} on the left, and S_{n+1} on the right. All dotted lines are long non-branching paths.

Algorithm Stage One: Suppose now that we have already constructed T_n and S_n and wish to construct T_{n+1} and S_{n+1} . Suppose further that we are given a vertex v of T_n for which we wish to find a partner w in S_{n+1} so that $T - v$ and $S - w$ are isomorphic. We begin by building a tree $\hat{T}_n \not\cong T_n$ which has some vertex w such that $T_n - v \cong \hat{T}_n - w$. This can be done by taking the components of $T_n - v$ and arranging them suitably around the new vertex w .

We will take S_{n+1} to include S_n and \hat{T}_n , with the copies of red and blue leaves in \hat{T}_n also coloured red and blue respectively. As indicated on the right in Figure 3.1, we add

long non-branching paths to some blue leaf b of S_n and to some red leaf r of \hat{T}_n and join these paths at their other endpoints by some edge e_n . We also join two new leaves y and g to the endvertices of e_n . We colour the leaf y yellow and the leaf g green. To ensure that $T_{n+1} - v \cong S_{n+1} - w$, we take T_{n+1} to include T_n together with a copy \hat{S}_n of S_n , with its leaves coloured appropriately, and joined up in the same way, as indicated on the left in Figure 3.1. Note that, whilst \hat{S}_n and S_n are isomorphic as graphs, we make a distinction as we want to lift the partial isomorphisms between $T_n - v_i \cong S_n - w_i$ to these new graphs, and our notation aims to emphasize the natural inclusions $T_n \subseteq T_{n+1}$ and $S_n \subseteq S_{n+1}$.

Algorithm Stage Two: We now have committed ourselves to two targets which are seemingly irreconcilable: first, we promised to extend in the same way at each red or blue leaf of T_n and S_n , but we also need that $T_{n+1} - v \cong S_{n+1} - w$. The solution is to copy the same subgraph appearing beyond r in Fig. 3.1, including its coloured leaves, onto all the other red leaves of S_n and T_n . Similarly we copy the subgraph appearing beyond the blue leaf b of S_n onto all other blue leaves of S_n and T_n . In doing so, we create new red and blue leaves, and we will keep adding, step by step, further copies of the graphs appearing beyond r and b in Fig. 3.1 respectively onto all red and blue leaves of everything we have constructed so far.





FIGURE 3.2. A sketch of T_{n+1} and S_{n+1} after countably many steps.

After countably many steps we have dealt with all red and blue leaves, and it can be checked that both our targets are achieved. We take these new trees to be S_{n+1} and T_{n+1} . They are non-isomorphic, as after removing all long non-branching paths, T_{n+1} contains T_n as a component, whereas S_{n+1} does not.

3.2.2. The one-ended case. To construct a one-ended non-reconstructible graph, we initially follow the same strategy as in the tree case and build locally finite graphs G_n and H_n and some partial hypomorphisms between them. Simultaneously, however, we will also build one-ended locally finite graphs of a grid-like form $F_n \times \mathbb{N}$ (the Cartesian product of a locally finite tree F_n with a ray) which share certain symmetries with G_n and H_n . These will allow us to glue $F_n \times \mathbb{N}$ onto both G_n and H_n , in order to make

FIGURE 3.3. A sketch of G_1 (above) and H_1 (below).

them one-ended, without spoiling the partial isomorphisms. Let us illustrate this idea by explicitly describing the first few steps of the construction.

We start with two non-isomorphic graphs G_0 and H_0 , such that G_0 and H_0 each have exactly one red and one blue leaf. After stage one of our algorithm, our approximations to G_1 and H_1 as in Figure 3.1 contain, in each of G_0 , \hat{H}_0 , \hat{G}_0 and H_0 , one coloured leaf. In stage two, we add copies of these graphs recursively. It follows that the resulting graphs G'_1 and H'_1 have the global structure of a double ray, along which parts corresponding to copies of G_0 , \hat{H}_0 , \hat{G}_0 and H_0 appear in a repeating pattern. Crucially, however, each graph G'_1 and H'_1 has infinitely many yellow and green leaves, which appear in an alternating pattern extending to infinity in both directions along the double ray.

Consider the minor F_1 of G'_1 obtained by collapsing every subgraph corresponding to G_0 , \hat{H}_0 , \hat{G}_0 and H_0 to a single point. Write $\psi_G: G'_1 \rightarrow F_1$ for the quotient map. Then F_1 is a double ray with alternating coloured leaves hanging off it. Note that we could have started with H'_1 and obtained the same F_1 . In other words, F_1 approximates the global structures of both G'_1 and H'_1 . Consider the one-ended grid-like graph $F_1 \times \mathbb{N}$, where we let $F_1 \times \{0\}$ inherit the colours from F_1 . We now form G_1 and H_1 by gluing $F_1 \times \mathbb{N}$ onto G'_1 , by identifying corresponding coloured vertices y and $\psi_G(y)$, and similarly for H'_1 .¹ Since the coloured leaves contained both ends of our graphs in their closure, the graphs G_1 and H_1 are now one-ended.

It remains to check that our partial isomorphism $h_1: G'_1 - v_1 \rightarrow H'_1 - w_1$ guaranteed by step two can be extended to $G_1 - v_1 \rightarrow H_1 - w_1$. This can be done essentially because of the following property: let us write $\mathcal{L}(\cdot)$ for the set of coloured leaves. It can be checked that there is an automorphism $\pi_1: F_1 \rightarrow F_1$ such that the diagram

¹For technical reasons, in the actual construction we identify $\psi_G(y)$ with the corresponding base vertex of the leaf y in G'_1 . In this way the coloured leaves of G'_1 remain leaves, and we can continue our recursive construction.

$$\begin{array}{ccc}
\mathcal{L}(G'_1) & \xrightarrow{h_1 \upharpoonright \mathcal{L}(G_1)} & \mathcal{L}(H'_1) \\
\psi_G \downarrow & & \downarrow \psi_H \\
\mathcal{L}(F_1) & \xrightarrow{\pi_1 \upharpoonright \mathcal{L}(F_1)} & \mathcal{L}(F_1)
\end{array}$$

is colour-preserving and commutes. Hence, $\pi_1 \times \text{id}$ is an automorphism of $F_1 \times \mathbb{N}$ which is compatible with our gluing procedure, so it can be combined with h_1 to give us the desired isomorphism.

We are now ready to describe the general step. Instead of describing F_n as a minor of G_n , which no longer works naïvely at later steps, we will directly build F_n by recursion, so that it satisfies the properties of the above diagram.

Suppose at step n we have constructed locally finite graphs G_n and H_n , and also a locally finite tree F_n where some leaves are coloured in one of two colours. Furthermore, suppose we have a family of isomorphisms

$$\mathcal{H}_n = \{h_x : G_n - x \rightarrow H_n - \varphi(x) : x \in X_n\},$$

for some subset $X_n \subseteq V(G_n)$, a family of isomorphisms $\Pi_n = \{\pi_x : F_n \rightarrow F_n : x \in X_n\}$, and colour-preserving bijections $\psi_{G_n} : \mathcal{L}(G_n) \rightarrow \mathcal{L}(F_n)$ and $\psi_{H_n} : \mathcal{L}(H_n) \rightarrow \mathcal{L}(F_n)$ such that the corresponding commutative diagram from above holds for each x . We construct G'_{n+1} and H'_{n+1} according to stages one and two of the previous algorithm. As before our isomorphisms h_x will lift to isomorphisms between $G'_{n+1} - x$ and $H'_{n+1} - \varphi(x)$.




FIGURE 3.4. The auxiliary graph \tilde{F}_n .

Algorithm Stage Three. As indicated in Figure 3.4, we take two copies F_n^G and F_n^H of F_n , and glue them together mimicking stage one of the algorithm, i.e. connect $\psi_{G_n}(r)$ in F_n^G by a path of length three to $\psi_{H_n}(b)$ in F_n^H , and attach two new leaves coloured yellow and green in the middle of the path. Call the resulting graph \tilde{F}_n . We then apply stage two of the algorithm to this graph, gluing again and again onto every blue vertex a copy of the graph of \tilde{F}_n behind $\psi_{H_n}(b)$, and similarly for every red leaf, to obtain a tree F_{n+1} . Since this procedure is, in structural terms, so similar to the construction of G'_{n+1} and H'_{n+1} , it can be shown that we do obtain a colour-preserving commuting diagram of the form

$$\begin{array}{ccc}
\mathcal{L}(G'_{n+1}) & \xrightarrow{h_x \upharpoonright \mathcal{L}(G'_{n+1})} & \mathcal{L}(H'_{n+1}) \\
\psi_{G_{n+1}} \downarrow & & \downarrow \psi_{H_{n+1}} \\
\mathcal{L}(F_{n+1}) & \xrightarrow{\pi_x \upharpoonright \mathcal{L}(F_{n+1})} & \mathcal{L}(F_{n+1})
\end{array}$$

As before, this means that we can indeed glue together G'_{n+1} and $F_{n+1} \times \mathbb{N}$, and H'_{n+1} and $F_{n+1} \times \mathbb{N}$ to obtain one-ended graphs G_{n+1} and H_{n+1} as desired.

At the end of our construction, after countably many steps, we have built two graphs G and H which are hypomorphic, and for the same reasons as in the tree case the two graphs will not be isomorphic. Further, since all G_n and H_n are one-ended, so will be G and H .

3.2.3. The countably-ended case. In order to produce hypomorphic graphs with countably many ends we follow the same procedure as for the one-ended case, except that we start with one-ended (non-isomorphic) graphs G_0 and H_0 .

After the first and second stage of our algorithm, the resulting graphs G'_1 and H'_1 will again consist of infinitely many copies of G_0 and H_0 glued together along a double ray. After gluing $F_1 \times \mathbb{N}$ to these graphs as before, we obtain graphs with one thick end, with many coloured leaves tending to that end, as well as infinitely many thin ends, coming from the copies of G_0 and H_0 , each of which contained a ray. These thin ends will eventually be rays, and so have no coloured leaves tending towards them. This guarantees that in the next step, when we glue $F_2 \times \mathbb{N}$ onto G'_2 and H'_2 , the thin ends will not be affected, and that all the other ends in the graph will be amalgamated into one thick end.

Then, in each stage of the construction, the graphs G_n and H_n will have exactly one thick end, again with many coloured leaves tending towards it, and infinitely many thin ends each of which is eventually a ray. This property lifts to the graphs G and H constructed in the limit: they will have one thick end and infinitely many ends which are eventually rays. However, since G and H are countable, there can only be countably many of these rays. Hence the two graphs G and H have countably many ends in total, and as before they will be hypomorphic but not isomorphic.

3.3. Closure with respect to promises

A *bridge* in a graph G is an edge $e = \{x, y\}$ such that x and y lie in different components of $G - e$. Given a directed bridge $\vec{e} = \vec{x}\vec{y}$ in some graph $G = (V, E)$, we denote by $G(\vec{e})$ the unique component of $G - e$ containing the vertex y . We think of $G(\vec{e})$ as a rooted graph with root y .

DEFINITION 3.3.1 (Promise structure). A *promise structure* $\mathcal{P} = (G, \vec{P}, \mathcal{L})$ is a triple consisting of:

- a graph G ,
- $\vec{P} = \{\vec{p}_i : i \in I\}$ a set of directed bridges $\vec{P} \subseteq \vec{E}(G)$, and
- $\mathcal{L} = \{L_i : i \in I\}$ a set of pairwise disjoint sets of leaves of G .

We insist further that, if the component $G(\vec{p}_i)$ consists of a single leaf $c \in L_j$, then $i = j$.

Often, when the context is clear, we will not make a distinction between \mathcal{L} and the set $\bigcup_i L_i$, for notational convenience.

We call an edge $\vec{p}_i \in \vec{P}$ a *promise edge*, and leaves $\ell \in L_i$ *promise leaves*. A promise edge $\vec{p}_i \in \vec{P}$ is called a *placeholder-promise* if the component $G(\vec{p}_i)$ consists of a single leaf $c \in L_i$, which we call a *placeholder-leaf*. We write

$$\mathcal{L}_p = \{L_i : \vec{p}_i \text{ a placeholder-promise}\} \text{ and } \mathcal{L}_q = \mathcal{L} \setminus \mathcal{L}_p.$$

Given a leaf ℓ in G , there is a unique edge $q_\ell \in E(G)$ incident with ℓ , and this edge has a natural orientation \vec{q}_ℓ towards ℓ . Informally, we think of $\ell \in L_i$ as the ‘promise’ that if we extend G to a graph $H \supset G$, we will do so in such a way that $H(\vec{q}_\ell) \cong H(\vec{p}_i)$.

DEFINITION 3.3.2 (Leaf extension). Given an inclusion $H \supseteq G$ of graphs and a set L of leaves of G , H is called a *leaf extension*, or more specifically an *L -extension*, of G , if:

- every component of H contains precisely one component of G , and
- every component of $H - G$ is adjacent to a unique vertex l of G , and we have $l \in L$.

In [36], given a promise structure $\mathcal{P} = (G, \vec{P}, \mathcal{L})$, it is shown how to construct a graph $\text{cl}(G) \supset G$ which has the following properties.

PROPOSITION 3.3.3 (Closure w.r.t a promise structure, cf. [36, Proposition 3.3]). *Let G be a graph and let $(G, \vec{P}, \mathcal{L})$ be a promise structure. Then there is a graph $\text{cl}(G)$, called the closure of G with respect to \mathcal{P} , such that:*

- (cl.1) $\text{cl}(G)$ is an \mathcal{L}_q -extension of G ,
- (cl.2) for every $\vec{p}_i \in \vec{P}$ and all $\ell \in L_i$,

$$\text{cl}(G)(\vec{p}_i) \cong \text{cl}(G)(\vec{q}_\ell)$$

are isomorphic as rooted graphs.

Since the existence of $\text{cl}(G)$ is crucial to our proof, we briefly remind the reader how to construct such a graph. As a first approximation, in order to try to achieve ((cl.2)), we glue a copy of the component $G(\vec{p}_i)$ onto each leaf $\ell \in L_i$, for each $i \in I$. We call this the *1-step extension* $G^{(1)}$ of G . If there were no promise leaves in the component $G(\vec{p}_i)$, then the promises in L_i would be satisfied. However, if there are, then we have grown $G(\vec{p}_i)$ by adding copies of various $G(\vec{p}_j)$ s behind promise leaves appearing in $G(\vec{p}_i)$.

However, remembering all promise leaves inside the newly added copies of $G(\vec{p}_i)$ we glued behind each $\ell \in L_i$, we continue this process indefinitely, growing the graph one step

at a time by gluing copies of (the original) $G(\vec{p}_i)$ to promise leaves ℓ' which have appeared most recently as copies of $\ell \in L_i$. After a countable number of steps the resulting graph $\text{cl}(G)$ satisfies Proposition 3.3.3. We note also that the maximum degree of $\text{cl}(G)$ equals that of G .

DEFINITION 3.3.4 (Promise-respecting map). Let G be a graph, $\mathcal{P} = (G, \vec{P}, \mathcal{L})$ be a promise structure on G , and let T_1 and T_2 be two components of G .

Given $x \in T_1$ and $y \in T_2$, a bijection $\varphi: T_1 - x \rightarrow T_2 - y$ is \vec{P} -respecting (with respect to \mathcal{P}) if the image of $L_i \cap T_1$ under φ is $L_i \cap T_2$ for all i .

We can think of \mathcal{P} as defining a $|\vec{P}|$ -colouring on some sets of leaves. Then a mapping is \vec{P} -respecting if it preserves leaf colours.

Suppose that \vec{p}_i is a placeholder promise, and $G = H^{(0)} \subseteq H^{(1)} \subseteq \dots$ is the sequence of 1-step extensions whose direct limit is $\text{cl}(G)$. Then, if we denote by $L_i^{(n)}$ the set of promise leaves associated with \vec{p}_i in $H^{(n)}$, it follows that $L_i^{(n)} \supseteq L_i^{(n-1)}$ since $G(\vec{p}_i)$ is just a single vertex $c_i \in L_i$. For every placeholder promise $\vec{p}_i \in \vec{P}$, we define $\text{cl}(L_i) = \bigcup_n L_i^{(n)}$.

DEFINITION 3.3.5 (Closure of a promise structure). The *closure* of the promise structure $(G, \vec{P}, \mathcal{L})$ is the promise structure $\text{cl}(\mathcal{P}) = (\text{cl}(G), \text{cl}(\vec{P}), \text{cl}(\mathcal{L}))$, where:

- $\text{cl}(\vec{P}) = \{\vec{p}_i: \vec{p}_i \in \vec{P} \text{ is a placeholder-promise}\}$,
- $\text{cl}(\mathcal{L}) = \{\text{cl}(L_i): \vec{p}_i \in \vec{P} \text{ is a placeholder-promise}\}$.

PROPOSITION 3.3.6 ([36, Proposition 3.3]). *Let G be a graph and let $(G, \vec{P}, \mathcal{L})$ be a promise structure. Then $\text{cl}(G)$ satisfies:*

(cl.3) *for every $\vec{p}_i \in \vec{P}$ and every $\ell \in L_i$,*

$$\text{cl}(G)(\vec{p}_i) \cong \text{cl}(G)(\vec{q}_\ell)$$

are isomorphic as rooted graphs, and this isomorphism is $\text{cl}(\vec{P})$ -respecting with respect to $\text{cl}(\mathcal{P})$.

It is precisely this property (cl.3) of the promise closure that will allow us to maintain partial hypomorphisms during our recursive construction.

The last two results of this section serve as preparation for growing G_{n+1} , H_{n+1} and F_{n+1} ‘in parallel’, as outlined in the third stage of the algorithm in §3.2.2. If $\mathcal{L} = \{L_i: i \in I\}$ and $\mathcal{L}' = \{L'_i: i \in I\}$, we say a map $\psi: \bigcup \mathcal{L} \rightarrow \bigcup \mathcal{L}'$ is *colour-preserving* if $\psi(L_i) \subseteq L'_i$ for every i .

LEMMA 3.3.7. *Let $(G, \vec{P}, \mathcal{L})$ and $(G', \vec{P}', \mathcal{L}')$ be promise structures, and let $G = H^{(0)} \subseteq H^{(1)} \subseteq \dots$ and $G' = H'^{(0)} \subseteq H'^{(1)} \subseteq \dots$ be 1-step extensions approximating their respective closures.*

Assume that $\vec{P} = \{\vec{p}_1, \dots, \vec{p}_k\}$ and $\vec{P}' = \{\vec{r}_1, \dots, \vec{r}_k\}$, and that there is a colour-preserving bijection

$$\psi: \bigcup \mathcal{L} \rightarrow \bigcup \mathcal{L}'$$

such that (recall that $\mathcal{L}(\cdot)$ is the set of leaves of a graph that are in \mathcal{L})

$$\psi \upharpoonright G(\vec{p}_i): \mathcal{L}(G(\vec{p}_i)) \rightarrow \mathcal{L}'(G'(\vec{r}_i))$$

is still a colour-preserving bijection for all $\vec{p}_i \in \vec{P}$.

Then for each $i \leq k$ there is a sequence of colour-preserving bijections

$$\alpha_n^i: \mathcal{L}(H^{(n)}(\vec{p}_i)) \rightarrow \mathcal{L}'(H'^{(n)}(\vec{r}_i))$$

such that α_{n+1}^i extends α_n^i .

PROOF. Fix i . We proceed by induction on n . Put $\alpha_0^i := \psi \upharpoonright G(\vec{p}_i)$.

Now suppose that α_n^i exists. To form $H^{(n+1)}(\vec{p}_i)$, we glued a copy of $G(\vec{p}_j)$ to each $\ell \in L_j^{(n)} \cap H^{(n)}(\vec{p}_i)$ for all $j \leq k$, and to construct $H'^{(n+1)}(\vec{r}_i)$, we glued a copy of $G'(\vec{r}_j)$ to each $\ell' \in L_j'^{(n)} \cap H'^{(n)}(\vec{r}_i)$ for all $j \leq k$, in both cases keeping all copies of promise leaves.

By assumption, the second part can be phrased equivalently as: we glued on a copy of $G'(\vec{r}_j)$ to each $\alpha_n^i(\ell)$ for $\ell \in L_j^{(n)} \cap H^{(n)}(\vec{r}_i)$. Thus, we can now combine the bijections $\alpha_n^i(\ell)$ with all the individual bijections ψ between all newly added $G(\vec{p}_j)$ and $G'(\vec{r}_j)$ to obtain a bijection α_{n+1}^i as desired. \square

COROLLARY 3.3.8. *In the above situation, for each i there is a colour-preserving bijection α^i between $\mathcal{L}(\text{cl}(G)(\vec{p}_i))$ and $\mathcal{L}'(\text{cl}(G')(\vec{r}_i))$ with respect to the promise closures $\text{cl}(\mathcal{P})$ and $\text{cl}(\mathcal{P}')$.*

PROOF. Put $\alpha^i = \bigcup_n \alpha_n^i$. Because all α_n^i respected all colours, they respect in particular the placeholder promises which make up $\text{cl}(\mathcal{P})$ and $\text{cl}(\mathcal{P}')$. \square

3.4. Thickening the graph

In this section, we lay the groundwork for the third stage of our algorithm, as outlined in §3.2.2. Our aim is to clarify how gluing a one-ended graph F onto a graph G affects automorphisms and the end-space of the resulting graph.

DEFINITION 3.4.1 (Gluing sum). Given two graphs G and F , and a bijection ψ with $\text{dom}(\psi) \subseteq V(G)$ and $\text{ran}(\psi) \subseteq V(F)$, the *gluing sum of G and F along ψ* , denoted by $G \oplus_\psi F$, is the quotient graph $(G \cup F) / \sim$ where $v \sim \psi(v)$ for all $v \in \text{dom}(\psi)$.

Our first lemma of this section explains how a partial isomorphism from $G_n - x$ to $H_n - \varphi(x)$ in our construction can be lifted to the gluing sum of G_n and H_n with a graph F respectively.

LEMMA 3.4.2. *Let G , H and F be graphs, and consider two gluing sums $G \oplus_{\psi_G} F$ and $H \oplus_{\psi_H} F$ along partial bijections ψ_G and ψ_H . Suppose there exists an isomorphism $h: G - x \rightarrow H - y$ that restricts to a bijection between $\text{dom}(\psi_G)$ and $\text{dom}(\psi_H)$.*

Then h extends to an isomorphism $(G \oplus_{\psi_G} F) - x \rightarrow (H \oplus_{\psi_H} F) - y$ provided there is an automorphism π of F such that $\pi \circ \psi_G(v) = \psi_H \circ h(v)$ for all $v \in \text{dom}(\psi_G)$.

PROOF. We verify that the map

$$\hat{h}: (G \oplus_{\psi_G} F) - x \rightarrow (H \oplus_{\psi_H} F) - y, \quad v \mapsto \begin{cases} h(v) & \text{if } v \in G - x, \text{ and} \\ \pi(v) & \text{if } v \in F \end{cases}$$

is a well-defined isomorphism. It is well-defined, since if $v \sim \psi_G(v)$ in $G \oplus_{\psi_G} F$, then $\hat{h}(v) \sim \hat{h}(\psi_G(v))$ in $H \oplus_{\psi_H} F$ by assumption on π . Moreover, since h and π are isomorphisms, it follows that \hat{h} is an isomorphism, too. \square

For the remainder of this section, all graphs are assumed to be locally finite. A *ray* in a graph G is a one-way infinite path. Given a ray R , then for any finite vertex set $S \subseteq V(G)$ there is a unique component $C(R, S)$ of $G - S$ containing a tail of R . An *end* in a graph is an equivalence class of rays under the relation

$$R \sim R' \Leftrightarrow \text{for every finite vertex set } S \subseteq V(G) \text{ we have } C(R, S) = C(R', S).$$

We denote by $\Omega(G)$ the set of ends in the graph G , and write $C(\omega, S) := C(R, S)$ with $R \in \omega$. Let $\Omega(\omega, S) = \{\omega' : C(\omega', S) = C(\omega, S)\}$. The singletons $\{v\}$ for $v \in V(G)$ and sets of the form $C(\omega, S) \cup \Omega(\omega, S)$ generate a compact metrizable topology on the set $V(G) \cup \Omega(G)$, which is known in the literature as $|G|$.² This topology allows us to talk about the closure of a set of vertices $X \subseteq V(G)$, denoted by \bar{X} . Write $\partial(X) = \bar{X} \setminus X = \bar{X} \cap \Omega(G)$ for the boundary of X : the collection of all ends in the closure of X . Then an end $\omega \in \Omega(G)$ lies in $\partial(X)$ if and only if for every finite vertex set $S \subseteq V(G)$, we have $|X \cap C(\omega, S)| = \infty$. Therefore $\Omega(G) = \partial(X)$ if and only if for every finite vertex set $S \subseteq V(G)$, every infinite component of $G - S$ meets X infinitely often. In this case we say that X is *dense* for $\Omega(G)$.

Finally, an end $\omega \in \Omega(G)$ is *free* if for some S , the set $\Omega(\omega, S) = \{\omega\}$. Then $\Omega'(G)$ denotes the *non-free* (or limit-)ends. Note that $\Omega'(G)$ is a closed subset of $\Omega(G)$.

LEMMA 3.4.3. *For locally finite connected graphs G and F , consider the gluing sum $G \oplus_{\psi} F$ for a partial bijection ψ . If F is one-ended and $\text{dom}(\psi)$ is infinite, then $\Omega(G \oplus_{\psi} F) \cong \Omega(G) / \partial(\text{dom}(\psi))$.*

²Normally $|G|$ is defined on the 1-complex of G together with its ends, but for our purposes it will be enough to just consider the subspace $V(G) \cup \Omega(G)$. See the survey paper of Diestel [53] for further details.

PROOF. Note first that for locally finite graphs G and F , also $G \oplus_\psi F$ is locally finite. Observe further that all rays of the unique end of F are still equivalent in $G \oplus_\psi F$, and so $G \oplus_\psi F$ has an end $\hat{\omega}$ containing the single end of F .

We are going to define a continuous surjection $f: \Omega(G) \rightarrow \Omega(G \oplus_\psi F)$ with the property that f has precisely one non-trivial fibre, namely $f^{-1}(\hat{\omega}) = \partial(\text{dom}(\psi))$. It then follows from definition of the quotient topology that f induces a continuous bijection from the compact space $\Omega(G)/\partial(\text{dom}(\psi))$ to the Hausdorff space $\Omega(G \oplus_\psi F)$, which, as such, is necessarily a homeomorphism.

The mapping f is defined as follows. Given an end $\omega \in \Omega(G) \setminus \partial(\text{dom}(\psi))$, there is a finite $S \subseteq V(G)$ such that $C(\omega, S) \cap \text{dom}(\psi) = \emptyset$, and so $C = C(\omega, S)$ is also a component of $(G \oplus_\psi F) - S$, which is disjoint from F . Define f to be the identity between $\Omega(G) \cap \overline{C}$ and $\Omega(G \oplus_\psi F) \cap \overline{C}$, while for all remaining ends $\omega \in \Omega(G) \cap \overline{\text{dom}(\psi)}$, we put $f(\omega) = \hat{\omega}$.

To see that this assignment is continuous at $\omega \in \Omega(G) \cap \overline{\text{dom}(\psi)}$, it suffices to show that $C := C(\omega, S) \subseteq G - S$ is a subset of $C' := C(\hat{\omega}, S) \subseteq (G \oplus_\psi F) - S$ for any finite set $S \subseteq G \oplus_\psi F$. To see this inclusion, note that by choice of ω , we have $|\text{dom}(\psi) \cap C| = \infty$. At the same time, since F is both one-ended and locally finite, $F - S$ has precisely one infinite component D and $F - D$ is finite, so as ψ is a bijection, there is $v \in \text{dom}(\psi) \cap C$ with $\psi(v) \in D$ (in fact, there are infinitely many such v). Since v and $\psi(v)$ get identified in $G \oplus_\psi F$, we conclude that $C \cup D$ is connected in $(G \oplus_\psi F) - S$, and hence that $C \cup D \subseteq C'$ as desired.

Finally, to see that f is indeed surjective, note first that the fact that $\text{dom}(\psi)$ is infinite implies that $\overline{\text{dom}(\psi)} \cap \Omega(G) \neq \emptyset$, and so $\hat{\omega} \in \text{ran}(f)$. Next, consider an end $\omega \in \Omega(G \oplus_\psi F)$ different from $\hat{\omega}$. Find a finite separator $S \subseteq V(G \oplus_\psi F)$ such that $C(\omega, S) \neq C(\hat{\omega}, S)$. It follows that $\text{dom}(\psi) \cap C(\omega, S)$ is finite. So there is a finite $S' \supseteq S$ such that $C := C(\omega, S') \neq C(\hat{\omega}, S')$ and $\text{dom}(\psi) \cap C = \emptyset$. So by definition, f is a bijection between $\Omega(G) \cap \overline{C}$ and $\Omega(G \oplus_\psi F) \cap \overline{C}$, so $\omega \in \text{ran}(f)$. \square

COROLLARY 3.4.4. *Under the above assumptions, if $\text{dom}(\psi)$ is dense for $\Omega(G)$, then $G \oplus_\psi F$ is one-ended.*

COROLLARY 3.4.5. *Under the above assumptions, if $\overline{\text{dom}(\psi)} \cap \Omega(G) = \Omega(G)$, then $G \oplus_\psi F$ has at most one non-free end.*

We remark that more direct proofs for Corollaries 3.4.4 and 3.4.5 can be given that do not need the full power of Lemma 3.4.3.

3.5. The construction

3.5.1. Preliminary definitions. In the precise statement of our construction in §3.5.2, we are going to employ the following notation.

DEFINITION 3.5.1 (Mii-path). A path $P = v_0, v_1, \dots, v_n$ in a graph G is called *internally isolated* if $\deg_G(v_i) = 2$ for all internal vertices v_i for $0 < i < n$. The path P is *maximal internally isolated* (or *mii* for short) if in addition $\deg_G(v_0) \neq 2 \neq \deg_G(v_n)$. An infinite path $P = v_0, v_1, v_2, \dots$ is *mii* if $\deg_G(v_0) \neq 2$ and $\deg_G(v_i) = 2$ for all $i \geq 1$.

DEFINITION 3.5.2 (Mii-spectrum). The *mii-spectrum* of G is

$$\Sigma(G) := \{k \in \mathbb{N} : G \text{ contains an mii-path of length } k\}.$$

If $\Sigma(G)$ is finite, we let $\sigma_0(G) = \max \Sigma(G)$ and $\sigma_1(G) = \max(\Sigma(G) \setminus \{\sigma_0(G)\})$.

LEMMA 3.5.3. *Let e be an edge of a locally finite graph G . If $\Sigma(G)$ is finite, then $\Sigma(G - e)$ is finite.*

PROOF. Observe first that every vertex of degree ≤ 2 in any graph can lie on at most one mii-path.

We now claim that for an edge $e = xy$, there are at most two finite mii-paths in $G - e$ which are not subpaths of finite mii-paths of G .

Indeed, if $\deg x = 3$ in G , then x can now be the interior vertex of one new finite mii-path in $G - e$. And if $\deg x = 2$ in G , then x can now be end-vertex of one new finite mii-path in $G - e$ (this is relevant if x lies on an infinite mii-path of G). The argument is for y is the same, so the claim follows. \square

DEFINITION 3.5.4 (Spectrally distinguishable). Given two graphs G and H , we say that G and H are *spectrally distinguishable* if there is some $k \geq 3$ such that $k \in \Sigma(G) \Delta \Sigma(H) = \Sigma(G) \setminus \Sigma(H) \cup \Sigma(H) \setminus \Sigma(G)$.

Note that being spectrally distinguishable is a strong certificate for being non-isomorphic.

DEFINITION 3.5.5 (k -ball). For G a subgraph of H , and $k > 0$, the k -ball $\text{Ball}_H(G, k)$ is the induced subgraph of H on the set of vertices at distance at most k of some vertex of G .

DEFINITION 3.5.6 (proper Mii-extension; infinite growth). Let G be a graph, B a subset of leaves of G , and H a component of G .

- A graph $\hat{G} \supset H$ is an *mii-extension* of H at B to length k if $\text{Ball}_{\hat{G}}(H, k)$ can be obtained from H by adjoining, at each vertex $l \in B \cap V(H)$, a new path of length k starting at l , and a new leaf whose only neighbour is l .³
- A leaf l in a graph G is *proper* if the unique neighbour of l in G has degree ≥ 3 . An mii-extension is called *proper* if every leaf in B is proper.
- An mii-extension \hat{G} of G is *of infinite growth* if every component of $\hat{G} - G$ is infinite.

³We note that this is a slightly different definition of an mii-extension to that in [36].

3.5.2. The back-and-forth construction. Our aim in this section is to prove our main theorem announced in the introduction.

THEOREM 2.1.6. There are two hypomorphic connected one-ended infinite graphs G and H with maximum degree five such that G is not isomorphic to H .

To do this we shall recursively construct, for each $n \in \mathbb{N}$,

- disjoint rooted connected graphs G_n and H_n ,
- disjoint sets R_n and B_n of proper leaves of the graph $G_n \cup H_n$,
- trees F_n ,
- disjoint sets R'_n and B'_n of leaves of F_n ,
- bijections $\psi_{G_n}: V(G_n) \cap (R_n \cup B_n) \rightarrow R'_n \cup B'_n$ and $\psi_{H_n}: V(H_n) \cap (R_n \cup B_n) \rightarrow R'_n \cup B'_n$,
- finite sets $X_n \subseteq V(G_n)$ and $Y_n \subseteq V(H_n)$, and bijections $\varphi_n: X_n \rightarrow Y_n$,
- a family of isomorphisms $\mathcal{H}_n = \{h_{n,x}: G_n - x \rightarrow H_n - \varphi_n(x): x \in X_n\}$,
- a family of automorphisms $\Pi_n = \{\pi_{n,x}: F_n \rightarrow F_n: x \in X_n\}$,
- a strictly increasing sequence of integers $k_n \geq 2$,

such that for all $n \in \mathbb{N}$:⁴

- (†1) $G_{n-1} \subseteq G_n$ and $H_{n-1} \subseteq H_n$ as induced subgraphs,
- (†2) the vertices of G_n and H_n all have degree at most 5,
- (†3) the vertices of F_n all have degree at most 3,
- (†4) the root of G_n is in R_n and the root of H_n is in B_n ,
- (†5) $\sigma_0(G_n) = \sigma_0(H_n) = k_n$,
- (†6) G_n and H_n are spectrally distinguishable,
- (†7) G_n and H_n have at most one end,
- (†8) $\Omega(G_n \cup H_n) \subseteq \overline{R_n \cup B_n}$,
- (†9) (a) G_n is a (proper) mii-extension of infinite growth of G_{n-1} at $R_{n-1} \cup B_{n-1}$ to length $k_{n-1} + 1$, and
(b) $\text{Ball}_{G_n}(G_{n-1}, k_{n-1} + 1)$ does not meet $R_n \cup B_n$,
- (†10) (a) H_n is a (proper) mii-extension of infinite growth of H_{n-1} at $R_{n-1} \cup B_{n-1}$ to length $k_{n-1} + 1$, and
(b) $\text{Ball}_{H_n}(H_{n-1}, k_{n-1} + 1)$ does not meet $R_n \cup B_n$,
- (†11) there are enumerations $V(G_n) = \{t_j: j \in J_n\}$ and $V(H_n) = \{s_j: j \in J_n\}$ such that
 - $J_{n-1} \subseteq J_n \subseteq \mathbb{N}$,
 - $\{t_j: j \in J_n\}$ extends the enumeration $\{t_j: j \in J_{n-1}\}$ of $V(G_{n-1})$, and similarly for $\{s_j: j \in J_n\}$,
 - $|\mathbb{N} \setminus J_n| = \infty$,

⁴If the statement involves an object indexed by $n - 1$ we only require that it holds for $n \geq 1$.

- $\{0, 1, \dots, n\} \subseteq J_n$,
- (†12) $\{t_j, s_j : j \leq n\} \cap (R_n \cup B_n) = \emptyset$,
- (†13) the finite sets of vertices X_n and Y_n satisfy $|X_n| = n = |Y_n|$, and
- $X_{n-1} \subseteq X_n$ and $Y_{n-1} \subseteq Y_n$,
 - $\varphi_n \upharpoonright X_{n-1} = \varphi_{n-1}$,
 - $\{t_j : j \leq \lfloor (n-1)/2 \rfloor\} \subseteq X_n$ and $\{s_j : j \leq \lfloor n/2 \rfloor - 1\} \subseteq Y_n$,
 - $(X_n \cup Y_n) \cap (R_n \cup B_n) = \emptyset$,
- (†14) the families of isomorphisms \mathcal{H}_n satisfy
- $h_{n,x} \upharpoonright (G_{n-1} - x) = h_{n-1,x}$ for all $x \in X_{n-1}$,
 - the image of $R_n \cap V(G_n)$ under $h_{n,x}$ is $R_n \cap V(H_n)$,
 - the image of $B_n \cap V(G_n)$ under $h_{n,x}$ is $B_n \cap V(H_n)$ for all $x \in X_n$.
- (†15) the families of automorphisms Π_n satisfy
- $\pi_{n,x} \upharpoonright R'_n$ is a permutation of R'_n for each $x \in X_n$,
 - $\pi_{n,x} \upharpoonright B'_n$ is a permutation of B'_n for each $x \in X_n$,
 - for each $x \in X_n$, the following diagram commutes:

$$\begin{array}{ccc}
 \mathcal{L}(G_n) & \xrightarrow{h_{n,x} \upharpoonright \mathcal{L}(G_n)} & \mathcal{L}(H_n) \\
 \psi_{G_n} \downarrow & & \downarrow \psi_{H_n} \\
 \mathcal{L}(F_n) & \xrightarrow{\pi_{n,x} \upharpoonright \mathcal{L}(F_n)} & \mathcal{L}(F_n)
 \end{array}$$

I.e. for every $\ell \in \mathcal{L}(G_n) := V(G_n) \cap (R_n \cup B_n)$ we have $\pi_{n,x}(\psi_{G_n}(\ell)) = \psi_{H_n}(h_{n,x}(\ell))$.

3.5.3. The construction yields the desired non-reconstructible one-ended graphs. By property (†1), we have $G_0 \subseteq G_1 \subseteq G_2 \subseteq \dots$ and $H_0 \subseteq H_1 \subseteq H_2 \subseteq \dots$. Let G and H be the union of the respective sequences. Then both G and H are connected, and as a consequence of (†2), both graphs have maximum degree 5.

We claim that the map $\varphi = \bigcup_n \varphi_n$ is a hypomorphism between G and H . Indeed, it follows from (†11) and (†13) that φ is a well-defined bijection from $V(G)$ to $V(H)$. To see that φ is a hypomorphism, consider any vertex x of G . This vertex appears as some t_j in our enumeration of $V(G)$, so the map

$$h_x = \bigcup_{n > 2j} h_{n,x} : G - x \rightarrow H - \varphi(x),$$

is a well-defined isomorphism by (†14) between $G - x$ and $H - \varphi(x)$.

Now suppose for a contradiction that there exists an isomorphism $f : G \rightarrow H$. Then $f(t_0)$ is mapped into H_n for some $n \in \mathbb{N}$. Properties (†5) and (†9) imply that after deleting all mii-paths in G of length $> k_n$, the connected component C of t_0 is a leaf extension of G_n adding one further leaf to every vertex in $V(G_n) \cap (R_n \cup B_n)$. Similarly, properties (†5) and (†10) imply that after deleting all mii-paths in H of length $> k_n$, the connected component D of $f(t_0)$ is a leaf-extension of H_n adding one further leaf to every

vertex in $V(H_n) \cap (R_n \cup B_n)$. Note that f restricts to an isomorphism between C and D . However, since C and D are proper extensions, we have $\Sigma(C) \Delta \Sigma(G_n) \subseteq \{1, 2\}$ and $\Sigma(D) \Delta \Sigma(H_n) \subseteq \{1, 2\}$. Hence, since G_n and H_n are spectrally distinguishable by $(\dagger 6)$, so are C and D , a contradiction. We have established that G and H are non-isomorphic reconstructions of each other.

Finally, for G being one-ended, we now show that for every finite vertex separator $S \subseteq V(G)$, the graph $G - S$ has only one infinite component (the argument for H is similar). Suppose for a contradiction $G - S$ has two infinite components C_1 and C_2 . Consider n large enough such that $S \subseteq V(G_n)$. Since G_k is one-ended for all k by $(\dagger 7)$, we may assume that $C_1 \cap G_k$ falls apart into finite components for all $k \geq n$. Since C_1 is infinite and connected, it follows from $(\dagger 9)(b)$ that C_1 intersects $G_{n+1} - G_n$. But since G_{n+1} is an mii-extension of G_n of infinite growth by $(\dagger 9)(a)$, we see that that $C_1 \cap (G_{n+1} - G_n)$ contains an infinite component, a contradiction.

3.5.4. The base case: there are finite rooted graphs G_0 and H_0 satisfying requirements $(\dagger 1)$ – $(\dagger 15)$. Choose a pair of spectrally distinguishable, equally sized graphs G_0 and H_0 with maximum degree ≤ 5 and $\sigma_0(G_0) = \sigma_0(H_0) = k_0$. Pick a proper leaf each as roots $r(G_0)$ and $r(H_0)$ for G_0 and H_0 , and further proper leaves $\ell_b \in G_0$ and $\ell_r \in H_0$.




FIGURE 3.5. A possible choice for the finite rooted graphs G_0 and H_0 .

Define $R_0 = \{r(G_0), \ell_r\}$ and $B_0 = \{r(H_0), \ell_b\}$. We take F_0 to be two vertices x and y joined by an edge, with $R'_0 = \{x\}$ and $B'_0 = \{y\}$ and take ψ_{G_0} to be the unique bijection sending $R_0 \cap G_0$ to R'_0 and $B_0 \cap G_0$ to B'_0 , and similarly for ψ_{H_0} .



FIGURE 3.6. F_0 .

Let $J_0 = \{0, 1, \dots, |G_0| - 1\}$ and choose enumerations $V(G_0) = \{t_j : j \in J_0\}$ and $V(H_0) = \{s_j : j \in J_0\}$ with $t_0 \neq r(G_0)$ and $s_0 \neq r(H_0)$. Finally we let $X_0 = Y_0 = \mathcal{H}_0 = \emptyset$. It is a simple check that conditions $(\dagger 1)$ – $(\dagger 15)$ are satisfied.

3.5.5. The inductive step: set-up. Now, assume that we have constructed graphs G_k and H_k for all $k \leq n$ such that $(\dagger 1)$ – $(\dagger 15)$ are satisfied up to n . If $n = 2m$ is even, then we have $\{t_j : j \leq m - 1\} \subseteq X_n$ and in order to satisfy $(\dagger 13)$ we have to construct G_{n+1} and H_{n+1} such that the vertex t_m is taken care of in our partial hypomorphism. Similarly,

if $n = 2m + 1$ is odd, then we have $\{s_j : j \leq m - 1\} \subseteq Y_n$ and we have to construct G_{n+1} and H_{n+1} such that the vertex s_m is taken care of in our partial hypomorphism. Both cases are symmetric, so let us assume in the following that $n = 2m$ is even.

Now let v be the vertex with the least index in the set $\{t_j : j \in J_n\} \setminus X_n$, i.e.

$$(14) \quad v = t_i \text{ for } i = \min \{j : t_j \in V(G_n) \setminus X_n\}.$$

Then by assumption (†13), v will be t_m , unless t_m was already in X_n anyway. In any case, since $|X_n| = |Y_n| = n$, it follows from (†11) that $i \leq n$, so by (†12), v does not lie in our leaf sets $R_n \cup B_n$, i.e.

$$(15) \quad v \notin R_n \cup B_n.$$

In the next sections, we will demonstrate how to obtain graphs $G_{n+1} \supset G_n$, $H_{n+1} \supset H_n$ and F_{n+1} with $X_{n+1} = X_n \cup \{v\}$ and $Y_{n+1} = Y_n \cup \{\varphi_{n+1}(v)\}$ satisfying (†1)–(†10) and (†13)–(†15).

After we have completed this step, since $|\mathbb{N} \setminus J_n| = \infty$, it is clear that we can extend our enumerations of G_n and H_n to enumerations of G_{n+1} and H_{n+1} as required, making sure to first list some new elements that do not lie in $R_{n+1} \cup B_{n+1}$. This takes care of (†11) and (†12) and completes the step $n \mapsto n + 1$.

3.5.6. The inductive step: construction. We will construct the graphs G_{n+1} and H_{n+1} in three steps. First, in §3.5.6.1 we construct graphs $G'_{n+1} \supset G_n$ and $H'_{n+1} \supset H_n$ such that there is a vertex $\varphi_{n+1}(v) \in H'_{n+1}$ with $G'_{n+1} - v \cong H'_{n+1} - \varphi_{n+1}(v)$. This first step essentially follows the argument from [36, §4.6]. We will also construct a graph F_{n+1} via a parallel process.

Secondly, in §3.5.6.2 we will show that there are well-behaved maps from the coloured leaves of G'_{n+1} and H'_{n+1} to $F_{n+1} \times \mathbb{N}$, such that analogues of (†14) and (†15) hold for G'_{n+1} , H'_{n+1} and F_{n+1} , giving us control over the corresponding gluing sum.

Lastly, in §3.5.6.3, we do the actual gluing process and define all objects needed for step $n + 1$ of our inductive construction.

3.5.6.1. *Building the auxiliary graphs.* Given the two graphs G_n and H_n , we extend each of them through their roots as indicated in Figure 3.7 to graphs \tilde{G}_n and \tilde{H}_n respectively.

Since v is not the root of G_n , there is a unique component of $G_n - v$ containing the root, which we call $G_n(r)$. Let $G_n(v)$ be the induced subgraph of G_n on the remaining vertices, including v . We remark that if v is not a cutvertex of G_n , then $G_n(v)$ is just a single vertex v . Since $\sigma_0(G_n) = k_n$ by (†5) and $\deg(v) \leq 5$ by (†2), it follows from an iterative application of Lemma 3.5.3 that $\Sigma(G_n(r))$ and $\Sigma(G_n(v))$ are finite. Let $k = \tilde{k}_n = \max\{\sigma_0(G_n), \sigma_0(G_n(r)), \sigma_0(G_n(v)), \sigma_0(H_n)\} + 1$.

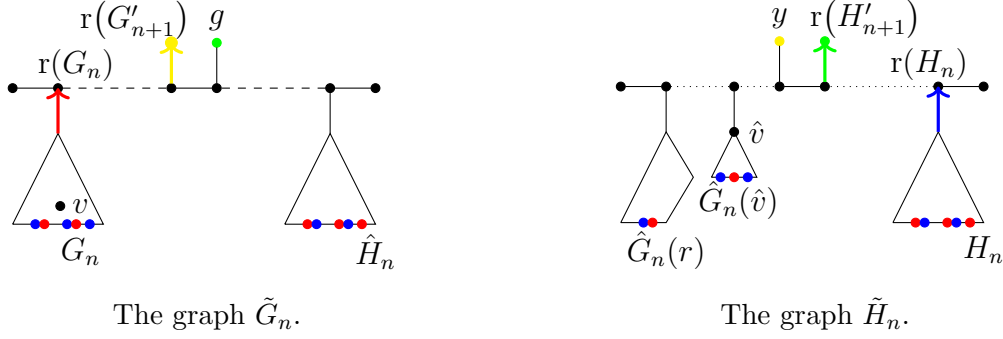


FIGURE 3.7. All dotted lines are mii-paths of length at least $k + 1 = \tilde{k}_n + 1$.

To obtain \tilde{G}_n , we extend G_n through its root $r(G_n) \in R_n$ by a path

$$r(G_n) = u_0, u_1, \dots, u_{p-1}, u_p = r(\hat{H}_n)$$

of length $p = 4(\tilde{k}_n + 1) + 1$, where at its last vertex u_p we glue a rooted copy \hat{H}_n of H_n (via an isomorphism $\hat{z} \leftrightarrow z$), identifying u_p with the root of \hat{H}_n .

Next, we add two additional leaves at u_0 and u_p , so that $\deg(r(G_n)) = 3 = \deg(r(\hat{H}_n))$. Further, we add a leaf $r(G'_{n+1})$ at u_{2k+2} , which will be our new root for the next tree G'_{n+1} ; and another leaf g at u_{2k+3} . This completes the construction of \tilde{G}_n .

The construction of \tilde{H}_n is similar, but not entirely symmetric. For its construction, we extend H_n through its root $r(H_n) \in B_n$ by a path

$$r(H_n) = v_p, v_{p-1}, \dots, v_1, v_0 = r(\hat{G}_n(r))$$

of length p , where at its last vertex v_0 we glue a copy $\hat{G}_n(r)$ of $G_n(r)$, identifying v_0 with the root of $\hat{G}_n(r)$. Then, we take a copy $\hat{G}_n(\hat{v})$ of $G_n(v)$ and connect \hat{v} via an edge to v_{k+1} .

Finally, as before, we add two leaves at v_0 and v_p so that $\deg(r(\hat{G}_n(r))) = 3 = \deg(r(H_n))$. Next, we add a leaf $r(H'_{n+1})$ to v_{2k+3} , which will be our new root for the next tree H'_{n+1} ; and another leaf y to v_{2k+2} . This completes the construction of \tilde{H}_n .

By the induction assumption, certain leaves of G_n have been coloured with one of the two colours in $R_n \cup B_n$, and also some leaves of H_n have been coloured with one of the two colours in $R_n \cup B_n$. In the above construction, we colour leaves of \hat{H}_n , $\hat{G}_n(r)$ and $\hat{G}_n(\hat{v})$ accordingly:

$$(16) \quad \begin{aligned} \tilde{R}_n &= \left(R_n \cup \left\{ \hat{z} \in \hat{H}_n \cup \hat{G}_n(r) \cup \hat{G}_n(\hat{v}) : z \in R_n \right\} \right) \setminus \left\{ r(G_n), r(\hat{G}_n(r)) \right\}, \\ \tilde{B}_n &= \left(B_n \cup \left\{ \hat{z} \in \hat{H}_n \cup \hat{G}_n(r) \cup \hat{G}_n(\hat{v}) : z \in B_n \right\} \right) \setminus \left\{ r(H_n), r(\hat{H}_n) \right\}. \end{aligned}$$

Now put $M_n := \tilde{G}_n \cup \tilde{H}_n$ and consider the following promise structure $\mathcal{P} = (M_n, \vec{P}, \mathcal{L})$ on M_n , consisting of four promise edges $\vec{P} = \{\vec{p}_1, \vec{p}_2, \vec{p}_3, \vec{p}_4\}$ and corresponding leaf sets

$\mathcal{L} = \{L_1, L_2, L_3, L_4\}$, as follows:

- \vec{p}_1 pointing in G_n towards $r(G_n)$, with $L_1 = \tilde{R}_n$,
- \vec{p}_2 pointing in H_n towards $r(H_n)$, with $L_2 = \tilde{B}_n$,
- \vec{p}_3 pointing in \tilde{G}_n towards $r(G'_{n+1})$, with $L_3 = \{r(G'_{n+1}), y\}$,
- \vec{p}_4 pointing in \tilde{H}_n towards $r(H'_{n+1})$, with $L_4 = \{r(H'_{n+1}), g\}$.

Note that our construction so far has been tailored to provide us with a \vec{P} -respecting isomorphism

$$(18) \quad h: \tilde{G}_n - v \rightarrow \tilde{H}_n - \hat{v}.$$

Consider the closure $\text{cl}(M_n)$ with respect to the above defined promise structure \mathcal{P} . Since $\text{cl}(M_n)$ is a leaf-extension of M_n , it has two connected components, just as M_n . We now define

$$(19) \quad \begin{aligned} G'_{n+1} &= \text{the component containing } G_n \text{ in } \text{cl}(M_n), \\ H'_{n+1} &= \text{the component containing } H_n \text{ in } \text{cl}(M_n). \end{aligned}$$

It follows that $\text{cl}(M_n) = G'_{n+1} \cup H'_{n+1}$. Further, since \vec{p}_3 and \vec{p}_4 are placeholder promises, $\text{cl}(M_n)$ carries a corresponding promise structure, cf. Def. 3.3.5. We define

$$(20) \quad R_{n+1} = \text{cl}(L_3) \quad \text{and} \quad B_{n+1} = \text{cl}(L_4).$$

Lastly, set

$$(21) \quad \begin{aligned} X_{n+1} &= X_n \cup \{v\}, \\ Y_{n+1} &= Y_n \cup \{\hat{v}\}, \\ \varphi_{n+1} &= \varphi_n \cup \{(v, \hat{v})\}, \\ k_{n+1} &= 2(\tilde{k}_n + 1). \end{aligned}$$

We now build F_{n+1} in a similar fashion to the above procedure. That is, we take two copies of F_n and join them pairwise through their roots as indicated in Figure 3.7 to form a graph \tilde{F}_n . We consider the graph $N_n = \tilde{F}_n \cup \hat{\tilde{F}}_n$, and take F_{n+1} to be one of the components of $\text{cl}(N_n)$ (unlike for $\text{cl}(M_n)$, both components of $\text{cl}(N_n)$ are isomorphic).




FIGURE 3.8. The graph $N_n = \tilde{F}_n \cup \hat{\tilde{F}}_n$.

More precisely we take two copies of F_n , which we will denote by F_n^G and F_n^H . We extend F_n^G through the image of the $r(G_n)$ under the bijection ψ_{G_n} by a path

$$\psi_{G_n}(r(G_n)) = u_0, u_1, u_2, u_3 = \psi_{H_n}(r(H_n))$$

of length three, where $\psi_{G_n}(r(G_n))$ is taken in F_n^G and $\psi_{H_n}(r(H_n))$ is taken in F_n^H . Further, we add a leaf x at u_1 , and another leaf y at u_2 . We will consider the graph $N_n = \tilde{F}_n \cup \hat{F}_n$ as in Figure 3.8 formed by taking two disjoint copies of \tilde{F}_n .

By the induction assumption, certain leaves of F_n have been coloured with one of the two colours in $R'_n \cup B'_n$. In the above construction, we colour leaves of $F_n^G, F_n^H, \hat{F}_n^G$ and \hat{F}_n^H accordingly:

$$(22) \quad \begin{aligned} \tilde{R}'_n &= \left\{ w \in F_n^G \cup F_n^H \cup \hat{F}_n^G \cup \hat{F}_n^H : w \in R'_n \right\} \setminus \left\{ \psi_{G_n}(r(G_n)), \psi_{G_n}(\widehat{r(G_n)}) \right\} \\ \tilde{B}'_n &= \left\{ w \in F_n^G \cup F_n^H \cup \hat{F}_n^G \cup \hat{F}_n^H : w \in B'_n \right\} \setminus \left\{ \psi_{H_n}(r(H_n)), \psi_{H_n}(\widehat{r(H_n)}) \right\}. \end{aligned}$$

Now consider the following promise structure $\mathcal{P}' = (N_n, \vec{P}', \mathcal{L}')$ on N_n , consisting of four promise edges $\vec{P}' = \{\vec{r}_1, \vec{r}_2, \vec{r}_3, \vec{r}_4\}$ and corresponding leaf sets $\mathcal{L}' = \{L'_1, L'_2, L'_3, L'_4\}$, as follows:

$$(23) \quad \begin{aligned} &\bullet \vec{r}_1 \text{ pointing in } F_n^G \text{ towards } \psi_{G_n}(r(G_n)), \text{ with } L'_1 = \tilde{R}'_n, \\ &\bullet \vec{r}_2 \text{ pointing in } \hat{F}_n^H \text{ towards } \psi_{H_n}(r(H_n)), \text{ with } L'_2 = \tilde{B}'_n, \\ &\bullet \vec{r}_3 \text{ pointing in } \tilde{F}_n \text{ towards } x, \text{ with } L'_3 = \{x, \hat{x}\}, \\ &\bullet \vec{r}_4 \text{ pointing in } \hat{F}_n \text{ towards } \hat{y}, \text{ with } L'_4 = \{y, \hat{y}\}. \end{aligned}$$

Consider the closure $\text{cl}(N_n)$ with respect to the promise structure \mathcal{P}' defined above. Since $\text{cl}(N_n)$ is a leaf-extension of N_n , it has two connected components, and we define F_{n+1} to be the component containing F_n^G in $\text{cl}(N_n)$. Since \vec{r}_3 and \vec{r}_4 are placeholder promises, $\text{cl}(N_n)$ carries a corresponding promise structure, cf. Def. 3.3.5. We define

$$(24) \quad R'_{n+1} = \text{cl}(L'_3) \cap F_{n+1} \quad \text{and} \quad B'_{n+1} = \text{cl}(L'_4) \cap F_{n+1}.$$

3.5.6.2. *Extending maps.* In order to glue $F_{n+1} \times \mathbb{N}$ onto G'_{n+1} and H'_{n+1} we will need to show that that analogues of (†14) and (†15) hold for G'_{n+1}, H'_{n+1} and F_{n+1} . Our next lemma is essentially [36, Claim 4.13], and is an analogue of (†14). We briefly remind the reader of the details, as we need to know the nature of our extensions in our later claims.

LEMMA 3.5.7. *There is a family of isomorphisms $\mathcal{H}'_{n+1} = \{h'_{n+1,x} : x \in X_{n+1}\}$ witnessing that $G'_{n+1} - x$ and $H'_{n+1} - \varphi_{n+1}(x)$ are isomorphic for all $x \in X_{n+1}$, such that $h'_{n+1,x}$ extends $h_{n,x}$ for all $x \in X_n$.*

PROOF. The graphs G'_{n+1} and H'_{n+1} defined in (19) are obtained from \tilde{G}_n and \tilde{H}_n by attaching at every leaf in \tilde{R}_n a copy of the rooted graph $\text{cl}(M_n)(\vec{p}_1)$, and by attaching at every leaf in \tilde{B}_n a copy of the rooted graph $\text{cl}(M_n)(\vec{p}_2)$ by (cl.2).

From (18) we know that there is a \vec{P} -respecting isomorphism

$$h: \tilde{G}_n - v \rightarrow \tilde{H}_n - \varphi_{n+1}(v).$$

In other words, h maps promise leaves in $L_i \cap V(\tilde{G}_n)$ bijectively to the promise leaves in $L_i \cap V(\tilde{H}_n)$ for all $i = 1, 2, 3, 4$.

There is for each $\ell \in \tilde{R}_n \cup \tilde{B}_n \cup \{r(G_n), r(H_n)\}$ a $\text{cl}(\vec{P})$ -respecting isomorphism of rooted graphs

$$(25) \quad f_\ell: \text{cl}(M_n)(\vec{q}_\ell) \cong \text{cl}(M_n)(\vec{p}_i)$$

given by (cl.3) for $\ell \in (\tilde{R}_n \cup \tilde{B}_n)$, where i equals blue or red depending on whether $\ell \in \tilde{R}_n$ or \tilde{B}_n , and for the roots of G_n and H_n we have $\vec{q}_r = \vec{p}_i$ and the isomorphism is the identity. Hence, for each ℓ ,

$$f_{h(\ell)}^{-1} \circ f_\ell: \text{cl}(M_n)(\vec{q}_\ell) \cong \text{cl}(M_n)(\vec{q}_{h(\ell)})$$

is a $\text{cl}(\vec{P})$ -respecting isomorphism of rooted graphs. By combining the isomorphism h between $\tilde{G}_n - v$ and $\tilde{H}_n - \varphi_{n+1}(v)$ with these isomorphisms between each $\text{cl}(M_n)(\vec{q}_\ell)$ and $\text{cl}(M_n)(\vec{q}_{h(\ell)})$ we get a $\text{cl}(\vec{P})$ -respecting isomorphism

$$h'_{n+1,v}: G'_{n+1} - v \rightarrow H'_{n+1} - \varphi_{n+1}(v).$$

To extend the old isomorphisms $h_{n,x}$ (for $x \in X_n$), note that G'_{n+1} and H'_{n+1} are obtained from G_n and H_n by attaching at every leaf in R_n a copy of the rooted graph $\text{cl}(M_n)(\vec{p}_1)$, and similarly by attaching at every leaf in B_n a copy of the rooted graph $\text{cl}(M_n)(\vec{p}_2)$. By induction assumption (†14), for each $x \in X_n$ the isomorphism

$$h_{n,x}: G_n - x \rightarrow H_n - \varphi_n(x)$$

maps the red leaves of G_n bijectively to the red leaves of H_n , and the blue leaves of G_n bijectively to the blue leaves of H_n . Thus, by (25),

$$f_{h_{n,x}(\ell)}^{-1} \circ f_\ell: \text{cl}(M_n)(\vec{q}_\ell) \cong \text{cl}(M_n)(\vec{q}_{h_{n,x}(\ell)})$$

are $\text{cl}(\vec{P})$ -respecting isomorphisms of rooted graphs for all $\ell \in (R_n \cup B_n) \cap V(G_n)$. By combining the isomorphism $h_{n,x}$ between $G_n - x$ and $H_n - \varphi_n(x)$ with these isomorphisms between each $\text{cl}(M_n)(\vec{q}_\ell) = G'_{n+1}(\vec{q}_\ell)$ and $\text{cl}(M_n)(\vec{q}_{h_{n,x}(\ell)}) = H'_{n+1}(\vec{q}_{h_{n,x}(\ell)})$, we obtain a $\text{cl}(\vec{P})$ -respecting extension

$$h'_{n+1,x}: G'_{n+1} - x \rightarrow H'_{n+1} - \varphi_n(x). \quad \square$$

Our next claim should be seen as an approximation to property (†15). Recall that $\text{cl}(N_n)$ has two components $F_{n+1} \cong \hat{F}_{n+1}$.

LEMMA 3.5.8. *There are colour-preserving bijections*

$$\psi_{G'_{n+1}}: V(G'_{n+1}) \cap (R_{n+1} \cup B_{n+1}) \rightarrow R'_{n+1} \cup B'_{n+1},$$

$$\psi_{H'_{n+1}}: V(H'_{n+1}) \cap (R_{n+1} \cup B_{n+1}) \rightarrow \hat{R}'_{n+1} \cup \hat{B}'_{n+1},$$

and a family of isomorphisms

$$\hat{\Pi}_{n+1} = \left\{ \hat{\pi}_{n+1,x}: F_{n+1} \rightarrow \hat{F}_{n+1}: x \in X_{n+1} \right\}$$

such that for each $x \in X_{n+1}$ the following diagram commutes.

$$\begin{array}{ccc} \mathcal{L}(G'_{n+1}) & \xrightarrow{h'_{n+1,x} \upharpoonright \mathcal{L}(G'_{n+1})} & \mathcal{L}(H'_{n+1}) \\ \psi_{G'_{n+1}} \downarrow & & \downarrow \psi_{H'_{n+1}} \\ \mathcal{L}'(F_{n+1}) & \xrightarrow{\hat{\pi}_{n+1,x} \upharpoonright \mathcal{L}'(F_{n+1})} & \mathcal{L}'(\hat{F}_{n+1}) \end{array}$$

PROOF. *Defining $\psi_{G'_{n+1}}$ and $\psi_{H'_{n+1}}$.* By construction, we can combine the maps ψ_{G_n} and ψ_{H_n} to obtain a natural colour-preserving bijection

$$\psi: \mathcal{L}(M_n) \rightarrow \mathcal{L}'(N_n),$$

which satisfies the assumptions of Lemma 3.3.7. Thus, by Corollary 3.3.8, there are bijections

$$\alpha^i: \mathcal{L}(\text{cl}(M_n)(\vec{p}_i)) \rightarrow \mathcal{L}'(\text{cl}(N_n)(\vec{r}_i))$$

which are colour-preserving with respect to the promise structures $\text{cl}(\mathcal{P})$ and $\text{cl}(\mathcal{P}')$ on $\text{cl}(M_n)$ and $\text{cl}(N_n)$, respectively.

We now claim that ψ extends to a colour-preserving bijection (w.r.t. $\text{cl}(\mathcal{P})$)

$$\text{cl}(\psi): \mathcal{L}(\text{cl}(M_n)) \rightarrow \mathcal{L}'(\text{cl}(N_n)).$$

Indeed, by (cl.3), for every $\ell \in \tilde{R}'_n \cup \tilde{B}'_n$, there is a \vec{P}' -respecting rooted isomorphism

$$(26) \quad g_\ell: \text{cl}(N_n)(\vec{q}_\ell) \rightarrow \text{cl}(N_n)(\vec{r}_i),$$

where i equals blue or red depending on whether $\ell \in \tilde{R}'_n$ or \tilde{B}'_n . As in the case of (25) we define the maps g_r with $\vec{q}_r = \vec{r}_i$ for the roots of F_n^G and \hat{F}_n^H respectively to be the identity. Together with the rooted isomorphisms f_ℓ from (25), it follows that for each $\ell \in \tilde{R}_n \cup \tilde{B}_n \cup \{\text{r}(G_n), \text{r}(H_n)\}$, the map

$$\psi_\ell = g_{\psi(\ell)}^{-1} \circ \alpha^i \circ f_\ell: \mathcal{L}(\text{cl}(M_n)(\vec{q}_\ell)) \rightarrow \mathcal{L}(\text{cl}(N_n)(\vec{q}_{\psi(\ell)}))$$

is a colour-preserving bijection. Now combine ψ with the individual ψ_ℓ to obtain $\text{cl}(\psi)$. We then put

$$\psi_{G'_{n+1}} = \text{cl}(\psi) \upharpoonright G'_{n+1} \quad \text{and} \quad \psi_{H'_{n+1}} = \text{cl}(\psi) \upharpoonright H'_{n+1}.$$

Defining isomorphisms $\hat{\Pi}_{n+1}$. To extend the old isomorphisms $\pi_{n,x}$, given by the induction assumption, note that by (cl.2), F_{n+1} is obtained from F_n by attaching at every leaf in R'_n a copy of the rooted graph $F_{n+1}(\vec{r}_1)$, and similarly by attaching at every leaf in B'_n a copy of the rooted graph $F_{n+1}(\vec{r}_2)$. For each $x \in X_n$ let us write $\hat{\pi}_{n,x}$ for the map

sending each $z \in F_n^G$ to the copy of $\pi_{n,x}(z)$ in \hat{F}_n^H . By the induction assumption (†15), for each $x \in X_n$ the isomorphism

$$\hat{\pi}_{n,x}: F_n^G \rightarrow \hat{F}_n^H$$

preserves the colour of red and blue leaves. Thus, using the maps g_ℓ from (26), the mappings

$$g_{\hat{\pi}_{n,x}(\ell)}^{-1} \circ g_\ell: \text{cl}(N_n)(\vec{q}_\ell) \cong \text{cl}(N_n)(\vec{q}_{\hat{\pi}_{n,x}(\ell)})$$

are $\text{cl}(\vec{P}')$ -respecting isomorphisms of rooted graphs for all $\ell \in R'_n \cup B'_n$. By combining the isomorphism $\pi_{n,x}$ with these isomorphisms between each $F_{n+1}(\vec{q}_\ell)$ and $\hat{F}_{n+1}(\vec{q}_{\hat{\pi}_{n,x}(\ell)})$, we obtain a $\text{cl}(\vec{P}')$ -respecting extension

$$\hat{\pi}_{n+1,x}: F_{n+1} \rightarrow \hat{F}_{n+1}.$$

For the new isomorphism $\hat{\pi}_{n+1,v}: F_{n+1} \rightarrow \hat{F}_{n+1}$, we simply take the ‘identity’ map which extends the map sending each $z \in \tilde{F}_n$ to $\hat{z} \in \hat{F}_n$.

The diagram commutes. To see that the new diagram above commutes, for each $x \in X_n$ it suffices to check that for all $\ell \in (R_n \cup B_n) \cap V(G_n)$ we have

$$\hat{\pi}_{n+1,x} \circ \psi_{G'_{n+1}} \upharpoonright \mathcal{L}(G'_{n+1}(\vec{q}_\ell)) = \psi_{H'_{n+1}} \circ h'_{n+1,x} \upharpoonright \mathcal{L}(G'_{n+1}(\vec{q}_\ell)),$$

which by construction of $\text{cl}(\psi)$ above is equivalent to showing that

$$\hat{\pi}_{n+1,x} \circ \psi_\ell = \psi_{h_{n,x}(\ell)} \circ h'_{n+1,x}.$$

By definition of ψ_ℓ this holds if and only if

$$\hat{\pi}_{n+1,x} \circ g_{\psi(\ell)}^{-1} \circ \alpha^i \circ f_\ell = g_{\psi(h_{n,x}(\ell))}^{-1} \circ \alpha^i \circ f_{h_{n,x}(\ell)} \circ h'_{n+1,x}.$$

Now by construction of $\hat{\pi}_{n+1,x}$ and $h'_{n+1,x}$, we have

$$\hat{\pi}_{n+1,x} \circ g_{\psi(\ell)}^{-1} = g_{\hat{\pi}_{n,x}(\psi(\ell))}^{-1} \quad \text{and} \quad f_{h_{n,x}(\ell)} \circ h'_{n+1,x} = f_\ell.$$

Hence, the above is true if and only if

$$g_{\hat{\pi}_{n,x}(\psi(\ell))}^{-1} \circ \alpha^i \circ f_\ell = g_{\psi(h_{n,x}(\ell))}^{-1} \circ \alpha^i \circ f_\ell.$$

Finally, this last line holds since $\psi(\ell) = \psi_{G_n}(\ell)$ and $\psi(h_{n,x}(\ell)) = \psi_{H_n}(h_{n,x}(\ell))$ by definition of ψ , and because

$$\hat{\pi}_{n,x} \circ \psi_{G_n}(\ell) = \psi_{H_n} \circ h_{n,x}(\ell)$$

by the induction assumption.

For $\hat{\pi}_{n+1,v}$ we see that, as above, it will be sufficient to show that for all $\ell \in (\tilde{R}_n \cup \tilde{B}_n) \cap V(\tilde{G}_n)$ we have

$$\hat{\pi}_{n+1,v} \circ \psi_\ell = \psi_{h'_{n+1,v}(\ell)} \circ h'_{n+1,v},$$

which reduces as before to showing that,

$$g_{\hat{\pi}_{n+1,v}(\psi(\ell))}^{-1} \circ \alpha^i \circ f_\ell = g_{\psi(h'_{n+1,v}(\ell))}^{-1} \circ \alpha^i \circ f_\ell.$$

Recall that, $\hat{\pi}_{n+1,v}$ sends each v to \hat{v} and also, since $h'_{n+1,v} \upharpoonright \tilde{G}_n = h$, the image of every leaf $\ell \in (\tilde{R}_n \cup \tilde{B}_n) \cap V(\tilde{G}_n)$ is simply $\hat{\ell} \in \hat{G}_n(v) \cup \hat{G}_n(r)$. Hence we wish to show that

$$g_{(\psi(\hat{\ell}))}^{-1} \circ \alpha^i \circ f_\ell = g_{\psi(\hat{i})}^{-1} \circ \alpha^i \circ f_\ell,$$

that is,

$$(\psi(\hat{\ell})) = \psi(\hat{i}),$$

which follows from the construction of ψ . \square

3.5.6.3. Gluing the graphs together. Let us take the cartesian product of F_{n+1} with a ray, which we simply denote by $F_{n+1} \times \mathbb{N}$. If we identify F_{n+1} with the subgraph $F_{n+1} \times \{0\}$, then we can interpret both $\psi_{G'_{n+1}}$ and $\psi_{H'_{n+1}}$ as maps from $\mathcal{L}(G'_{n+1})$ and $\mathcal{L}(H'_{n+1})$ to a set of vertices in $F_{n+1} \times \mathbb{N}$, under the natural isomorphism between \hat{F}_{n+1} and F_{n+1} .

Instead of using the function $\psi_{G'_{n+1}}$ directly for our gluing operation, we identify, for every leaf l in $\mathcal{L}(G'_{n+1})$ the *unique neighbour* of l with $\psi_{G'_{n+1}}(l)$. Formally, define a bijection $\chi_{G_{n+1}}$ between the neighbours of $\mathcal{L}(G'_{n+1})$ and $\mathcal{L}'(F_{n+1})$ via

$$(27) \quad \chi_{G_{n+1}} = \left\{ (z_1, z_2) : \exists l \in \mathcal{L}(G'_{n+1}) \text{ s.t. } z_1 \in N(l) \text{ and } \psi_{G'_{n+1}}(l) = z_2 \right\},$$

and similarly

$$(28) \quad \chi_{H_{n+1}} = \left\{ (z_1, z_2) : \exists l \in \mathcal{L}(H'_{n+1}) \text{ s.t. } z_1 \in N(l) \text{ and } \psi_{H'_{n+1}}(l) = z_2 \right\}.$$

Since two promise leaves in G'_{n+1} or H'_{n+1} are never adjacent to the same vertex, $\chi_{G_{n+1}}$ and $\chi_{H_{n+1}}$ are indeed bijections. Moreover, since all promise leaves were proper, the vertices in the domain of $\chi_{G_{n+1}}$ and $\chi_{H_{n+1}}$ have degree at least 3. Using our notion of gluing-sum (see Def. 3.4.1), we now define

$$(29) \quad G_{n+1} := G'_{n+1} \oplus_{\chi_{G_{n+1}}} (F_{n+1} \times \mathbb{N}) \text{ and } H_{n+1} := H'_{n+1} \oplus_{\chi_{H_{n+1}}} (F_{n+1} \times \mathbb{N}).$$

We consider R_{n+1} , B_{n+1} , X_{n+1} and Y_{n+1} as subsets of G_{n+1} and H_{n+1} in the natural way. Then $\psi_{G_{n+1}}$ and $\psi_{H_{n+1}}$ can be taken to be the maps $\psi_{G'_{n+1}}$ and $\psi_{H'_{n+1}}$, again identifying \hat{F}_{n+1} with F_{n+1} in the natural way. We also take the roots of G_{n+1} and H_{n+1} to be the roots of G'_{n+1} and H'_{n+1} respectively

This completes the construction of graphs G_{n+1} , H_{n+1} , and F_{n+1} , the coloured leaf sets R_{n+1} , B_{n+1} , R'_{n+1} , and B'_{n+1} , the bijections $\psi_{G_{n+1}}$ and $\psi_{H_{n+1}}$, as well as $\varphi_{n+1} : X_{n+1} \rightarrow Y_{n+1}$, and $k_{n+1} = 2(\hat{k}_n + 1)$. In the next section, we show the existence of families of isomorphisms \mathcal{H}_{n+1} and Π_{n+1} , and verify that (†1)–(†15) are indeed satisfied for the $(n+1)^{\text{th}}$ instance.

3.5.7. The inductive step: verification.

LEMMA 3.5.9. *We have $G_n \subseteq G_{n+1}$, $H_n \subseteq H_{n+1}$, $\Delta(G_{n+1}), \Delta(H_{n+1}) \leq 5$, $\Delta(F_{n+1}) \leq 3$, and the roots of G_{n+1} and H_{n+1} are in R_{n+1} and B_{n+1} respectively.*

PROOF. We note that $G_n \subseteq G'_{n+1}$ by construction. Hence, it follows that

$$G_n \subseteq G'_{n+1} \subseteq G'_{n+1} \oplus_{\chi_{G_{n+1}}} (F_{n+1} \times \mathbb{N}) = G_{n+1},$$

and similarly for H_n . Since we glued together degree 3 and degree 2 vertices, and $\Delta(G_n), \Delta(H_n) \leq 5$ and $\Delta(F_n) \leq 3$, it is clear that the same bounds hold for $n + 1$. Finally, since the root of \tilde{G}_n was a placeholder promise, and R_{n+1} was the corresponding set of promise leaves in $\text{cl}(\tilde{G}_n)$, it follows that the root of G'_{n+1} is in R_{n+1} , and hence so is the root of G_{n+1} . A similar argument shows that the root of H_{n+1} is in B_{n+1} . \square

LEMMA 3.5.10. *We have $\sigma_0(G_{n+1}) = \sigma_0(H_{n+1}) = k_{n+1}$.*

PROOF. By construction we have that $\sigma_0(\tilde{G}_n) = \sigma_0(\tilde{H}_n) = k_{n+1}$. Since G'_{n+1} and H'_{n+1} are realised as components of the promise closure of M_n , and this was a proper extension, it is a simple check that $\sigma_0(G'_{n+1}) = \sigma_0(H'_{n+1}) = k_{n+1}$. Also note that $F_{n+1} \times \mathbb{N}$ has no mii-paths of length bigger than two, since the vertices of degree two in $F_{n+1} \times \mathbb{N}$ are precisely those of the form $(\ell, 0)$ with ℓ a leaf of F_{n+1} .

Since $G'_{n+1} \oplus_{\chi_{G_{n+1}}} (F_{n+1} \times \mathbb{N})$ is formed by gluing a set of degree-two vertices of $F_{n+1} \times \mathbb{N}$ to a set of degree-three vertices in G'_{n+1} , it follows that $\sigma_0(G_{n+1}) = k_{n+1}$ as claimed. A similar argument shows that $\sigma_0(H_{n+1}) = k_{n+1}$. \square

LEMMA 3.5.11. *The graphs G_{n+1} and H_{n+1} are spectrally distinguishable.*

PROOF. Since in \tilde{G}_n we have that all long mii-paths except for those of length k_{n+1} are contained inside G_n or \hat{H}_n , it follows from our induction assumption (†5) that $\sigma_1(\tilde{G}_n) = k_n$. However, in \tilde{H}_n , we attached $\hat{G}_n(\hat{v})$ to generate an mii-path of length $\tilde{k}_n + 1$ in \tilde{H}_n (see Fig. 3.7), implying that

$$\sigma_1(\tilde{H}_n) = \tilde{k}_n + 1 > k_n = \sigma_1(\tilde{G}_n).$$

As before, since the promise closures G'_{n+1} and H'_{n+1} are proper extensions of \tilde{G}_n and \tilde{H}_n , they are spectrally distinguishable. Lastly, since $F_{n+1} \times \mathbb{N}$ has no leaves and no mii-paths of length bigger than two, the same is true for G_{n+1} and H_{n+1} . \square

LEMMA 3.5.12. *The graphs G_{n+1} and H_{n+1} have exactly one end, and $\Omega(G_{n+1} \cup H_{n+1}) \subseteq \overline{R_{n+1} \cup B_{n+1}}$.*

PROOF. By the induction assumption (†8), we know that $\Omega(G_n \cup H_n) \subseteq \overline{R_n \cup B_n}$.

CLAIM. The set $R_{n+1} \cup B_{n+1}$ is dense for G'_{n+1} .

Consider a finite $S \subseteq V(G'_{n+1})$. We have to show that any infinite component C of $G'_{n+1} - S$ has non-empty intersection with $R_{n+1} \cup B_{n+1}$.

Let us consider the global structure of G'_{n+1} as being roughly that of an infinite regular tree, as in Figure 3.2. Specifically, we imagine a copy of G_n at the top level, at the next level are the copies of G_n and H_n that come from a blue or red leaf in the top level, at the next level are the copies attached to blue or red leaves from the previous level, and so on.

With this in mind, it is evident that either C contains an infinite component from some copy of $H_n - S$ or $G_n - S$, or C contains an infinite ray from this tree structure. In the first case, we have $|C \cap (R_n \cup B_n)| = \infty$ by induction assumption. Since any vertex from $R_n \cup B_n$ has a leaf from $R_{n+1} \cup B_{n+1}$ within distance $k_{n+1} + 1$ (cf. Figure 3.7), it follows that C also meets $R_{n+1} \cup B_{n+1}$ infinitely often. In the second case, the same conclusion follows, since between each level of our tree structure, there is a pair of leaves in $R_{n+1} \cup B_{n+1}$. This establishes the claim.

CLAIM. The set $R_{n+1} \cup B_{n+1}$ is dense for H'_{n+1} .

The proof of the second claim is entirely symmetric to the first claim.

To complete the proof of the lemma, observe that $F_{n+1} \times \mathbb{N}$ is one-ended, and with $R_{n+1} \cup B_{n+1}$, also $\text{dom}(\chi_{G_{n+1}}) \cup \text{dom}(\chi_{H_{n+1}})$ is dense for $G'_{n+1} \cup H'_{n+1}$ by our claims. So by Corollary 3.4.4, the graphs G_{n+1} and H_{n+1} have exactly one end. Moreover, since $R_{n+1} \cup B_{n+1}$ meets both graphs infinitely, it follows immediately that it is dense for $G_{n+1} \cup H_{n+1}$. \square

LEMMA 3.5.13. *The graph G_{n+1} is a proper mii-extension of infinite growth of G_n at $R_n \cup B_n$ to length $k_n + 1$, and $\text{Ball}_{G_{n+1}}(G_n, k_n + 1)$ does not meet $R_{n+1} \cup B_{n+1}$. Similarly, H_{n+1} is a proper mii-extension of infinite growth of H_n at $R_n \cup B_n$ to length $k_n + 1$, and $\text{Ball}_{H_{n+1}}(H_n, k_n + 1)$ does not meet $R_{n+1} \cup B_{n+1}$. Hence, (†9) and (†10) are satisfied at stage $n + 1$.*

PROOF. We will just prove the statement for G_{n+1} , as the corresponding proof for H_{n+1} is analogous.

Since G'_{n+1} is an $\left((\tilde{R}_n \cup \tilde{B}_n) \cap V(\tilde{G}_n)\right)$ -extension of \tilde{G}_n , it follows that G'_{n+1} is an

$$(30) \quad \left(\left((\tilde{R}_n \cup \tilde{B}_n) \cap V(G_n)\right) \cup r(G_n)\right) = \left((R_n \cup B_n) \cap V(G_n)\right)\text{-extension of } G_n.$$

However, from the construction of the closure of a graph it is clear that that G'_{n+1} is also an L' -extension of the supergraph K of G_n formed by gluing a copy of $\tilde{G}_n(\vec{p}_1)$ to every leaf in $R_n \cap V(G_n)$ and a copy of $\tilde{H}_n(\vec{p}_2)$ to every leaf in $B_n \cap V(G_n)$, where L' is defined as the set of inherited promise leaves from the copies of $\tilde{G}_n(\vec{p}_1)$ and $\tilde{H}_n(\vec{p}_2)$.

However, we note that every promise leaf in $\tilde{G}_n(\vec{p}_1)$ and $\tilde{H}_n(\vec{p}_2)$ is at distance at least $\tilde{k}_n + 1$ from the respective root, and so $\text{Ball}_{G'_{n+1}}(G_n, \tilde{k}_n) = \text{Ball}_K(G_n, \tilde{k}_n)$. However, $\text{Ball}_K(G_n, \tilde{k}_n)$ can be seen immediately to be an mii-extension of G_n at $R_n \cup B_n$ to length \tilde{k}_n , and since $\tilde{k}_n \geq k_n + 1$ it follows that $\text{Ball}_{G'_{n+1}}(G_n, k_n + 1)$ is an mii-extension of G_n at $R_n \cup B_n$ to length $k_n + 1$ as claimed.

Finally, we note that $R_{n+1} \cup B_{n+1}$ is the set of promise leaves $\text{cl}(\mathcal{L}_n)$. By the same reasoning as before, $\text{Ball}_{G'_{n+1}}(G_n, k_n + 1)$ contains no promise leaf in $\text{cl}(\mathcal{L}_n)$, and so does not meet $R_{n+1} \cup B_{n+1}$ as claimed. Furthermore, it doesn't meet any neighbours of $R_{n+1} \cup B_{n+1}$.

Recall that G_{n+1} is formed by gluing a set of vertices in $(F_{n+1} \times \mathbb{N})$ to neighbours of vertices in $R_{n+1} \cup B_{n+1}$. However, by the above claim, $\text{Ball}_{G'_{n+1}}(G_n, k_n + 1)$ does not meet any of the neighbours of $R_{n+1} \cup B_{n+1}$ and so $\text{Ball}_{G_{n+1}}(G_n, k_n + 1) = \text{Ball}_{G'_{n+1}}(G_n, k_n + 1)$, and the claim follows.

Finally, to see that G_{n+1} is a leaf extension of G_n of infinite growth, it suffices to observe that $G_{n+1} - G_n$ consists of one component only, which is a superset of the infinite graph $F_n \times \mathbb{N}$. \square

LEMMA 3.5.14. *There is a family of isomorphisms*

$$\mathcal{H}_{n+1} = \{h_{n+1,x}: G_{n+1} - x \rightarrow H_{n+1} - \varphi_{n+1}(x): x \in X_{n+1}\},$$

such that

- $h_{n+1,x} \upharpoonright (G_n - x) = h_{n,x}$ for all $x \in X_n$,
- the image of $R_{n+1} \cap V(G_{n+1})$ under $h_{n+1,x}$ is $R_{n+1} \cap V(H_{n+1})$,
- the image of $B_{n+1} \cap V(G_{n+1})$ under $h_{n+1,x}$ is $B_{n+1} \cap V(H_{n+1})$ for all $x \in X_{n+1}$.

PROOF. Recall that Lemma 3.5.7 shows that there exists such a family of isomorphisms between G'_{n+1} and H'_{n+1} . Furthermore, we have that

$$G_{n+1} := G'_{n+1} \oplus_{\chi_{G_{n+1}}} (F_{n+1} \times \mathbb{N}) \text{ and } H_{n+1} := H'_{n+1} \oplus_{\chi_{H_{n+1}}} (F_{n+1} \times \mathbb{N}).$$

where it is easy to check that $\chi_{G_{n+1}}$ and $\chi_{H_{n+1}}$ satisfy the assumptions of Lemma 3.4.2, since the functions $\psi_{G'_{n+1}}$ and $\psi_{H'_{n+1}}$ do by Lemma 3.5.8.

More precisely, given $x \in X_{n+1}$ and $h'_{n+1,x}$, it follows from Lemma 3.5.8 that

$$\chi_{H_{n+1}} \circ h'_{n+1,x} \circ \chi_{G_{n+1}}$$

extends to an isomorphism $\pi_{n+1,x}$ of F_{n+1} . Hence, by Lemma 3.4.2, $h'_{n+1,x}$ extends to an isomorphism $h_{n+1,x}$ from $G_{n+1} - x$ to $H_{n+1} - y$. That this isomorphism satisfies the three properties claimed follows immediately from Lemma 3.5.7 and the fact that $h_{n+1,x} \upharpoonright (G_n - x) = h'_{n+1,x} \upharpoonright (G_n - x)$. \square

LEMMA 3.5.15. *There exist bijections*

$$\psi_{G_{n+1}}: V(G_{n+1}) \cap (R_{n+1} \cup B_{n+1}) \rightarrow R'_{n+1} \cup B'_{n+1}$$

and

$$\psi_{H_{n+1}}: V(H_{n+1}) \cap (R_{n+1} \cup B_{n+1}) \rightarrow R'_{n+1} \cup B'_{n+1},$$

and a family of isomorphisms

$$\Pi_{n+1} = \{\pi_{n+1,x}: F_{n+1} \rightarrow F_{n+1}: x \in X_{n+1}\},$$

such that

- $\pi_{n+1,x} \upharpoonright R'_{n+1}$ is a permutation of R'_{n+1} for each x ,
- $\pi_{n+1,x} \upharpoonright B'_{n+1}$ is a permutation of B'_{n+1} for each x , and

- for each $x \in X_{n+1}$, the corresponding diagram commutes:

$$\begin{array}{ccc}
 \mathcal{L}(G_{n+1}) & \xrightarrow{h_{n+1,x} \upharpoonright \mathcal{L}(G_{n+1})} & \mathcal{L}(H_{n+1}) \\
 \psi_{G_{n+1}} \downarrow & & \downarrow \psi_{H_{n+1}} \\
 \mathcal{L}(F_{n+1}) & \xrightarrow{\pi_{n+1,x} \upharpoonright \mathcal{L}(F_{n+1})} & \mathcal{L}(F_{n+1})
 \end{array}$$

I.e. for every $\ell \in V(G_{n+1}) \cap (R_{n+1} \cup B_{n+1})$ we have $\pi_{n+1,x}(\psi_{G_{n+1}}(\ell)) = \psi_{H_{n+1}}(h_{n+1,x}(\ell))$.

PROOF. Since $R_{n+1}, B_{n+1} \subseteq G'_{n+1} \cup H'_{n+1}$, and $h_{n+1,x}$ extends $h'_{n+1,x}$ for each $x \in X_{n+1}$, this follows immediately from Lemma 3.5.8 after identifying \hat{F}_{n+1} with F_{n+1} . \square

This completes our recursive construction, and hence the proof of Theorem 3.1.4 is complete.

3.6. A non-reconstructible graph with countably many ends

In this section we will prove Theorem 3.1.5. Since the proof will follow almost exactly the same argument as the proof of Theorem 3.1.4, we will just indicate briefly here the parts which would need to be changed, and how the proof is structured.

The proof follows the same back and forth construction as in Section 3.5.2, however instead of starting with finite graphs G_0 and H_0 we will start with two infinite graphs, each containing one free end. For example we could start with the graphs in Figure 3.9.




FIGURE 3.9. A possible choice for G_0 and H_0 , where the dots indicate a ray.

The induction hypotheses remain the same, with the exception of (†7) and (†8) which are replaced by

(†7') G_n and H_n have exactly one limit end and infinitely many free ends when $n \geq 1$,
and

(†8') $\overline{R_n \cup B_n} \cap \Omega(G_n \cup H_n) = \Omega'(G_n \cup H_n)$.

The arguments of Section 3.5.5 will then go through mutatis mutandis: for the proof of the analogue of Lemma 3.5.12, use Corollary 3.4.5 instead of Corollary 3.4.4.

To show that the construction then yields the desired non-reconstructible pair of graphs with countably many ends, we have to check that (†7') holds for the limit graphs G and H . It is clear that since $\overline{R_n \cup B_n} \cap \Omega(G_n \cup H_n) = \Omega'(G_n \cup H_n)$, every free end in a graph G_n or H_n remains free in the limit. Moreover, a similar argument to that in Section 3.5.3

shows that any pair of rays in G or H which were not in a free end in some G_n or H_n are equivalent in G or H , respectively.

However, since the end space of a locally finite connected graph is a compact metrizable space, and therefore has a countable dense subset, such a graph has at most countably many free ends, since they are isolated in $\Omega(G)$. Hence, both G and H have at most countably many free ends, and one limit end, and so both graphs have countably many ends.

CHAPTER 4

Topological ubiquity of trees

Let \triangleleft be a relation between graphs. We say 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 .

The *Ubiquity Conjecture* of Andreae, a well-known open problem in the theory of infinite graphs, asserts that every locally finite connected graph is ubiquitous with respect to the minor relation.

In this paper, which is the first of a series of papers making progress towards the Ubiquity Conjecture, we show that all trees are ubiquitous with respect to the topological minor relation, irrespective of their cardinality. This answers a question of Andreae from 1979.

4.1. 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 [85, 86] 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. [9, 168, 109], examples of which, due to Andreae [16], are depicted in Figures 4.1 and 4.2 below.




FIGURE 4.1. A graph which is not \subseteq -ubiquitous.




FIGURE 4.2. 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 [15].

THE UBIQUITY CONJECTURE. *Every locally finite connected graph is \preceq -ubiquitous.*

In [15], 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 [148]. In particular, it is still an open question whether or not there exists a countable connected graph which is not \preceq -ubiquitous.

In his most recent paper on ubiquity to date, Andreae [16] exhibited infinite families of locally finite graphs for which the ubiquity conjecture holds. The present paper is the first in a series of papers [32, 33, 34] making further progress towards the ubiquity conjecture, with the aim being to show that all graphs of bounded tree-width are ubiquitous.

As a first step towards this, we in particular need to deal with infinite trees, for which one even gets affirmative results regarding ubiquity under the topological minor relation. Halin showed in [87] that all trees of maximum degree 3 are \leq -ubiquitous. Andreae improved this result to show that all *locally finite* trees are \leq -ubiquitous [10], and asked if his result could be extended to arbitrary trees [10, p. 214]. Our main result of this paper answers this question in the affirmative.

THEOREM 4.1.1. *Every tree is ubiquitous with respect to the topological minor relation.*

The proof will use some results about the well-quasi-ordering of trees under the topological minor relation of Nash-Williams [125] and Laver [111], as well as some notions about the topological structure of infinite graphs [55]. Interestingly, most of the work in proving Theorem 4.1.1 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 [143].

To explain our strategy, let us fix some notation. When H is a subdivision of G we write $G \leq^* H$. Then, $G \leq \Gamma$ means that there is a subgraph $H \subseteq \Gamma$ which is a subdivision of G , that is, $G \leq^* H$. 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 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 4.1.2. *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 4.1.2 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 4.1.1 however, it is true that arbitrary trees are also \preceq -ubiquitous, although the proof involves some extra technical difficulties that we will deal with in a later paper [34]. 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 [32, Theorem 1.7].

This paper is structured as follows: In Section 4.2, 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 [54]. Next, Sections 4.3 to 4.5 introduce the key ingredients for our main ubiquity result. Section 4.3, extending ideas from Andreae's [10], lists three useful corollaries of Nash-Williams' and Laver's result that (labelled) trees are well-quasi-ordered under the topological minor relation, Section 4.4 investigates under which conditions a given family of disjoint rays can be rerouted onto another family of disjoint rays, and Section 4.5 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 Theorem 4.1.2, in Section 4.6. Finally, Section 4.7 contains the induction argument establishing our main result, Theorem 4.1.1.

4.2. Preliminaries

DEFINITION 4.2.1. 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 [54, 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. [54, Section 12.2].

DEFINITION 4.2.2 (Ends of a graph, cf. [54, 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 vertex disjoint $v-R$ -paths in Γ .

DEFINITION 4.2.3. 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 $vPwQx$.

4.3. Well-quasi-orders and κ -embeddability

DEFINITION 4.3.1. 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 4.3.2 (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 4.3.3. *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 4.3.4 (Nash-Williams [125]). *The relation \leq_r is a well-quasi order on the set of rooted trees.*

THEOREM 4.3.5 (Laver [111]). *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 $(T_1, c_1), (T_2, c_2), \dots$ is a sequence of rooted trees with 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 $t \in T_i$.¹*

Together with Lemma 4.3.3 these results give us the following three corollaries:

DEFINITION 4.3.6. 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 4.3.7. *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 4.3.4 the set $\mathcal{T} = \{T_{t'} : t' \in N^+(t)\}$ is well-quasi-ordered by \leq_r and so the claim follows by Lemma 4.3.3 applied to \mathcal{T} , \leq_r , and κ . \square

COROLLARY 4.3.8. *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 .

¹In 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.

PROOF. Consider a labelling $c: T_t \rightarrow [2]$ mapping t to 1, and all remaining vertices of T_t to 2. By Theorem 4.3.5, 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 4.3.3 to \mathcal{T} and \leq_r with $\kappa = \aleph_0$. \square

DEFINITION 4.3.9 (Self-similarity). A ray $R = r_1 r_2 r_3 \dots$ 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 $H(R) \subseteq R$.

COROLLARY 4.3.10. *Let (T, v) be an infinite rooted tree and let $R = r_1 r_2 r_3 \dots$ 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 .*²

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 4.3.5, 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 4.3.3, 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

4.4. 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 4.4.1 (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 4.4.2 (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 $i \in I$.

²A slightly weaker statement, without the additional condition that $H(R) \subseteq R$ appeared in [10, Lemma 1].

LEMMA 4.4.3 (Weak linking lemma). *Let Γ be a graph and $\epsilon \in \Omega(\Gamma)$. 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 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 [54, 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, which is stated in slightly more generality than needed in this paper. Broadly the idea is that if we have a family of disjoint rays $(R_i: i \in [n])$ tending to an end ϵ and a number $a \in \mathbb{N}$, then there is some fixed number $N = N(a, n)$ such that if we have N disjoint graphs H_i , each with a specified ray S_i 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 a of the graphs H_i .

LEMMA 4.4.4 (Strong linking lemma). *Let Γ be a graph and $\epsilon \in \Omega(\Gamma)$. Let X be a finite set of vertices, $a, n \in \mathbb{N}$, 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 $N = N(\mathcal{R}, X, a)$ with the following property: For every collection $(H_j: j \in [N])$ of vertex disjoint subgraphs of Γ , all disjoint from X and each including a specified ray S_j in ϵ , there is a set $A \subseteq [N]$ of size a 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 $\bigcup_{k \in A} H_k$.

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 $N = N_0 + n^3 + a$.

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 family $(H_j : j \in [n^3 + a])$ is 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 contains $x'_i \pmod{n}$,
- \hat{K}_i meets exactly n of the H_j , that is $|\{j \in [n^3 + a] : \hat{K}_i \cap H_j \neq \emptyset\}| = n$, and
- \hat{K}_i avoids X' .

Suppose we have done so for all $i < m$. Let $X_m = X' \cup \bigcup_{i < m} V(\hat{K}_i)$. We will build inductively for $t = 0, \dots, n$ increasing connected subgraphs \hat{K}_m^t that meet $R_i \pmod{n}$, 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 $T(R_m \pmod{n}, X_m)$ meets some H_j not met by \hat{K}_m^t then there is some initial vertex $z_t \in T(R_m \pmod{n}, X_m)$ where it does so and we set $\hat{K}_m^{t+1} := \hat{K}_m^t \cup T(R_m \pmod{n}, X_m) z_t$. Otherwise we may assume $T(R_m \pmod{n}, X_m)$ does not meet any such H_j . In this case, let $j \in [n^3 + a]$ be such that $\hat{K}_m^t \cap H_j = \emptyset$. Since $R_m \pmod{n}$ and S_j belong to the same end ϵ , there is some path P between $T(R_m \pmod{n}, X_m)$ and $T(S_j, X_m)$ which avoids X_m . Since this path meets some H_k with $k \in [n^3 + a]$ 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 $T(R_m \pmod{n}, X_m)$ such that \hat{K}_m^{t+1} is connected and contains $x'_m \pmod{n}$. 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

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

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 , and hence avoids some R_q . Furthermore, there are n distinct \hat{K}_i such that $i = q \pmod{n}$, all of which are disjoint. Hence there is some \hat{K}_r with $r = q \pmod{n}$ 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 [54, Proposition 8.4.1], there is a family of disjoint paths $\mathcal{P} = (P_i: i \in [n])$ in Γ' from x'_i to $y_{\sigma(i)}$. Furthermore, since $|J| \leq n^3$ there is some subset $A \subseteq [n^3 + a]$ of size a such that H_k is disjoint from K for each $k \in A$.

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 $(S_j)_{j \in [n^3 + a]}$ 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 $\bigcup_{k \in A} H_k$. □

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 4.4.5. *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. □

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

LEMMA 4.4.6. *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 . □

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

For showing 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 need to 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 in this series [32, 33, 34].

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 .

DEFINITION 4.5.1 (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 of \mathcal{F} $H \in \bigcup \mathcal{F}$.
- 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*.³
- 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

³A similar notion of *thick* and *thin families* was also introduced by Andreae in [10] (in German) and in [16]. 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 series of papers, is new.

$\varphi_{F'}: F' \rightarrow \Psi(F')$ 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 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 4.5.2 (cf. [10, Lemma 3] or [16, Lemma 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} .

PROOF. \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 $|\hat{F}| \leq k$ for all $F \in \mathcal{F}$. Let $n \in \mathbb{N}$. As \mathcal{F} is thick, there is a layer $F \in \mathcal{F}$ satisfying $|F| \geq n + k$. Thus $|F \setminus \hat{F}| \geq n + k - k = n$. \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. E.g. we find a subtribe in which either all members contain an ϵ -ray, or none of them contain such a ray.

LEMMA 4.5.3 (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

In this series of papers we will be interested in graph relations such as \subseteq , \leq and \preceq . 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 4.5.4. *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 , so by Lemma 4.5.2 a thick flat subtribe \mathcal{F}'' 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 4.5.3, 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 $\mathcal{F}_{n+1} \subseteq \mathcal{F}_n$ 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, [55], 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 4.5.5. *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 the sets $F_X = \{H \in F : H \not\subseteq C(X, \epsilon)\} \subseteq F$ for $F \in \mathcal{F}$ 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 G -tribe $\mathcal{F}' = \{F \setminus F_X : F \in \mathcal{F}\}$ is a thick flat subtribe of \mathcal{F} by Lemma 4.5.2, 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

4.6. Countable subtrees

In this section we prove Theorem 4.1.2. 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 4.5.4, 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 4.6.1, the proof of Theorem 4.1.2 splits into two cases according to whether the number of ϵ -dominating vertices in Γ is finite (Section 4.6.2) or infinite (Section 4.6.3).

4.6.1. 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 4.6.1. 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 4.3.8. We will also consider every component of S^* as a rooted tree given by the induced tree order from T .

DEFINITION 4.6.2. 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 $Ext(S^*) \subseteq E(S^*)$ for the set of extension edges.

Consider the forest $S^* - Ext(S^*)$ obtained from S^* by removing all extension edges. Since every ray in S^* must contain an extension edge by Corollary 4.3.10, each component of $S^* - Ext(S^*)$ is a locally finite rayless tree and so is finite (this argument is inspired by [10, Lemma 2]). We enumerate the components of $S^* - Ext(S^*)$ as S_0^*, S_1^*, \dots in such a way that for every $n \geq 0$, 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 $\partial(S_n) \subseteq Ext(S^*)$. 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 4.5.3 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

$$\partial_\epsilon(S_n) := \{e \in \partial(S_n) : R_e \text{ converges to } \epsilon \text{ according to } \mathcal{F}\},$$

$$\partial_{-\epsilon}(S_n) := \{e \in \partial(S_n) : R_e \text{ is cut from } \epsilon \text{ according to } \mathcal{F}\},$$

$$V_\epsilon(S_n) := \{t \in V(S_n) : t \text{ dominates } \epsilon \text{ according to } \mathcal{F}\}, \text{ and}$$

$$V_{-\epsilon}(S_n) := \{t \in V(S_n) : t \text{ is cut from } \epsilon \text{ according to } \mathcal{F}\}.$$

- 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 4.6.3 (*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 4.6.2, 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} . If we define recursively for each $k \in \mathbb{N}$ $\hat{T}_k = \hat{T}_{k-1}(\hat{T}_1)$ 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 $k \in \mathbb{N}$

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

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 $H^{(k)}(T \setminus T_{v_0}) = H(T \setminus T_{v_0})$ and $H^{(k)}(T_{v_0}) = H(\hat{T}_k)$. 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 $K_H = \{k_1^H, k_2^H, \dots\} \subseteq \mathbb{N}$ such that $\{\{H^{(k)}\} : k \in K_H\}$ agrees about $\partial(S_{n+1})$. Consider the thick subtribe $\mathcal{F}'_n = \{F'_i : F \in \mathcal{F}_n, i \in \mathbb{N}\}$ of \mathcal{F}_n with

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

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

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

By definition of extension edges, cf. Definition 4.6.2, 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

$$(\dagger) \quad \tilde{H}(R_g) \text{ doesn't tend to } \epsilon.$$

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 4.5.3 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 (\dagger) , 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 4.3.8, 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 $\tilde{\mathcal{F}}_n$, 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

4.6.2. Only finitely many vertices dominate ϵ . We first note as in Lemma 4.5.4, 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 4.6.4 (Bounder, extender). Suppose that some thick T -tribe \mathcal{F} which is concentrated at ϵ agrees about S_n for some given $n \in \mathbb{N}$, and $Q_1^n, Q_2^n, \dots, Q_n^n$ are disjoint subdivisions of $S_n^{-\epsilon}$ (note, $S_n^{-\epsilon}$ depends on \mathcal{F}).

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

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

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

To prove Theorem 4.1.2, 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 $S_{r-1}^{-\epsilon} \subseteq S_r^{-\epsilon}$.⁴
- (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 $n = 0$ it 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 4.6.3 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 $X_{n+1} := X_n$.

⁴Note that since ϵ is undominated, every thick T -tribe agrees about the fact that $V_\epsilon(S_i) = \emptyset$ for all $i \in \mathbb{N}$.

Otherwise, we have $f \in \partial_\epsilon(S_{k-1})$. By Lemma 4.5.5(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 4.4.3 there is some linkage $\mathcal{P} \subseteq C(X_n, \epsilon)$ from

$$\mathcal{E}_n \text{ to } (H_j(R_f): j \in [N]),$$

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 4.6.3, 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

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

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

By Lemma 4.6.3(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 4.6.3(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 4.4.4 to find a linkage $\mathcal{P} \subseteq C(X_n, \epsilon)$ from

$$\mathcal{E}_n \text{ to } (H_i(R_{e_0}): i \in [N]),$$

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 4.6.3(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} [54, Exercise 19, Chapter 8], in which case proving ubiquity becomes trivial:

LEMMA 4.6.5. *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 $G \subseteq K_{\aleph_0}$. Hence, we have $\aleph_0 \cdot G \subseteq \aleph_0 \cdot K_{\aleph_0} \subseteq K_{\aleph_0}$. \square

4.6.3. 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 4.6.3 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 4.4.5 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 4.4.5 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 can 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, let us apply Lemma 4.4.6 to $Y := \{y_{e,i}^{n+1} : e \in \partial_\epsilon(S_{k+1}), i \leq k\}$ 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 4.1.2. □

4.7. The induction argument

We consider T as a rooted tree with root r . In Section 4.6 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 4.1.1.

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 $T_{t''}$ (Definition 4.3.6). By Corollary 4.3.7, the number of children of t which are not κ -embeddable is less than κ .

DEFINITION 4.7.1 (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 4.7.2. 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 inductively $S'(n+1)$ to consist of $S'(n)$ together with all non- κ -embeddable children contained in S for all vertices of $S'(n)$. It is clear that $\bigcup_{n \in \mathbb{N}} S'(n)$ is the κ -closure of S' . If κ_n denotes the size of $S'(n)$, then $\kappa_n < \kappa$ by induction with Corollary 4.3.7. 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 4.7.3. 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 4.7.2.

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. [143], but can be read without knowledge of the paper.

DEFINITION 4.7.4 (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 4.7.5. 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 4.7.2 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 4.7.6. *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 4.7.3 can be further extended to an S -horde.

So let us assume that μ is singular, and write $\gamma = \text{cf}(\mu)$. By Lemma 4.7.5, 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 4.1.1.

THEOREM 4.7.7. *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 4.1.2, 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 4.7.2 and so there is an S' -horde in Γ by Theorem 4.1.2. This can be extended to a T -horde in Γ by Lemma 4.7.6 with $\mu = |T|$. \square

Ubiquity of graphs with non-linear end structure

A graph G is said to be \preceq -ubiquitous, where \preceq is the minor relation between graphs, if whenever Γ is a graph with $nG \preceq \Gamma$ for all $n \in \mathbb{N}$, then one also has $\aleph_0 G \preceq \Gamma$, where αG is the disjoint union of α many copies of G . A well-known conjecture of Andreae is that every locally finite connected graph is \preceq -ubiquitous.

In this paper we give a sufficient condition on the structure of the ends of a graph G which implies that G is \preceq -ubiquitous. In particular this implies that the full grid is \preceq -ubiquitous.

5.1. Introduction

This paper is the second in a series of papers making progress towards a conjecture of Andreae on the *ubiquity* of graphs. 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 [85] 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 [16] for some particularly simple examples). In [15] 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 very large cardinality are known [148]. In particular, the question of whether there exists a countable graph which is not \preceq -ubiquitous remains open. Most importantly, however, Andreae [15] 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 [16] 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 5.1.1 (Andreae, [16, 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.*

Note that every end in such a graph G must have degree¹ one.

Andreae's proof employs deep results about well-quasi-orderings of labelled (infinite) trees [110]. Interestingly, the way these tools are used does not require the extra condition in Theorem 5.1.1 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 in our series, [33], to show that this is possible, i.e. that all connected, locally finite graphs of finite tree-width are \preceq -ubiquitous.




FIGURE 5.1. 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: Let $[n] = \{1, 2, \dots, n\}$. In the proof we often have two families of disjoint rays $\mathcal{R} = (R_i : i \in [n])$ and $\mathcal{S} = (S_j : j \in [m])$ in Γ , which we may assume all converge¹ to a common end of Γ , and we wish to find a *linkage* between \mathcal{R} and \mathcal{S} , that is, an injective function $\sigma : [n] \rightarrow [m]$ 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

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

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 5.1). 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 $n \leq m$, there is enough connectivity between the rays in Γ so that such a linkage always exists.

¹A 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 5.2.

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 quite possible that for certain choices of σ no suitable linkage exists. Consider for example the case where Γ is the *half grid* (briefly denoted 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.

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 $G - V(R_i)$.

Thus 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. Since ends of degree at most two are automatically linear, every end of a graph with nowhere-linear end structure must have degree at least three.

Our main theorem in this paper is the following.

THEOREM 5.1.2. *Every locally finite connected graph with nowhere-linear end structure is \preceq -ubiquitous.*

Roughly, if we assume that every end of G has nonlinear structure, then the fact that $nG \preceq \Gamma$ for all $n \in \mathbb{N}$ allows us to deduce that Γ must also have some end with a sufficiently complicated structure that we can always find suitable linkages for all σ as above. In fact, this property is so strong that we do not need to follow Andreae's strategy for such graphs. We can use the linkages to directly build a K_{\aleph_0} -minor of Γ , and it follows that $\aleph_0 G \preceq \Gamma$.

In later papers in the series, we shall need to make more careful use of the ideas developed here. We shall analyse the possible kinds of linkages which can arise between two families of rays converging to a given end. If some end of Γ admits many different kinds of linkages, then we can again find a K_{\aleph_0} -minor. If not, then we can use the results of the present paper to show that certain ends of G are linear. This extra structure allows us to carry out an argument like that of Andreae, but using only the limited collection of these maps σ which we know to be present. This technique will be key to extending Theorem 5.1.1 in [33].

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, let G be the full grid, a graph not previously known to be ubiquitous. The *full grid* (briefly denoted by $\mathbb{Z}\square\mathbb{Z}$) is analogously defined as the half

grid but with $\mathbb{Z} \times \mathbb{Z}$ as vertex set. The grid 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 5.1.2 has the following corollary:

COROLLARY 5.1.3. *The full grid is \preceq -ubiquitous.*

Using an argument similar in spirit to that of Halin [87], we also establish the following theorem in this paper:

THEOREM 5.1.4. *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 5.1.4 implies that all countable trees are \preceq -ubiquitous, see Corollary 5.7.4. We remark that while all trees are ubiquitous with respect to the topological minor relation, [31], the problem whether all uncountable trees are \preceq -ubiquitous has remained open, and we hope to resolve this in a paper in preparation [34].

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 a subgraph of the half grid respectively. Hence, Theorems 5.1.2 and 5.1.4 together have the following corollary.

THEOREM 5.1.5. *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 5.1.4. In our last theorem, we show that DR_k is ubiquitous for all $k \in \mathbb{N}$.

THEOREM 5.1.6. *The k -fold dominated ray DR_k is \preceq -ubiquitous for every $k \in \mathbb{N}$.*

The paper is structured as follows: In Section 5.2 we introduce some basic terminology for talking about minors. In Section 5.3 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 5.4 and 5.5 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 Section 5.6 we re-introduce some concepts from [31] and show that we may assume that the G -minors in Γ are *concentrated* towards some end ϵ of Γ . In Section 5.7 we use the results of the previous section to prove Theorem 5.1.4 and finally in Section 5.8 we prove Theorem 5.1.2 and its corollaries.

5.2. Preliminaries

In our graph theoretic notation we generally follow the textbook of Diestel [54]. 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 $(b, d) \in E(H)$ or $(a, c) \in E(G)$ and $b = d$.

DEFINITION 5.2.1. 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 5.2.2 (Ends of a graph, cf. [54, 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 of all sizes of sets containing vertex disjoint ϵ -rays. 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 $R \in \omega$ such that there are infinitely many $v-R$ -paths in Γ that are vertex disjoint except from v .

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

PROPOSITION 5.2.3. [54, Proposition 8.2.1] *An infinite connected graph contains either a ray or a vertex of infinite degree.*

PROPOSITION 5.2.4. [54, 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 5.2.5 (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 $(v, w) \in E(G)$ 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 5.2.6 (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 5.2.7 (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 $(v, w) \in E(G)$, the unique edge between the branch sets $H(v)$ and $H(w)$ is also the unique edge between $H'(v)$ and $H'(w)$.

DEFINITION 5.2.8 (Tidiness). An *IG* (H, φ) is called *tidy* if

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

Note that every *IG* H 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 5.2.9 (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 *IG* 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 $(u, v) \in E(M)$.

Note that if H is tidy, then $H(M)$ will be tidy. Given a ray $R \subseteq G$ and a tidy *IG* H in a graph Γ , the restriction $H(R)$ is a one-ended tree, and so every ray in $H(R)$ 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 5.2.10 (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 well shall see, $H^\downarrow(R)$ will be a ray.

LEMMA 5.2.11. *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 $(v_i, w_i) \in E(H)$ 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 $P_i \subseteq H(x_{i+1})$ 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

5.3. The Ray Graph

DEFINITION 5.3.1 (Ray graph). Given a finite family of disjoint rays $\mathcal{R} = (R_i : i \in I)$ in a graph Γ the *ray graph* $RG_\Gamma(\mathcal{R}) = RG_\Gamma(R_i : i \in I)$ 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 $RG(\mathcal{R})$ for $RG_\Gamma(\mathcal{R})$.

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

LEMMA 5.3.2. *Let G be a graph and let $\mathcal{R} = (R_i : i \in I)$ be a finite family of disjoint rays in G . Then $RG_G(\mathcal{R})$ is connected if and only if all rays in \mathcal{R} tend to a common end $\omega \in \Omega(G)$.*

LEMMA 5.3.3. *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 $RG_G(\mathcal{R}) = RG_H(\mathcal{R}')$.*

LEMMA 5.3.4. *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 $RG_H(\mathcal{R})$ is a subgraph of $RG_G(\mathcal{R} \cup \mathcal{S})[I]$. In particular, if all rays in \mathcal{R} tend to a common end in H , then $RG_G(\mathcal{R} \cup \mathcal{S})[I]$ 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 5.3.5. *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, ℓ) is an edge in $RG(R_j: j \in [n])$. However, by Lemma 5.3.2 $RG(R_j: j \in [n])$ 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 $RG(\mathcal{R})$ is the path $12 \dots n$. It follows that, for each i , $(k, \ell) \notin E(RG(\mathcal{R}))$ whenever $1 \leq k < i < \ell \leq n - 1$, and so by definition of $RG(\mathcal{R})$ there is no infinite collection of vertex disjoint paths from R_k to R_ℓ in $G - V(R_i)$. Therefore R_k and R_ℓ belong to different ends of $G - V(R_i)$. \square

DEFINITION 5.3.6 (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 5.3.7 (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 Γ , where the initial vertex of each R_i is denoted x_i . A family $\mathcal{P} = (P_i: i \in I)$ of paths in Γ 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 \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 [31], which asserts the existence of linkages.

LEMMA 5.3.8 (Weak linking lemma). *Let Γ be a graph, $\omega \in \Omega(\Gamma)$ 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 .*

5.4. 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 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 5.4.1 (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 $RG(\mathcal{R})$ 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 5.4.2. *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 $RG(S_j: j \in [m])$, 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. We first note 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 will be 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 $RG(S_j: j \in [m])$ 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 $(t, \sigma(t)) \in E(RG(S_j: j \in [m]))$. 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” [102] or “token reconfiguration” [43]. 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 [102] 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 5.4.9. 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 5.4.3. *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 5.4.4. 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 5.4.3, the pebble-permutation group of (G, X) is a subgroup of the symmetric group S_k .

LEMMA 5.4.5. *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 5.4.6. *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 5.4.3 Y is an achievable game state in the X -pebble-pushing game. \square

LEMMA 5.4.7. *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 therefore know by Lemma 5.4.6 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 5.4.8. Given a graph G , a path $x_1x_2 \dots x_m$ in G is a *bare path* if $d_G(x_i) = 2$ for all $2 \leq i \leq m - 1$.

LEMMA 5.4.9. *Let G be a finite connected graph with vertex set V which is not k -pebble-win and with $|V| \geq k + 2$. Then there is a bare path $P = p_1p_2 \dots p_n$ 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. Since G is not k -pebble-win, by Lemma 5.4.7 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 5.4.3 $X^{(ij)}$ 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 $P = p_1p_2 \dots p_m$ 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 .

However, 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, $c_{Z'}(z'_i) = c_Z(p_1) = r$ and $c_{Z'}(z'_j) = c_Z(p_m) = b$, 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 . We wish to show that G is a cycle. We first make the following claim:

CLAIM 9. 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 9. 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

5.5. Pebbly ends

DEFINITION 5.5.1 (Pebbly). Let Γ be a graph and ω an end of Γ . We say ω is *pebbly* 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 $RG(\mathcal{R})$ is k -pebble-win. If for some k there is no such family \mathcal{R} , we say ω is *not k -pebble-win*.

The following is an immediate corollary of Lemma 5.4.9.

COROLLARY 5.5.2. *Let ω be an end of a graph Γ which is not k -pebble-win 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 $RG(R_i : i \in [m])$ such that $|[m] \setminus V(P)| \leq k$. Furthermore, either each edge in P is a bridge in $RG(R_i : i \in [m])$, or $RG(R_i : i \in [m])$ is a cycle.*

Hence, if an end in Γ is not pebbly, then we have some constraint on the behaviour of rays towards this ends. In a later paper [33] we will investigate more precisely what can be said about the structure of the graph towards this end. For now, the following lemma allows us to easily find any countable graph as a minor of a graph with a pebbly end.

LEMMA 5.5.3. *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 $k \in \mathbb{N}$, $RG(\mathcal{R}_k)$ 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} by choosing some bijection $\sigma : \mathbb{N} \cup \mathbb{N}^{(2)} \rightarrow \mathbb{N}$ such that $\sigma(i, j) > \sigma(i), \sigma(j)$ 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 5.3.8, 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 $RG(\mathcal{R}_r)$ is r -pebble-win,

it follows that $RG(T_i: i \in [m_r])$ is also r -pebble-win. Furthermore, since by Lemma 5.3.2 $RG(T_i: i \in [m_r])$ is connected, there is some neighbour $w \in [m_r]$ of u in $RG(T_i: i \in [m_r])$.

Let us first assume that $w \notin [r]$. Since $RG(T_i: i \in [m_r])$ 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 5.4.2 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 $RG(T_i: i \in [m_r])$ 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 $RG(T_i: i \in [m_r])$ 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 5.4.2 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 $\hat{y} \in T_w x'_v$ 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 5.5.4. *Let Γ be a graph with a pebbly end ω and let G be a countable graph. Then $\aleph_0 G \preceq \Gamma$.*

5.6. 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 [31].

DEFINITION 5.6.1 (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 -subtribe¹ 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 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} . 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 [31].

LEMMA 5.6.2 (Removing a thin subtribe, [31, 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 5.6.3 (Pigeon hole principle for thick G -tribes, [31, 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 lemma, it is necessary that G is connected, so that every member of the G -tribe is a connected graph.

¹When G is clear from the context we will often refer to a G -subtribe as simply a subtribe.

LEMMA 5.6.4 ([31, 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 5.6.5 ([31, 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 5.6.6. *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 5.6.5(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 ourself to thick G -tribes which are concentrated at thick ends.

LEMMA 5.6.7. *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 know by Proposition 5.2.4 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 [82, 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 5.6.6. \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 5.3.2 and 5.3.4, 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 5.6.8. *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 we know by Proposition 5.2.4 that Γ contains a topological K_{\aleph_0} minor, in which case $\aleph_0 G \preceq \Gamma$, 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

5.7. 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 5.2). The following theorem of Halin is one of the cornerstones of infinite graph theory.




FIGURE 5.2. The hexagonal half grid \mathbb{H} .

THEOREM 5.7.1 (Halin, see [54, Theorem 8.2.6]). *Whenever a graph Γ contains a thick end, then $\mathbb{H} \leq \Gamma$.* \square

In [87], 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.

LEMMA 5.7.2 ([87, (4) in Section 3]). *$\aleph_0\mathbb{H}$ is a topological minor of \mathbb{H} .*

THEOREM 5.1.4. *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 5.6.4 we may assume there is an end ϵ of Γ and a thick G -tribe \mathcal{F} which is concentrated at ϵ . By Lemma 5.6.7 we may assume that ϵ is thick. Hence $\mathbb{H} \leq \Gamma$ by Theorem 5.7.1, and with Lemma 5.7.2 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 5.7.3. *\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

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

from which the result follows. \square

COROLLARY 5.7.4. *All countable trees are ubiquitous with respect to the minor relation.*

PROOF. This is an immediate consequence of Lemma 5.7.3 and Theorem 5.1.4. \square

5.8. Proof of main results

LEMMA 5.8.1. *Let ϵ be a non-pebbly end of Γ and let \mathcal{F} be a 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}' 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 IGS, 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 rays which tend to ϵ in Γ and by Lemma 5.3.3 and Lemma 5.3.4 $RG_\Gamma(\mathcal{S})$ contains disjoint subgraphs K_1, K_2, \dots, K_{k+2} such that $K_i \cong RG_G(R_j^i : j \in [m_i])$. However, by Corollary 5.5.2, there is a set X of vertices of size at most k such that $RG_\Gamma(\mathcal{S}) - X$ is a bare path P . However, then some $K_i \subseteq P$ is a path, a contradiction.

Since \mathcal{F} is the union of \mathcal{F}' and \mathcal{F}'' where $\mathcal{F}'' = \{F'' : F \in \mathcal{F}\}$ with

$$F'' = \{H : H \in F \text{ and } \omega_H \text{ is linear}\},$$

it follows that \mathcal{F}'' is thick. \square

THEOREM 5.1.2. *Every locally finite connected graph with nowhere-linear end structure is \preceq -ubiquitous.*

PROOF. Let Γ be a graph such that $nG \preceq \Gamma$ holds for every $n \in \mathbb{N}$. Hence, Γ contains a thick G -tribe \mathcal{F} . By Lemmas 5.6.4 and 5.6.6 we may assume that \mathcal{F} is strongly concentrated at an end ϵ of Γ and so by Lemma 5.6.8 we may assume that for every $H \in \bigcup \mathcal{F}$ there is an end $\omega_H \in \Omega(G)$ such that $H(\omega_H) = \epsilon$.

Since ω_H is not linear for each $H \in \bigcup \mathcal{F}$, it follows by Lemma 5.8.1 that ϵ is pebbly, and hence by Corollary 5.5.4 $\aleph_0 G \preccurlyeq \Gamma$. \square

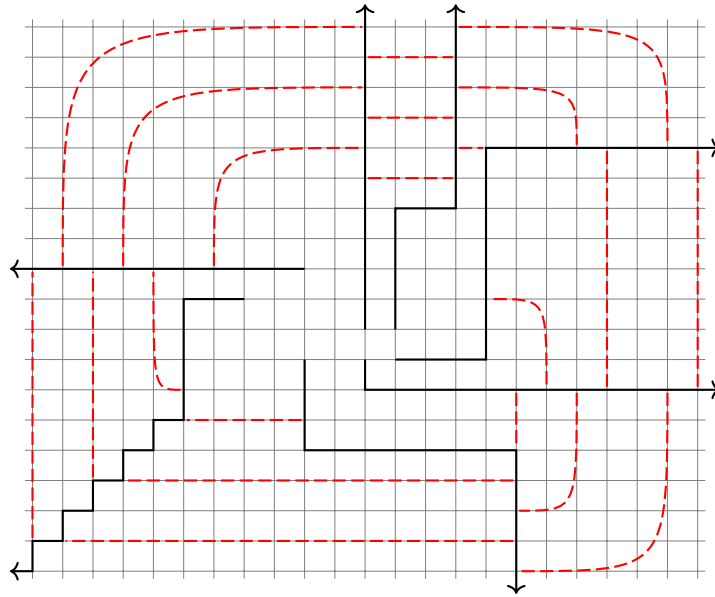


FIGURE 5.3. The ray graphs in the full grid are cycles.

COROLLARY 5.1.3. *The full grid is \preccurlyeq -ubiquitous.*

PROOF. Let G be the full grid. Since $G - R$ has at most one end for any ray $R \in G$, by Lemma 5.3.2 the ray graph $RG(\mathcal{R})$ is 2-connected for any finite family of three or more rays. Hence, by Theorem 5.1.2 G is \preccurlyeq -ubiquitous \square

REMARK 5.8.2. *In fact, every ray graph in the full grid is a cycle (see Figure 5.3).*

THEOREM 5.1.5. *For every locally finite connected graph G , both $G \square \mathbb{Z}$ and $G \square \mathbb{N}$ are \preccurlyeq -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 \preccurlyeq -ubiquitous by Theorem 5.1.4. If G is a double ray then $G \square \mathbb{Z}$ is the full grid and thus \preccurlyeq -ubiquitous by Corollary 5.1.3. Otherwise let G' be a finite connected subgraph of G which is not a path. For any end ω of $G \square \mathbb{Z}$ there is a ray R of \mathbb{Z} such that all rays of the form $\{v\} \square R$ for $v \in V(G')$ go to ω . But then G' is a subgraph of $RG_{G \square \mathbb{Z}}(\{\{v\} \square R\}_{v \in V(G')})$, so this ray-graph is not a path, hence by Lemma 5.3.5 $G \square \mathbb{Z}$ has nowhere-linear end structure and is therefore \preccurlyeq -ubiquitous by Theorem 5.1.2. \square

Finally let us prove Theorem 5.1.6. Recall that for $k \in \mathbb{N}$ 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 5.2.3.

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 5.8.3. [54, 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 5.1.6. *The k -fold dominated ray DR_k is \preceq -ubiquitous for every $k \in \mathbb{N}$.*

PROOF. Note that if $k \leq 2$ then DR_k is a minor of the half grid, and hence ubiquity follows from Theorem 5.1.4.

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. By Lemma 5.6.6 we may assume that there is an end ϵ of Γ such that \mathcal{F} is 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 5.2.4. 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 5.5.4. So suppose for a contradiction that ϵ is not r -pebble-win for some $r \in \mathbb{N}$.

Let R be the ray as stated in the definition of DR_k and let $v_1, v_2, \dots, v_k \in V(DR_k)$ be the vertices adjacent to each vertex of R . For each $H \in \bigcup \mathcal{F}$ and each $i \in [k]$ we have the $H(v_i)$ 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 dominated R , U is infinite, and so by Proposition 5.8.3 $H(v_i)$ 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 $R_{H,i}$. Let $R_H = H^\downarrow(R)$ be the pullback of the ray R in H . 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 $RG_H(\mathcal{R}_H)$ each $R_{H,i}$ is adjacent to R_H . Hence $RG_H(\mathcal{R}_H)$ 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 Lemma 5.3.3 and Lemma 5.3.4 $RG_\Gamma(\mathcal{R})$ contains $r+1$ vertices whose degree is at least $k \geq 3$. However, by Corollary 5.5.2, there is a vertex set X of size at most r such that $RG_\Gamma(\mathcal{R}) - X$ is a bare path P . But then some vertex whose degree is at least 3 is contained in the bare path, a contradiction. \square

Ubiquity of locally finite graphs with extensive tree decompositions

A graph G is said to be \preceq -ubiquitous, where \preceq is the minor relation between graphs, if whenever Γ is a graph with $nG \preceq \Gamma$ for all $n \in \mathbb{N}$, then one also has $\aleph_0 G \preceq \Gamma$. A well-known conjecture of Andreae is that every locally finite graph is \preceq -ubiquitous.

In this paper we show that locally finite graphs admitting a certain type of tree-decomposition, which we call *extensive tree decomposition*, are \preceq -ubiquitous. In particular this includes all locally finite graphs of finite tree-width and locally finite graphs with finitely many ends, all of which are thin.

6.1. 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 $\aleph_\alpha G \triangleleft \Gamma$, where αG is the disjoint union of α many copies of G . A classic result of Halin [85, Satz 1] says that the ray 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 is also \subseteq -ubiquitous [86].

However, not all graphs are \subseteq -ubiquitous, and in fact even trees can fail to be \subseteq -ubiquitous (see for example [168]). 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 [109, 9]), where \leq is the topological minor relation, it was shown by Andreae that all rayless countable graphs [11] and all locally finite trees [10] are \leq -ubiquitous. The latter result was recently extended to the class of all trees by the authors [31].

In [15] Andreae initiated the study of ubiquity of graphs with respect to the minor relation, \preceq . He constructed a graph which was 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 very large cardinality are known [148]. In particular, the question 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 [16] 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 6.1.1 (Andreae, [16, 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 6.1.2 (Andreae, [16, Corollary 2]). *Let G be a locally finite, connected graph of finite tree-width such that every block of G is finite. Then G is \preceq -ubiquitous.*

Note, in particular, that if G is such a graph, then the degree of every end in G must be one.¹ In this paper we will extend Andreae's approach to prove that an even larger class of locally finite graphs is \preceq -ubiquitous, removing the assumption of finite blocks. Again, the exact definition of this class will be technical, but in particular it will imply the following results, extending Theorems 6.1.1 and 6.1.2:

THEOREM 6.1.3. *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 6.1.4. *Every locally finite, connected graph of finite tree-width is \preceq -ubiquitous.*

The proof uses in an essential way some known results about the well-quasi-ordering of graphs under the minor relation, including Thomas' result [149] that graphs of bounded tree width are well-quasi-ordered under the minor relation. Our methods, building on Andreae's, give a blueprint by which stronger results about the well-quasi-ordering of graphs can be used to prove the ubiquity of larger classes of graphs. A more precise discussion of this connection will be given in Section 6.10.

In Section 6.2 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.

6.2. Proof sketch

To give a flavour of the main ideas involved in the proof, let's begin by considering the case of a locally finite connected graph G with a single end ω , where ω has finite degree d (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, \mathcal{V}) 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 $\bigcup_{l \geq n} V_l$ by G_n for each n , the following holds:

¹A precise definitions of the ends of a graph and their degree can be found in Section 6.3.

- (1) each V_n is finite;
- (2) every vertex of G appears in only finitely many V_n ;
- (3) all the A_i begin in V_0 , and
- (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 $i \leq d$, e is an edge of A_i if and only if the edge representing it in this minor is.

Property 4 seems rather strong, and 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 6.5. The skeptical 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 then 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 .

We now show that we can suppose without loss of generality that all G -minors go to the same end ϵ of Γ . For suppose that there are two G -minors H and H' with $\epsilon_H \neq \epsilon_{H'}$. Since G is locally finite, we may assume that all branch sets of H and H' are finite. Thus there is a finite set X such that each of H and H' only uses vertices from one component of $\Gamma - X$. In any $(|X| + 2n)G$ -minor of Γ , only at most $|X|$ of the G -minors involved can meet X , and each of the remaining $2n$ must be included in some component of $G - X$. Without loss of generality at most n of them are in the component that meets H , and so $\Gamma - H$ has an nG -minor.

Thus there is a G -minor H_0 of Γ such that $\Gamma_1 := \Gamma - H_0$ still has an nG -minor for each n . If there are two G -minors going to different ends of Γ_1 then we may as above find a G -minor H_1 of Γ_1 such that $\Gamma_2 := \Gamma_1 - H_1$ has an nG -minor for any n . Proceeding in this way we either find infinitely many disjoint G -minors H_0, H_1, H_2, \dots , giving an $\aleph_0 G$ -minor, or else after finitely many steps we find a subgraph Γ_k of Γ which has an nG -minor for any n and in which all G -minors go to the same end ϵ .

So from now on we will assume that all G -minors of Γ go to the same end ϵ . From any G -minor H we obtain rays $H(A_i)$ corresponding to our marked rays A_i in G . We will call this collection 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 to 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. 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 .



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 throw away 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 6.2.1 (Linkage of families of rays). Let $\mathcal{R} = (R_i : i \in I)$ and $\mathcal{S} = (S_j : j \in J)$ be families of disjoint rays, where the initial vertex of each R_i is denoted x_i . A family of paths $\mathcal{P} = (P_i : i \in I)$, is a *linkage* from \mathcal{R} to \mathcal{S} if there is an injective function $\sigma : I \rightarrow J$ such that

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

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

²Where we use the notation as in [54], see also Definition 6.3.3.

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 [31]:

LEMMA 6.2.2 (Weak linking lemma [31, Lemma 4.3]). *Let Γ be a graph and $\omega \in \Omega(\Gamma)$. Then for any collections $\mathcal{R} = (R_1, \dots, R_n)$ 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 copies of V_{n+1} sufficiently far along the new rays that we can stick them on to 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 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 [31, Lemma 4.4] or 4.4.4 for a precise statement.

A more serious difficulty is that in order to stick the new 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 precisely 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 6.2.3 (Transition function). Let $\mathcal{R} = (R_i : i \in I)$ and $\mathcal{S} = (S_j : j \in J)$ be families of disjoint rays, where the initial vertex of each R_i is denoted x_i . 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 6.2.4 (Ray graph). Given a finite family of disjoint rays $\mathcal{R} = (R_i : i \in I)$ in a graph Γ the *ray-graph*, $RG_\Gamma(\mathcal{R}) = RG_\Gamma(R_i : i \in I)$ 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 $RG(\mathcal{R})$ for $RG_\Gamma(\mathcal{R})$.

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.

In Sections 6.6 and 6.7 we therefore carefully analyse the possibilities for how the ray graphs and transition functions associated to a given thick³ end may look. We find that there are just 3 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 [32]. 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 paths* have degree 2 in the ray graph). 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.

The structure of the ray graphs and transition functions can be used to get around the problem discussed above, by slightly strengthening the properties required of the rays in the recursive construction. More precisely, we want that the ray graph of a slightly larger family \mathcal{R} of disjoint rays, consisting of the $R_{m,i}^n$ and some extra ‘junk’ rays, should have all the $R_{m,i}^n$ on the central path, arranged in such a way that for each n and m the $R_{m,i}^n$ are consecutive in order from $R_{m,1}^n$ to $R_{m,k}^n$.

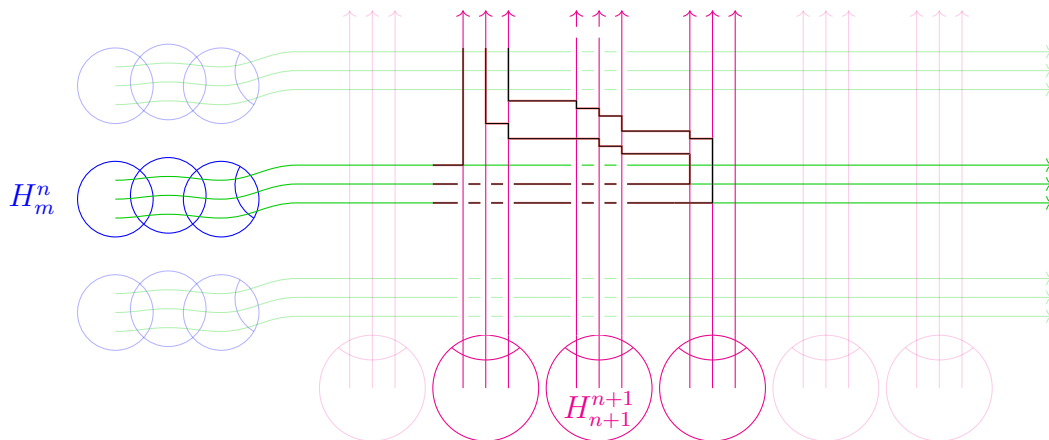
Of course, in order that this is possible we must first ensure that the A_i are arranged in order so that for every n we can find n disjoint G -minors H such that there is some ray graph in which, for each H , the rays $H(A_i)$ appear in order along the central path. Since there are only finitely many possible orders, there must be an order with this property.

Then our extra order assumptions ensure that, by transitioning between rays using edges of the ray graph, we can modify the linkage 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 A_i ’s, disjoint from all the other H_m^{n+1} and their bundles. Recall that the strong 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 modify it to one satisfying $(*)$ by diverting the rays along some of the paths between the new rays.

³An end is *thick* if there are infinitely many disjoint rays to it.

But then some of the rays through which we divert may be forced to intersect the rays emanating from H_{n+1}^{n+1} , if these rays from H_{n+1}^{n+1} lie between rays from the same bundle of some H_m^n .



However, we can get around this by using the paths between the rays in \mathcal{R} to jump between them *before* the linkage, so as to rearrange which bundles make use of (the tails of) which rays. More precisely, we first take a large but finite set of paths between the rays which is rich enough to allow us to rearrange which bundles end up where as much as possible. We collect these together in a *transition box*. Only then 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 bundles should be arranged in order that the rays emanating from H_{n+1}^{n+1} do not appear between rays from the same bundle, we can go back and perform a suitable rearrangement within the transition box, see Figure 6.1.

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 6.4.

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 . But for different values of δ , the ends $H(\delta)$ may well be different.

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 towards which the G -minors are *concentrated*, in the sense that for any finite X there are arbitrarily large families of G -minors in the same component of $G - X$ as ϵ . See Section 5.6 for details. In that section we introduce the term *tribe* for a collection of arbitrarily large families of disjoint G -minors.

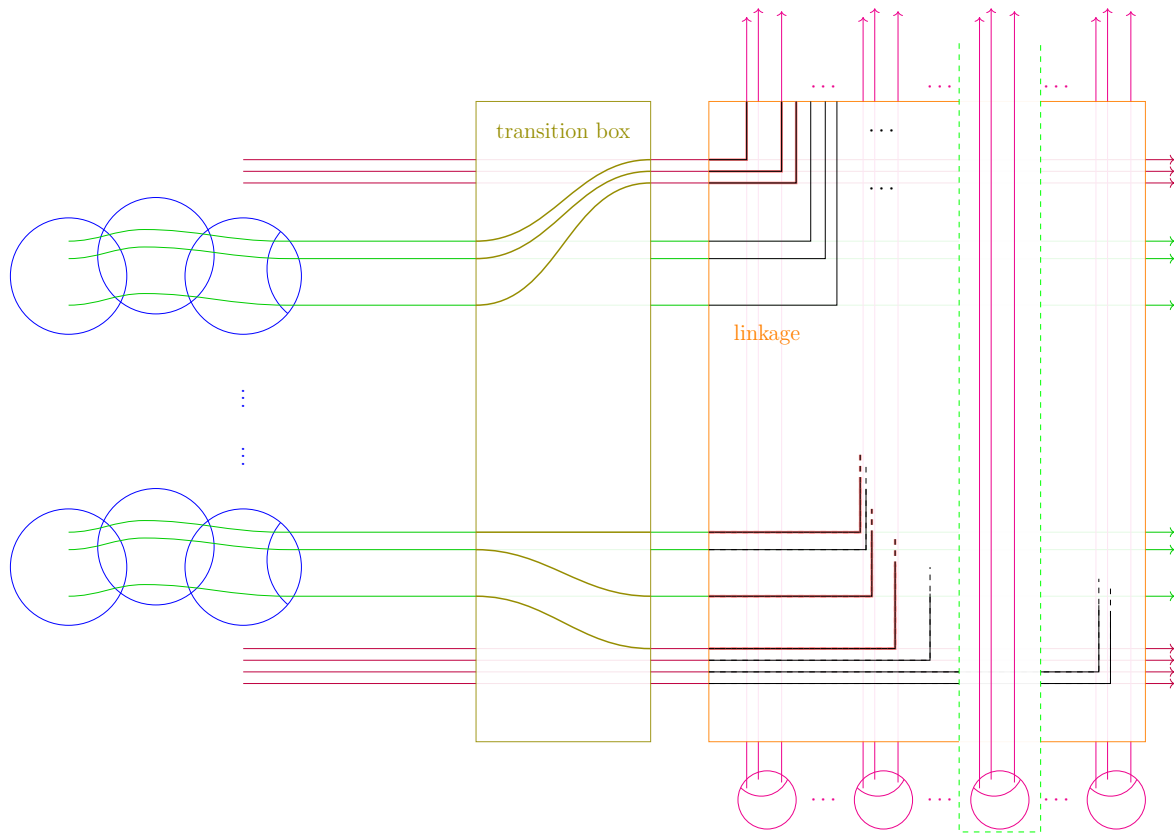


FIGURE 6.1. The transitioning strategy between the old and new bundles.

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 to ϵ corresponding to some canonical rays going to certain ends δ of G .

In order for this to work, we need some consistency about which $H(\delta)$ 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 [31], where they appear in slightly simpler circumstances.

The paper is structured as follows. In Section 5.2 we give precise definitions of some of the basic concepts we will be using, and prove some of their fundamental properties. In Section 6.4 we introduce extensive tree decompositions and in Section 6.5 we illustrate that many locally finite graphs can be given such decompositions. Sections 6.6 and 6.7 are devoted to the possible collections of ray graphs and transition functions between them which can occur in a thick end. In Section 5.6 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 4.6. We conclude with a discussion of the future outlook in Section 6.10.

6.3. Preliminaries

In this paper we follow the convention that 0 is not an element of the set \mathbb{N} of natural numbers.

For a graph $G = (V, E)$ and $W \subseteq V$ we write $G[W]$ for the induced subgraph. 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 $\deg_G(v_i) = 2$ for all inner vertices v_i for $0 < i < n$.

6.3.1. Rays and ends.

DEFINITION 6.3.1 (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 set \mathcal{R} of rays let $\text{init}(\mathcal{R})$ denote the set of initial vertices of the rays in \mathcal{R} .

DEFINITION 6.3.2 (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* , $T(R, X)$, is the unique infinite component of R in $G - X$.

DEFINITION 6.3.3 (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 $\overset{\circ}{v}P\overset{\circ}{w}$ for 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 paths without the vertex v .

Given two paths or rays P and Q which 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 $vPwQx$.

For a ray $R = r_0 r_1 \dots$ let R^- denote the tail $r_1 R$ of R starting at r_1 . Given a set \mathcal{R} of rays let \mathcal{R}^- denote the set $\{R^- : R \in \mathcal{R}\}$

DEFINITION 6.3.4 (Ends of a graph, cf. [54, 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 , denoted by $\deg(\epsilon) \in \mathbb{N} \cup \{\infty\}$, as the largest cardinality of a collection of vertex disjoint ϵ -rays. An end with finite/infinite degree is called *thin/thick*.

6.3.2. Inflated copies of graphs.

DEFINITION 6.3.5 (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 $(v, w) \in E(G)$ 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 6.3.6 (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 .

DEFINITION 6.3.7 (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 $v \in V(G) \cap V(G')$.

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 .

Note that since $H \subseteq H'$, for every edge $(v, w) \in E(G)$, the unique edge between the branch sets $H'(v)$ and $H'(w)$ is also the unique edge between $H(v)$ and $H(w)$.

DEFINITION 6.3.8 (Tidiness). An *IG* (H, φ) is called *tidy* if

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

Note that every IG H 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 6.3.9 (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 IG 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 $(u, v) \in E(M)$.

Note that if H is tidy, then $H(M)$ will be tidy. Given a ray $R \subseteq G$ and a tidy IG H in a graph Γ , the restriction $H(R)$ is a one-ended tree, and so every ray in $H(R)$ 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 6.3.10 (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 well shall see, $H^\downarrow(R)$ will be a ray.

LEMMA 6.3.11. *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 $(v_i, w_i) \in E(H)$ 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 $P_i \subseteq H(x_{i+1})$ 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

DEFINITION 6.3.12. 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 $(H^\downarrow(R) : R \in \mathcal{R})$.

DEFINITION 6.3.13. 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 $R \in \omega$.

It is an easy check that if two rays R and S in G are equivalent, then also $H^\downarrow(R)$ and $H^\downarrow(S)$ are rays (Lemma 6.3.11) which are equivalent in H , and hence also equivalent in Γ .

6.3.3. Transitional linkages and the strong linking lemma.

DEFINITION 6.3.14. We say a linkage is *transitional* if the function which it induces between the corresponding ray graphs is a transition function.

LEMMA 6.3.15. *Let Γ be a graph and $\epsilon \in \Omega(\Gamma)$. Then for any collections $\mathcal{R} = (R_1, \dots, R_n)$ and $\mathcal{S} = (S_1, \dots, S_n)$ of ϵ -rays in Γ there is a finite set X such that every linkage after X is transitional.*

PROOF. By definition, for every function $\sigma: [n] \rightarrow [n]$ 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 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 6.2.2 we will also need the following stronger linking lemma, which is a slight modification of [31, Lemma 4.4]:

LEMMA 6.3.16 (Strong linking lemma). *Let Γ be a graph and $\omega \in \Omega(\Gamma)$. 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 ω . Let $x_i = \text{init}(R_i)$ and $x'_i = \text{init}(T(R_i, X))$. 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 subgraphs of Γ , all disjoint from X and each including a specified ray S_j in ω , there is a $j \in [N]$ and a transitional 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_j .

PROOF. Let $Y \subseteq V(\Gamma)$ be a finite set as in Lemma 6.3.15. We apply the strong linking lemma established in [31, Lemma 4.4] to the set $X \cup Y$ to obtain this version of the strong linking lemma. \square

LEMMA AND DEFINITION 6.3.17. *Let Γ be a graph, $\epsilon \in \Omega(\Gamma)$, $X \subseteq V(\Gamma)$ be finite, and let $\mathcal{R} = (R_i: i \in I_1)$, $\mathcal{S} = (S_i: i \in I_2)$ be two finite families of disjoint ϵ -rays with $|I_1| \leq |I_2|$. Then there is a finite subgraph $Y \subseteq C(X, \epsilon)$ such that for any transition function σ between \mathcal{R} and \mathcal{S} there is a linkage \mathcal{P}_σ from \mathcal{R} to \mathcal{S} inducing σ with $\bigcup \mathcal{P}_\sigma \subseteq \Gamma[Y]$.*

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

PROOF. Let $\sigma: I_1 \rightarrow I_2$ be a transition function between \mathcal{R} and \mathcal{S} . By definition there is a linkage \mathcal{P}_σ from \mathcal{R} to \mathcal{S} after X which induces σ . Note that, since \mathcal{P}_σ is after X , it follows that $\bigcup \mathcal{P}_\sigma \subseteq C(X, \epsilon)$.

Let Φ be the set of all transition functions between \mathcal{R} and \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 6.3.18. 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 \mathcal{P}_2 a transitional linkage from \mathcal{R}_2 to \mathcal{R}_3 after $\bigcup \mathcal{P}_2$.

- (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 6.3.19. *Let $(R_i: i \in I)$ be a disjoint finite family of ϵ -rays, then the ray graph $RG(R_i: i \in I)$ is connected. Also, if R'_i is a tail of R_i for each $i \in I$, then $RG(R_i: i \in I) = RG(R'_i: i \in I)$. \square*

LEMMA 6.3.20 ([32, Lemma 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 $RG_H(\mathcal{R})$ is a subgraph of $RG_G(\mathcal{R} \cup \mathcal{S})[I]$. \square*

6.4. 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 [10, Lemma 2].

6.4.1. Separations and tree-decompositions of graphs.

DEFINITION 6.4.1. Let T be a tree with a root $v \in V(T)$. Given nodes $x, y \in V(T)$ let us denote by xTy the unique path in T between x and y , by T_x denote the component of $T - E(vTx)$ containing x , and by $\overline{T_x}$ the tree $T - T_x$.

Given an edge $e = tt' \in E(T)$, we say that t is the *lower vertex* of e , 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 .

DEFINITION 6.4.2. 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 $G[B]$.

A reader unfamiliar with tree-decompositions may also consult [54, §12.3].

DEFINITION 6.4.3 (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 collection of subsets of vertices $\mathcal{V} = (V_t \subseteq V(G): t \in V(T))$ 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 6.4.4 (Tree-width). Suppose (T, \mathcal{V}) is a tree-decomposition of a graph G . The *width* of (T, \mathcal{V}) is the number $\sup \{|V_t| - 1 : t \in V(T)\} \in \mathbb{N} \cup \{\infty\}$. The *tree-width* of a graph G is the least width of any tree-decomposition of G .

DEFINITION 6.4.5 (Separations induced by tree-decompositions). Given a tree-decomposition (T, \mathcal{V}) of a graph G , and an edge $e \in E(T)$, let

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

Then $(A(e), B(e))$ is a separation of G (cf. [54, 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 6.4.2.)

DEFINITION 6.4.6. 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 6.4.7 (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 set $\{R_{e,s} : s \in S(e)\}$ 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 an W is called a *witness for the self-similarity of $B(e)$ of distance at least n* .

DEFINITION 6.4.8 (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} , and
- for each $e \in E(T)$, the bough $B(e)$ is self-similar towards some end ω_e of G .

The following is the main result of this paper.

THEOREM 6.4.9. *Every locally finite connected graph admitting an extensive tree-decomposition is \preceq -ubiquitous.*

6.4.2. Self minors and push-outs. 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 6.4.10. 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 $x, y \in G$, $H_1(x) \cap H_2(y) \neq \emptyset$ 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 $H = \bigcup_{i \in [n]} P_i \cup H_1 \cup H_2$ and

$$H(v) = \varphi^{-1}(v) := \begin{cases} H_1(v_i) \cup V(P_i) \cup H_2(v_i) & \text{if } v = v_i \in A \cap B, \\ H_1(v) & \text{if } v \in A \setminus B, \\ H_2(v) & \text{if } v \in B \setminus A. \end{cases}$$

DEFINITION 6.4.11 (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 $e' \in R_e$ such that $\text{dist}(e^-, e'^-) \geq n$ and a subgraph $W \subseteq B(e')$ which is an $IG[\overline{B(e)}]$ such that $G' = G[A(e)] \oplus_{\mathcal{P}} W$, where $\mathcal{P} = (P_s : s \in S(e))$ is defined as the family of paths where P_s is the initial segment of $R_{e,s}$ up to the first point it meets $W(s)$.

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 $e' \in R_e$ such that $\text{dist}(e^-, e'^-) \geq n$ and a subgraph $W \subseteq H(B(e'))$ which is an $IG[\overline{B(e)}]$ such that

$$H' = H(G[A(e)]) \oplus_{\mathcal{P}} W$$

where $\mathcal{P} = (P_s : s \in S(e))$ is defined as the family of paths where P_s is the initial segment of $H^\downarrow(R_{e,s})$ up to the first point it meets $W(s)$.

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 6.4.12. *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 $G_W := G[A(e)] \oplus_{\mathcal{P}} W$ of G along e to depth n , where $\mathcal{P} = (P_s : s \in S(e))$ is defined as the family of paths where P_s is the initial segment of $R_{e,s}$ up to the first point it meets $W(s)$.*

PROOF. Given an edge $e \in E(T)$, by Definition 6.4.7 for every $n \in \mathbb{N}$ there is a witness W for the self-similarity of $B(e)$ of distance at least n along the ray R_e .

Explicitly there is a family of rays $(R_{e,s} : s \in S(e))$ such that for every $n \in \mathbb{N}$ there is an edge $e' \in E(T_{e^+})$ of distance at least n from e , and a subgraph $W \subseteq G[B(e')]$, such that

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

Since $(A(e), B(e))$ and $(A(e'), B(e'))$, and $W \subseteq B(e')$ it is clear that $W \cap G[A(e)] \subseteq S(e)$, and since

Let us define $\mathcal{P} = (P_s : s \in S(e))$ as in the statement of the lemma. It is clear that each P_s is from $G[A(e)](s)$ to $W(s)$, and is otherwise disjoint from $G[A(e)] \cup W$.

Furthermore, since $(A(e), B(e))$ and $(A(e'), B(e'))$ are nested separations of G , $A(e) \cap V(W) \subseteq S(e) \cap S(e')$. Hence if $W(s) \cap G[A(e)](s') \neq \emptyset$ it follows that $s' \in S(e) \cap S(e')$, and hence $s' \in V(R_{e,s'}) \cap S(e') \subseteq W(s')$, by Definition 6.4.7. In particular, $W(s) \cap G[A(e)](s') \neq \emptyset$ only if $s = s' \in S(e)$.

Hence, by Definitions 6.4.10 and 6.4.11, $G[A(e)] \oplus_{\mathcal{P}} W$ is well-defined and is indeed a push-out of G along e to depth n . \square

The existence of push-out 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 6.9.

6.5. 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. We will deduce the existence of such 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 sequences $x_1, x_2, \dots \in X$ there exists an $i < j$ such that $x_i \preceq x_j$. The following two alternative characterisations will be useful.

REMARK 6.5.1. *A simple Ramsey type argument shows that if \preceq is a well-quasi-order on X , then every sequence $x_1, x_2, \dots \in X$ contains an increasing subsequence $x_{i_1}, x_{i_2}, \dots \in X$. That is, an increasing 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 and $x_1, x_2, \dots \in X$, then 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 [136], proved over a series of 20 papers, shows that finite graphs are well-quasi-ordered under the minor relation. Thomas [149]

showed that for any $k \in \mathbb{N}$ the class of graphs with tree-width $\leq k$ is well-quasi-ordered by the minor relation.

We will use slight strengthenings of both of these result, Lemma 6.5.3 and Lemma 6.5.11, to show that our two classes of graphs admit extensive tree-decompositions.

In Section 6.10 we will discuss in more detail the connection between our proof and well-quasi-ordering, and indicate how stronger well-quasi-ordering results could be used to prove the ubiquity of larger classes of graphs.

6.5.1. Finitely many thin ends. We will consider the following strengthening of the minor relation.

DEFINITION 6.5.2. Given $\ell \in \mathbb{N}$ an ℓ -pointed graph is a graph G together with a point function $\pi: [\ell] \rightarrow V(G)$. 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 6.5.3. *The set of ℓ -pointed finite graphs is well-quasi-ordered under the relation \preceq_p .*

PROOF. This follows from a stronger statement Robertson and Seymour proved in [137, 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 6.5.4. *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 $i \in \mathbb{N}$:*

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

PROOF. See [85, Satz 3']. \square

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

LEMMA 6.5.5. *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 with root t_0 .*

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 collection of disjoint rays in G . Let (R', \mathcal{W}) be the tree-decomposition of G given by Lemma 6.5.4 where $R' = t'_0 t'_1 \dots$ a ray.

Without loss of generality (taking the union over the first few parts, 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 disjoint family of k rays and $|S(t'_{i-1} t'_i)| = k$ for each $i \in \mathbb{N}$, each vertex in $S(t'_{i-1} t'_i)$ is contained in a unique ray in \mathcal{R} .

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

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

By Lemma 6.5.3 and Remark 6.5.1 there is an n_0 such that for every $n \geq n_0$ there are infinitely many $m > n$ with $(G_n, \pi_n) \preceq_p (G_m, \pi_m)$.

Let $V_{t_0} := \bigcup_{i=0}^{n_0} W_{t'_i}$ and $V_{t_i} := W_{t'_{n_0+i}}$ for all $i \in \mathbb{N}$. We claim that $(R, (V_{t_i} : i \in \mathbb{N}))$ 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. It remains to show that for every $i \in \mathbb{N}$, the bough $B(t_{i-1} t_i)$ is self-similar.

Let $e = t_{i-1} t_i$. Let us label $\mathcal{R} = \{R_{e,s} : s \in S(e)\}$ where $R_{e,s}$ is the unique ray in \mathcal{R} containing s . We wish to show there is a witness W for the self-similarity of $B(e)$ of distance at least n for each $n \in Nbb$. Note that $B(e) = \bigcup_{j \geq 0} G_{n_0+i+j}$. By the choice of n_0 in Remark 6.5.1, there exists $m > i + n$ 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 $B(e)$.

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. In particular, $S(t'_{n_0+m_j-1} t'_{n_0+m_j}) \cup S(t'_{n_0+m_j} t'_{n_0+m_j+1}) \subseteq H_{m_j}$ for each $j \geq 0$.

Hence for, for every $j \in \mathbb{N}$ 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 $IG[B(e)]$.

To prove this claim, for each $j \in \mathbb{N}$ and each $s \in S(t_{j-1} t_j)$, let $R_{p(s)} \in \mathcal{R}$ be the unique ray with $s \in R_{p(s)}$. Then $W'(s) = H_{m_{j-1}}(s) \cup P_{p(s),j} \cup H_{m_j}(s)$ is a connected branch set. Indeed, by construction, every $P_{p,j}$ is a path from $\pi_{n_0+m_{j-1}}(k+p)$ to $\pi_{n_0+m_j}(p)$. And since the H_{m_j} are pointed minors of $G_{n_0+m_j}$, it follows that $\pi_{n_0+m_{j-1}}(k+p(s)) \in H_{m_{j-1}}(s)$ and $\pi_{n_0+m_j}(p(s)) \in H_{m_j}(s)$ are as desired.

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

LEMMA 6.5.6. *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 . Pick a finite set $X \subseteq V$ of vertices separating 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 $G_i := G[C_i \cup S]$. Then each G_i is a locally finite connected one-ended graph, with a thin end ω_i , and hence by Lemma 6.5.5 each of the G_i admits an extensive tree-decomposition (R^i, \mathcal{V}^i) with root $r^i \in V(R^i)$. 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 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

6.5.2. Finite tree-width.

DEFINITION 6.5.7. A rooted tree-decomposition (T, \mathcal{V}) of G is *lean* if for any $k \in \mathbb{N}$, any two nodes $t_1, t_2 \in V(T)$ and any $X_{t_1} \subseteq V_{t_1}, X_{t_2} \subseteq V_{t_2}$ such that $|X_{t_1}|, |X_{t_2}| \geq k$ 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 6.5.8. *Kříž and Thomas [105] showed that if G has tree-width $\leq m$ for some $m \in \mathbb{N}$, then G has a lean tree-decomposition of width $\leq m$.*

LEMMA 6.5.9. *If G is a connected locally finite graph and $(T, (V_t: t \in T))$ a lean tree-decomposition of G such that every V_t is finite, then there is a locally finite subtree S of T such that $(S, (V_t: t \in S))$ is also a lean tree-decomposition of G .*

PROOF. Pick an arbitrary root r of T . We will build recursively finite subtrees of T whose union will be the desired locally finite tree. Let $S_0 = L_0 = \{r\}$. For each $n \in \mathbb{N}$ let L_n be the set of leaves of S_n .

Consider some $t \in L_n$. Since V_t is finite and G is locally finite, the set \mathcal{C}_t of components of $G - V_t$ is finite. Then, for each edge e leaving T_n with $t = e^-$ we have, by the definition of a tree-decomposition, that there is some subset $\mathcal{C}_e \subseteq \mathcal{C}_t$ such that

$$\bigcup \mathcal{C}_e \subseteq B(e) \subseteq \bigcup \mathcal{C}_e \cup V_t.$$

For each of the finitely many sets $\mathcal{C} \subseteq \mathcal{C}_t$ appearing as some \mathcal{C}_e pick an arbitrary e which witnesses this. Let $E_t \subseteq E(T)$ be the set of all e chosen in this way, note that E_t is finite. Let S_{n+1} be $S_n \cup E_t$.

Finally, we let $S := \bigcup_{n \in \mathbb{N}} S_n$. It is simple to check that S is a locally finite tree and that $(S, \{V_t \mid t \in S\})$ is indeed a lean tree-decomposition of G . \square

LEMMA 6.5.10. *Let G be a locally finite, connected graph, and let (T, \mathcal{V}) be a lean tree-decomposition of G with root r and width $\leq m$, with T locally finite. Then there exists a lean tree-decomposition of G with width $\leq m$ 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. Let $D_0 := \{r\}$ and $(T_0, \mathcal{V}_0) := (T, \mathcal{V})$. For $i \in \mathbb{N}$ let

$$D_i := \{e \in E(T_{i-1}) : \text{dist}_{T_i}(r, e^-) = i\}.$$

Construct (T_i, \mathcal{V}_i) from $(T_{i-1}, \mathcal{V}_{i-1})$ by performing the following operation for each edge $e \in D_i$:

Let $t = t_e^+$ and let C_1, \dots, C_n be the connected components of $B(e)$. Replace the subtree T_t with nT_t . For each $s \in T_t$ there are k copies of s in nT_t which we will call s_1, \dots, s_k . For each $s \in T_t$ and $k \in [n]$ let $V_{s_k} := C_k \cap V_s$. Finally, let $\hat{T} = \bigcup_{i \in \mathbb{N}} T_i[\{t \in T_i \mid d_{T_i}(r, t) \leq i\}]$ and $\hat{\mathcal{V}} = (V_t \mid t \in \hat{T})$.

It is simple to check that $(\hat{T}, \hat{\mathcal{V}})$ is a tree-decomposition of width $\leq m$, that \hat{T} is locally finite, and by construction $B(e)$ is connected for each $e \in E(T)$. Furthermore, suppose $k \in \mathbb{N}$, $t_1, t_2 \in \hat{T}$ and $X_{t_1} \subseteq \hat{V}_{t_1}, X_{t_2} \subseteq \hat{V}_{t_2}$ are such that $|X_{t_1}|, |X_{t_2}| \geq k$. By construction, there are nodes t'_1 and t'_2 of T such that $X_{t_1} \subseteq \hat{V}_{t_1} \subseteq V_{t'_1}, X_{t_2} \subseteq \hat{V}_{t_2} \subseteq V_{t'_2}$. Thus, since (T, \mathcal{V}) is lean, either there is a vertex t' of T between t'_1, t'_2 such that $|V_{t'}| < k$ or there are k disjoint paths between X_{t_1} and X_{t_2} in G . However, in the first case, by construction, there also is a node t of \hat{T} between t_1 and t_2 such that $\hat{V}_t \subseteq V_{t'}$. Thus, $(\hat{T}, \hat{\mathcal{V}})$ is indeed lean.

Suppose there is an edge $e = st \in \hat{T}$, such that $B(e)$ is finite, but \hat{T}_t is infinite. Since $\hat{V}_x \subseteq B(e)$ for any vertex $x \in V(\hat{T}_t)$, the set $\{\hat{V}_x : x \in V(\hat{T}_t)\}$ is finite. Hence, there is a finite subtree $\bar{T}_t \subseteq \hat{T}_t$ which contains at least one node for each of these bags. Let us replace, for each minimal $e \in E(T)$ with $B(e)$ finite, the subtree \hat{T}_t with \bar{T}_t , to give a tree \bar{T} , and let $\bar{\mathcal{V}} = (V_t : t \in V(\bar{T}))$. Then, $(\bar{T}, \bar{\mathcal{V}})$ is a lean-tree decomposition with width $\leq m$ such that \bar{T} is locally finite and every bough $B(e)$ is connected. Moreover it has the following property

(\dagger) For every $t \in V(\bar{T})$, if \bar{T}_t is infinite, then so is $B(e)$.

Finally, suppose for a contradiction that there are vertices which appear in infinitely many parts of $(\bar{T}, \bar{\mathcal{V}})$. Let X be a \subseteq -maximal set of vertices appearing as a subset in infinitely many parts of $(\bar{T}, \bar{\mathcal{V}})$. Note that X is finite, since every part has size at most m . Since \bar{T} is locally finite and $(\bar{T}, \bar{\mathcal{V}})$ is a tree-decomposition, there is a ray R in \bar{T} such that X appears as a subset in every part corresponding to a node of R . We may assume without loss of generality that $R \subseteq \bar{T}_r$, where $r = \text{init}(R)$. Since for each $t \in R$ the subtree

\bar{T}_t contains a tail of R , it is infinite, and hence by (\dagger) $B(e)$ is infinite and $X \subseteq B(e)$ for every $e \in R$. Since $B(e)$ is connected, X has a neighbour in $B(e) \setminus X$. However, since G is locally finite, X has only finitely many neighbours, and by \subseteq -maximality of X each neighbour appears in only finitely many parts of $(\bar{T}, \bar{\mathcal{V}})$, and so in only finitely many sets $B(e)$ with $e \in R$. This contradicts the fact that X has a neighbour in every $B(e) \setminus X$. \square

LEMMA 6.5.11. *For all $k, \ell \in \mathbb{N}$ the class of ℓ -pointed graphs with tree-width $\leq k$ is well-quasi-ordered under the relation \preceq_p .*

PROOF. This is a consequence of a result of Thomas [149]. \square

LEMMA 6.5.12. *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 Lemma 6.5.9 there is a lean tree-decomposition (T, \mathcal{V}) of G with width m , such that T is a locally finite tree with root r . By Lemma 6.5.10 we may assume that every vertex is contained in only finitely many parts of this tree-decomposition.

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

Since (T, \mathcal{V}) is lean, there are d_ϵ disjoint paths between $S(t_{i_k-1}^\omega t_{i_k}^\omega)$ and $S(t_{i_{k+1}-1}^\omega t_{i_{k+1}}^\omega)$ for every $k \in \mathbb{N}$. Moreover, since each $S(t_{i_k-1}^\omega t_{i_k}^\omega)$ 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 set X separating some pair of rays R and R' . 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 $t \in T_{t_{i_k-1}^\epsilon}$. By construction R and R' have tails in $B(t_{i_{k+1}-1}^\omega t_{i_{k+1}}^\omega)$ 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 $R_j^\epsilon \cap S(t_{i_k-1}^\epsilon t_{i_k}^\epsilon)$.

By Lemma 6.5.11 and Remark 6.5.1, the sequence $(G[B(t_{i_k-1}^\epsilon t_{i_k}^\epsilon)], \pi_{i_k}^\epsilon)_{k \in \mathbb{N}_{>0}}$ has an increasing subsequence $(G[B(t_{i-1}^\epsilon t_i^\epsilon)], \pi_i^\epsilon)_{i \in I_\epsilon}$, i.e. for any $k, j \in I_\epsilon, k < j$ 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, also each set $\bigcup_{t \in C} V_t$ is finite.

Let us define a tree decomposition (T', \mathcal{V}') on $T' = T/\mathcal{C}(T^-)$ where $V_{t'}' = \bigcup_{t \in t'} V_t$. We claim this is an extensive tree-decomposition.

Clearly, T' is a locally finite tree, and each part of (T', \mathcal{V}') is finite, and every vertex of G is contained in only finitely many parts of the tree-decomposition. Give $e \in E(T')$ 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' . Since $G[B(e)] = G[B(t_{k-1}^\epsilon t_k^\epsilon)]$ and $G[B(e')] = G[B(t_{k'-1}^\epsilon t_{k'}^\epsilon)]$ we have $(G[B(e)], \pi_k^\epsilon) \preceq_p (G[B(e')], \pi_{k'}^\epsilon)$, witnessing the self-similarity of $B(e)$ towards ω_e with the rays $(R_{e,j} : j \in [d_\epsilon])$. \square

REMARK 6.5.13. *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.*

6.5.3. Special graphs. We note that, whilst Lemmas 6.5.6 and 6.5.12 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 6.5.14. *The grid $\mathbb{Z} \times \mathbb{Z}$ has an extensive tree-decomposition, as can be seen in Figure 6.2. 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 will self similar.*

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 [32].

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

QUESTION 6.5.15. Do all locally finite graphs admit an extensive tree-decomposition?

6.6. 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 of some host graph Γ . In particular we are interested in how the rays of these minors occupy a given end ϵ of Γ . In [32] we established the distinction between *pebbly* and *non-pebbly* ends, cf. Definition 6.6.4. We showed that




FIGURE 6.2. In the grid the boughs are self-similar.

the existence of a pebbly end of Γ already guarantees the existence of a K^{\aleph_0} -minor in Γ , and therefore the following corollary holds:

COROLLARY 6.6.1 ([32, Corollary 6.4]). *Let Γ be a graph with a pebbly end ω and let G be a countable graph. Then $\aleph_0 G \preceq \Gamma$.*

We will now analyse the structure of non-pebbly ends and give a description of their shape. For a fixed set of start vertices we will consider the possible families of disjoint rays with these start vertices. This shall be made precise in the definition of *polypods*, cf. Definition 6.6.7 below. We will investigate how these rays relate in terms of connecting paths between them and see that, due to the non-pebbly structure of the end, the structure of possible connections between the rays is somewhat restricted.

6.6.1. Pebble Pushing. Given a path P with end-vertices s and t we say *the orientation of P from s to t* to mean the total order 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.

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 3 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 $(f(x), f(y), f(z))$.

A permutation of a finite set X is a bijection $\nu: X \rightarrow X$. A sequence $(x_1 \dots x_n)$ of distinct elements of X is called a *cycle* of ν if $\nu(x_n) = x_1$ and $\nu(x_i) = x_{i+1}$ for all $i \in \{1, \dots, n-1\}$. In this case n is called the *length* of the cycle, a cycle of length 1 is called *trivial*. The term $(x_1 \dots x_n)$ is also used to denote the bijection ν which contains the cycle $(x_1 \dots x_n)$ and otherwise is the identity on $X \setminus \{x_1, \dots, x_n\}$. It is a well-known fact that every bijection can be written as a product of (disjoint) cycles.

We utilise the following results and definitions from [32].

DEFINITION 6.6.2 (Pebble-pushing game). Let $G = (V, E)$ be a graph. We call a tuple $(x_1, \dots, x_k) \in V^k$ a *game state (of order k)* 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 $Y = (y_1, \dots, y_k)$, we imagine k labelled pebbles placed on the vertices (y_1, \dots, y_k) . A move for a game state in the pebble-pushing game consists of 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 $Z = (z_1, \dots, z_k)$ 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 $X = (x_1, \dots, x_k)$ 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, \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 a X_i -move for all $i \in [n-1]$, that is if there 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 of order k .

LEMMA 6.6.3 ([32, Lemma 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 $RG(S_j: j \in [m])$, the map σ defined via $\sigma(i) := z_i$ for every $i \in [k]$ is a transition function⁴ from $(S_i: i \in [k])$ to $(S_j: j \in [m])$.

DEFINITION 6.6.4 (Pebbly ends). Let Γ be a graph and ω an end of Γ . We say ω is *pebbly* if for every k there is an $n \geq k$ and a family $\mathcal{R} = \{R_1, \dots, R_n\}$ of disjoint rays in ω such that $RG(R_i: i \in [n])$ is k -pebble-win. If for some k there is no such family \mathcal{R} we say ω is *not k -pebble-win*.

LEMMA 6.6.5 ([32, Lemma 6.3]). Let Γ be a graph and let $\omega \in \Omega(\Gamma)$ be a pebbly end. Then $K_{\aleph_0} \preceq \Gamma$.

⁴See Definition 6.2.3.

Recall that a path $P = v_0v_1 \dots v_n$ in a graph G is called a *bare* if all its inner vertices have degree 2 in G .

COROLLARY 6.6.6 ([32, Corollary 5.2]). *Let ω be an end of Γ which is not k -pebble-win 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 $P = p_1 \dots p_n$ in $RG(R_i : i \in [m])$ such that $n \geq m - k$. Furthermore, either each edge in P is a bridge, or $RG(R_i : i \in [m])$ is a cycle.*

6.6.2. Polypods.

DEFINITION 6.6.7. Given an end ω of a graph Γ , a *polypod* (for ω 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 rays to ω , where R_y begins at y and all the R_y are disjoint from X . Such a family (R_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 6.6.8. We say that a polypod (X, Y) for ω in Γ is *tight* if its connection graph is minimal amongst connection graphs of polypods for ω in Γ with respect to the spanning isomorphic subgraph relation, i.e. for no other polypod (X', Y') for ω 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 6.6.9. *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 $R_y \dot{v}_y$. 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 6.6.10. *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 paths and a path Q_n such that:

- Each P_y^n starts at y .
- Each P_y^n has length at least n .
- For $m \leq n$, the path P_y^n extends P_y^m .
- If $n \in A_{vw}$ then Q_n is a path from P_v^n to P_w^n .
- If $n \in A_{vw}$ then Q_n meets no P_y^m with $y \notin \{v, w\}$.
- All the Q_n are disjoint.
- All the P_y^n and all the Q_n are disjoint from X .
- For any 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 only meet the Q_m with $m \leq n$ in inner vertices of the P_y^n .

It is clear that if we can do this then we will obtain a family of tendrils by letting R_y be the union of all the P_y^n . 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 the whole of $K_{X,Y}$, as required.

So it remains to explain 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 of length at least 1 and containing all vertices of $Q_1 \cap R_y^1$. This choice of the P_y^1 and of Q_1 clearly satisfies the conditions above.

Now 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 . Then by Lemma 6.6.9 (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 $y P_y^n y_n S_{y_n}^{n+1}$ and let P_y^{n+1} be an initial segment of R_y^{n+1} long enough to include P_y^n , of length at least $n+1$, and containing all vertices of $Q_{n+1} \cap R_y^{n+1}$ as inner vertices. This completes the recursion step, and so the construction is complete. \square

LEMMA 6.6.11. *If (X, Y) is a polypod of order n for ω in Γ with connection graph $K_{X,Y}$ then for any set of n disjoint ω -rays $(R_i : i \in [n])$ in Γ , $RG(R_i : i \in [n]) \subsetneq K_{X,Y}$.*

PROOF. If we apply the Weak Linking Lemma 6.2.2 to the rays $(R_i: i \in [n])$ and a family of tendrils for (X, Y) , together with the finite set X , we obtain a family of tendrils for (X, Y) whose tails coincide with that of $(R_i: i \in [n])$. Hence, the ray graph of these tendrils is $RG(R_i: i \in [n])$ and so $RG(R_i: i \in [n]) \subseteq K_{X,Y}$. \square

COROLLARY AND DEFINITION 6.6.12. *Any two polypods for ω 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 ω in Γ as the n^{th} shape graph of the end ω . \square

6.6.3. Frames. Akin to the transition boxes defined in Lemma 6.3.17 we want to consider *frames*, finite subgraphs which are just large enough to include a linkage which, say, induces a transition function of the family of tendrils of some polypod. This will allow us to reason about transition functions in terms of graph automorphisms.

DEFINITION 6.6.13. 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 . We call the Y -frame *strait* if it has at most one weave graph and at most one weave pattern. 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 6.6.14. *Let (X, Y) be a polypod for ω 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 6.2.2 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 $RG(R_y: y \in Y)$ 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 $R_y\alpha(y)Q_y\beta(\pi_{\mathcal{Q}}(y))R_{\pi_{\mathcal{Q}}(y)}$ shows that $K_{\mathcal{Q}}$ is a subgraph of $K_{X,Y}$ and that the inverse image of any edge of $K_{X,Y}$ under $\pi_{\mathcal{Q}}$ is again an edge of $K_{X,Y}$, from which it follows that $\pi_{\mathcal{Q}}$ is an automorphism of $K_{X,Y}$. \square

COROLLARY 6.6.15. *Let (X, Y) be a polypod for ω 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. Pick a linkage $(P_y: y \in Y)$ from \mathcal{R} to itself after X inducing σ . Let L be a finite subgraph of Γ containing all 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 frame on L which exists by Lemma 6.6.14 has the desired properties. \square

LEMMA 6.6.16. *Let (X, Y) be a polypod for ω 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 $RG(R_y: y \in Y)$ 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 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 $RG(R_y: y \in Y)$ 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

$$X' \cup \bigcup_{e \in E(RG(R_y: y \in Y))} P_e \cup \bigcup_{f \in E(K_{X,Y})} Q_f.$$

Let the initial vertex of P'_y be $\gamma(y)$ and the end vertex be $\beta(y)$. Let $\pi(y)$ be the element of Y with $\beta(y)$ on $R_{\pi(y)}$. Let L be the subgraph of Γ containing all paths of the forms $S_y\gamma(y), R_{\pi(y)}\beta(y), P'_y, P_e$ and Q_f .

Letting α be the identity function on Y , it follows from Lemma 6.6.14 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 of the R_y . The two must be equal since whenever for two of the R_y there is any path between them, disjoint from all the other R_y and from X' , then there are infinitely many disjoint such paths. \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 6.6.17. 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 6.6.18. *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 $\ell \in \mathbb{N}$ 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 contains at least 3 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 are preserved by locally generated automorphisms.*

LEMMA 6.6.19. *Let (L, α, β) be a Y -frame which is K -spartan but not strait. Then each of its weave graphs includes a turnable cycle or a flippable edge of K and for any two of its weave patterns π and π' the automorphism $\pi^{-1} \cdot \pi'$ of K is locally generated.*

PROOF. Suppose not for a contradiction, and let (L, α, β) be a counterexample in which $|E(L)|$ is minimal. Note that, as L is not strait, there are either at least two weave patterns for L or there are at least two weave graphs for L . Thus, we can find weaves $\mathcal{P} = (P_y : y \in Y)$ and $\mathcal{Q} = (Q_y : y \in Y)$ such that either $K_{\mathcal{P}} \neq K_{\mathcal{Q}}$ or $\pi_{\mathcal{P}} \neq \pi_{\mathcal{Q}}$ and such that either $K_{\mathcal{Q}}$ includes no turnable cycle or flippable edge or $\pi_{\mathcal{P}}^{-1} \cdot \pi_{\mathcal{Q}}$ is not locally generated. Furthermore, by exchanging \mathcal{P} and \mathcal{Q} if necessary, we may assume that $K_{\mathcal{P}}$ is not a proper subgraph of $K_{\mathcal{Q}}$.

Each edge of L is in one 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 appearing in neither \mathcal{P} nor \mathcal{Q} are isolated and so may be ignored. Thus we may suppose that each edge appears in precisely one of \mathcal{P} or \mathcal{Q} , and that each vertex appears in both.

Let Z be the set of those $y \in Y$ such that $\alpha(y) \neq \beta(y)$. For any $z \in Z$ let $\gamma(z)$ be the second vertex of 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 $i \leq n$ 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, $(\alpha(y)\gamma(y)\hat{P}_y)_{y \in Y}$ 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_i} be $\alpha'(y_i)Q_{\sigma(y_i)}$ for any $y \in Y$. Then we have $\pi_{Q'} = \pi_Q \cdot \sigma$ and so $\sigma = \pi_Q^{-1} \cdot \pi_{Q'}$ is an automorphism of K . 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 so σ is local. Hence K_Q includes a turnable cycle or a flippable edge. By the minimality of $|E(L)|$ we know that $\pi_{P'}^{-1} \cdot \pi_{Q'}$ is locally generated and hence so is $\pi_{P'}^{-1} \cdot \pi_Q = \pi_{P'}^{-1} \cdot \pi_{Q'} \cdot \sigma^{-1}$. This is the desired contradiction. \square

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

LEMMA 6.6.20. *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 6.6.16 there is some $K_{X,Y}$ -spartan Y -frame for which both $K_{X,Y}$ and the ray graph 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 6.6.19 and so since it is a subgraph of $K_{X,Y}$ it must be the whole of $K_{X,Y}$. Similarly Lemma 6.6.19 together with Corollary 6.6.15 shows that all transition functions must be locally generated and so must preserve the orientation. \square

LEMMA 6.6.21. *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 6.6.15 and 6.6.19 any transition function must be a locally generated automorphism of $K_{X,Y}$, and so by Remark 6.6.18 it cannot move the vertices of the bare path, which are vertices of degree one or cutvertices. \square

6.7. 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. We fix a graph Γ together with a thick end ω of Γ . If ω is pebbly then $K_{\aleph_0} \preceq \Gamma$ by Lemma 6.6.5, and every locally finite graph G satisfies $\aleph_0 G \preceq K_{\aleph_0} \preceq \Gamma$.

So in the following we further restrict ourselves to the case that ω is not pebbly; for this section we fix a number N such that there is no family $(R_i: i \in [n])$ of disjoint rays with $n \geq N$ such that $RG(R_i: i \in [n])$ is 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 6.6.6: The grid-like and the half-grid-like case.

6.7.1. Grid-like ends. The first case is 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 ω is *grid-like* if the $(N + 2)^{\text{nd}}$ shape graph for ω is a cycle. For the rest of this subsection we will assume that ω 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 6.7.1. *Any ray graph K for a set $(R_i : i \in I)$ of ω -rays in Γ with $|I| \geq N + 2$ is a cycle.*

PROOF. Let $(T_y : y \in Y)$ be a family of tendrils for (X, Y) obtained by transitioning from the S_y to the R_i after X along a linkage, and let $\sigma : Y \rightarrow I$ be the function induced by this linkage. Then by Lemma 6.6.20 the ray graph of the T_y is the cycle $K_{X,Y}$. We know by Corollary 6.6.6 that K includes a bare path P such that $|V(P)| \geq |V(K)| - N$. Thus there are distinct vertices $y_1, y_2 \in Y$ with $\sigma(y_1), \sigma(y_2) \in P$ and no other vertex in the image of σ between them on P . Then for any other vertex y of Y there are paths from y to y_1 avoiding y_2 and from y to y_2 avoiding y_1 in $K_{X,Y}$. Hence there are paths from $\sigma(y)$ to each of $\sigma(y_1)$ and $\sigma(y_2)$ avoiding $\sigma(y_1)P\sigma(y_2)$. Thus none of the edges of $\sigma(y_1)P\sigma(y_2)$ is a bridge, so by Corollary 6.6.6 again K is a cycle. \square

We will now choose cyclic orientations of all these cycles 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 ω -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 6.7.2. *For any such family $(R_i : i \in I)$ of at least $N + 3$ disjoint ω -rays there is precisely one correct cyclic orientation of its ray graph.*

PROOF. That there is at least one is clear by Lemma 6.2.2. 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 6.6.3 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 doesn't affect the cyclic order, we may assume without loss of generality that $I' = I$. By Lemma 6.2.2 again, there is some transition function τ from the R_i to the S_y . By Lemma 6.6.20 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 6.7.3. *Any transition function between ray graphs on at least $N + 3$ rays preserves the correct orientations of the cycles.* \square

6.7.2. Half-grid-like ends. In this subsection we suppose that ω 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.

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 6.6.6 there is a bare path of length at least 2 in $K_{X,Y}$ of which all edges are bridges. Let y_1y_2 be any edge of that path. Without loss of generality we have $C^{y_1y_2}(K_{X,Y}) \neq \emptyset$.

Let $(R_i : i \in I)$ be a family of disjoint rays with $|I| \geq N + 3$ and let K be its ray graph.

REMARK 6.7.4. *For any transition function σ from the S_y to the R_i we have $\sigma[C^{y_1y_2}(K_{X,Y})] \subseteq C^{\sigma(y_1)\sigma(y_2)}(K)$ and $\sigma[C^{y_2y_1}(K_{X,Y})] \subseteq C^{\sigma(y_2)\sigma(y_1)}(K)$. Thus $\sigma[K_{X,Y}]$ and $K^{\sigma(y_1)\sigma(y_2)}$ meet precisely in $\sigma(y_1)$ and $\sigma(y_2)$.*

LEMMA 6.7.5. *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 6.6.3 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 6.7.4 we have $\sigma[C^{y_1y_2}(K_{X,Y})] \subseteq C^{\sigma(y_1)\sigma(y_2)}(K) \cap C^{\sigma'(y_1)\sigma'(y_2)}(K) = C^{\sigma(y_1)\sigma(y_2)}(K) \cap C^{\sigma(y_2)\sigma(y_1)}(K) = \emptyset$. So this is impossible, and $K^{\sigma(y_1)\sigma(y_2)} = P$. The last sentence of the lemma now follows from the definition of $K^{\sigma(y_1)\sigma(y_2)}$. \square

Now we fix a transition function σ_{\max} so that the path $P := K^{\sigma_{\max}(y_1)\sigma_{\max}(y_2)}$ is as long as possible. If $\sigma_{\max}[C^{y_1y_2}(K_{X,Y})]$ were a proper subset of $C^{\sigma_{\max}(y_1)\sigma_{\max}(y_2)}(K)$ then we would be able to use Lemma 6.6.3 to produce a transition function in which this path is longer. So we must have $\sigma_{\max}[C^{y_1y_2}(K_{X,Y})] = C^{\sigma_{\max}(y_1)\sigma_{\max}(y_2)}(K)$ and similarly $\sigma_{\max}[C^{y_2y_1}(K_{X,Y})] = C^{\sigma_{\max}(y_2)\sigma_{\max}(y_1)}(K)$.

We call P the *central* path of K and the orientation from $\sigma_{\max}(y_1)$ to $\sigma_{\max}(y_2)$ the *correct* orientation. In the following lemma we use this orientation to determine which vertices appear before which along P .

LEMMA 6.7.6. *For any two vertices v_1 and v_2 of K , there is a transition function $\sigma: K_{X,Y} \rightarrow K$ 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 6.6.3 to σ_{\max} . For the ‘only if’ direction, we begin by setting $c_1 = |C^{y_1 y_2}(K_{X,Y})|$ and $c_2 = |C^{y_2 y_1}(K_{X,Y})|$. We enumerate $C^{y_1 y_2}(K_{X,Y})$ as $y_3 \dots y_{c_1+2}$ and $C^{y_2 y_1}(K_{X,Y})$ 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 achievable in the $(\sigma_{\max}(y_1), \dots, \sigma_{\max}(y_{N+2}))$ pebble pushing game must have the following 3 properties, since they are preserved by any single move:

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

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 $\{x_1, \dots, x_{N+1}\} = \sigma[Y]$. By Lemma 6.6.3 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 6.6.21 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 $\sigma(y_2) = \sigma'(y_2) = x_2$. 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

COROLLARY 6.7.7. *Each ray graph contains a unique central path with a correct orientation and all transition functions between ray graphs send vertices of the central path to vertices of the central path and preserve the correct orientation.*

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, could be derived from earlier work of Diestel and Thomas [60], 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, if not more complicated, and so we have opted for a straightforward and self-contained presentation.

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

DEFINITION 6.7.8. Given a graph G , an end ω and three rays R, S, T in ω 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 $G - S$.

For the following, recall the definition of *ray graph* in Definition 6.2.4.

LEMMA 6.7.9. *Let G be a graph, ω an end of G and $(R_i)_{i \in I}$ a finite family of disjoint ω -rays. If, for some $i_1, i_2, j \in I$, the vertices i_1 and i_2 belong to different components of $RG((R_i)_{i \in I}) - j$, then R_j separates R_{i_1} from R_{i_2} .*

PROOF. If R_{i_1} and R_{i_2} belong to the same end of $G - V(R_j)$, there are infinitely many disjoint paths between R_{i_1} and R_{i_2} in $G - V(R_j)$. Hence, by the pigeonhole principle there are indices j_1 and j_2 belonging to different components of $RG((R_i)_{i \in I}) - j$, such that these disjoint paths induce infinitely many disjoint paths from R_{j_1} to R_{j_2} all disjoint from all other R_i . Thus there is an edge from j_1 to j_2 in $RG((R_i)_{i \in I})$ contradicting the assumption that j disconnects j_1 from j_2 . \square

LEMMA 6.7.10. *Consider three rays R, S, T belonging to the same end ω of some graph G . If 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 ω , there are infinitely many disjoint paths between them. As S separates R from T , 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

DEFINITION 6.7.11. Given a graph G and two (possibly infinite) vertex-sets X and Y , we say that an end ω of $G - X$ is a *sub-end* of an end ω' of $G - Y$ if every ray in ω has a tail in ω' .

DEFINITION 6.7.12. Let ω be a half-grid-like end, let R be an ω -ray. We say R is a *core ray* (of ω) if there is a finite family $\mathcal{R} = (R_i : i \in I)$ of disjoint ω -rays with $R = R_c$ for some $c \in I$ such that c lies on, but is not an endpoint, of the central path of \mathcal{R} .

LEMMA 6.7.13. *Let R be a core ray of ω . Then in $G - R$ the end ω splits into precisely two different ends. (That is, there are two ends ω' and ω'' of $G - R$ such that every ω -ray in $G \setminus V(R)$ is in ω' or ω'' .)*

PROOF. Let $\mathcal{R} = (R_i : i \in I)$ be a family witnessing that $R = R_c$ for some $c \in I$ is a core ray. Then there are exactly two ends in $G \setminus V(R)$ which contain rays in \mathcal{R} , since connected components of $RG(\mathcal{R})$ when we delete the vertex corresponding to R are equivalent sets of rays in $G \setminus V(R)$ and more over, no two of these connected components can belong to the same end of $G \setminus V(R)$ by Lemma 6.7.9.

Suppose there is a third end in $G \setminus V(R)$ that contains an ω -ray S . We first claim that there is a tail of S which is disjoint from \mathcal{R} . Indeed, clearly S is disjoint from R , and if S met $\bigcup \mathcal{R}$ infinitely often then it would meet some $R_i \in \mathcal{R}$ infinitely often, and hence lie in the same end of $G \setminus V(R)$ as R_i . Let us assume then that S is disjoint from \mathcal{R} .

Let us consider the ray graph $RG(\mathcal{R} \cup \{S\})$. Again, if S is adjacent to any ray except R in the ray graph, it would lie in the same end as some ray in \mathcal{R}_J in $G \setminus V(R)$.

Since S is an ω -ray the ray graph is connected, and hence S is adjacent to R , and R is still connected to its neighbours in $RG(\mathcal{R})$. However, $\mathcal{R} \cup \{S\}$ is also a family that witnesses that $R = R_c$ is a core ray and hence c has degree two in $RG(\mathcal{R} \cup \{S\})$, a contradiction. \square

Given a family of rays $(R_i)_{i \in I}$ witnessing that $R = R_c$ is a core ray, we denote by $\top(R, (R_i)_{i \in I})$ the end of $G - V(R)$ containing rays R_i satisfying $i < c$ and with $\perp(R, (R_i)_{i \in I})$ the end containing rays R_i satisfying $i > c$.

LEMMA 6.7.14. *Let R and S be disjoint core rays of ω . Let us suppose that ω splits in $G - S$ in ω'_S and ω''_S and in $G - R$ in ω'_R and ω''_R . If R belongs to ω'_S and S belongs to ω'_R , then ω''_S is a sub-end of ω'_R and ω''_R is a sub-end of ω'_S .*

PROOF. Let T be a ray in ω''_S . As R belongs to a different end of $G - S$ than T , there is a tail of T which is disjoint from R . Thus, we may assume that T and R are disjoint. As S separates R from T , by Lemma 6.7.10, R does not separate S from T , hence T belongs to ω'_R . \square

LEMMA AND DEFINITION 6.7.15. *Let $\mathcal{R}_1 = (R_i : i \in I_1)$, $\mathcal{R}_2 = (R_i : i \in I_2)$ be two finite families of disjoint ω -rays both witnessing that for some $c \in I_1 \cap I_2$ the ray R_c is a core ray in ω . Then $\top(R, (R_i)_{i \in I_1}) = \top(R, (R_i)_{i \in I_2})$ and $\perp(R, (R_i)_{i \in I_1}) = \perp(R, (R_i)_{i \in I_2})$.*

We therefore write $\top(\omega, R)$ for the end $\top(R, (R_i)_{i \in I_1})$ and $\perp(\omega, R)$ respectively, i.e. $\top(\omega, R)$ is the end of $G - R$ containing rays that appear on the central path of some ray graph before R according to the correct orientation and $\perp(\omega, R)$ is the end of $G - R$ containing rays that appear on the central path of some ray graph after R according to the correct orientation. Note that $\top(\omega, R) \cap \perp(\omega, R) = \emptyset$.

PROOF. Suppose, this is not the case, hence $\omega_1 := \top(R_c, (R_i)_{i \in I_1}) = \perp(R_c, (R_i)_{i \in I_2})$ and $\omega_2 := \perp(R_c, (R_i)_{i \in I_1}) = \top(R_c, (R_i)_{i \in I_2})$. Let $\mathcal{R}_{I_1\omega_1}$ be the set of rays in \mathcal{R}_1 belonging to ω_1 . Let $\mathcal{R}_{I_1\omega_2}, \mathcal{R}_{I_2\omega_1}$ and $\mathcal{R}_{I_2\omega_2}$ be defined accordingly. If $|\mathcal{R}_{I_1\omega_1}| > |\mathcal{R}_{I_2\omega_1}|$ we define \mathcal{R}_{ω_1} to be $\mathcal{R}_{I_1\omega_1}$, otherwise $\mathcal{R}_{\omega_1} = \mathcal{R}_{I_2\omega_1}$. Let \mathcal{R}_{ω_2} be defined similarly.

Let us consider $\mathcal{R} := \mathcal{R}_{\omega_1} \cup \mathcal{R}_{\omega_2} \cup \{R_c\}$. After replacing some of the rays with tails, this is a collection of disjoint rays, so let us assume that \mathcal{R} itself is a family of disjoint rays. There is a transition function from \mathcal{R}_{I_1} to \mathcal{R} mapping R_c to itself, every ray in $\mathcal{R}_{I_1\omega_1}$ to a ray in \mathcal{R}_{ω_1} and every ray in $\mathcal{R}_{I_1\omega_2}$ to a ray in \mathcal{R}_{ω_2} :

Consider a finite separator X separating ω_1 from ω_2 in $G - V(R_c)$. Consider linkages after X in $G - V(R_c)$ from \mathcal{R}_{ω_1} to \mathcal{R}_{ω_1} and from \mathcal{R}_{ω_2} to \mathcal{R}_{ω_2} . Pairs of such linkages can be combined to suitable linkages on G , inducing a transition function which is as desired.

Similarly there is a transition function from \mathcal{R}_{I_2} to \mathcal{R} mapping R_c to itself, every ray in $\mathcal{R}_{I_2\omega_1}$ to a ray in \mathcal{R}_{ω_1} and every ray in $\mathcal{R}_{I_2\omega_2}$ to a ray in \mathcal{R}_{ω_2} .

These transition functions preserve the central path, thus c lies on the central path of $RG(\mathcal{R})$. Moreover, \mathcal{R} also witness that R_c is a core ray. However, the first transition function shows that $\omega_1 = \top(R_c, \mathcal{R})$ whereas the second one shows that $\omega_2 = \top(R_c, \mathcal{R})$, contradicting the assumption that $\omega_1 \neq \omega_2$. \square

LEMMA AND DEFINITION 6.7.16. *Let $\text{core}(\omega)$ denote the set of core rays in ω . We define a partial order \leq_ω on $\text{core}(\omega)$ by*

$$R \leq_\omega S \text{ if and only if either } R = S,$$

$$\text{or } R \text{ and } S \text{ have disjoint tails } xR \text{ and } yS \text{ and } xR \in \top(\omega, yS)$$

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

PROOF. For the anti-symmetry let us suppose that R and S are disjoint rays such that $R \leq_\omega S$ and $S \leq_\omega R$. Hence, $R \in \top(\omega, S)$ as well as $S \in \top(\omega, R)$. Let \mathcal{R}_S be a family of rays witnessing that S is a core ray and \mathcal{R}_R a family witnessing that R is a core ray. By Lemma 6.7.14, $\perp(\omega, S)$ is a sub-end of $\top(\omega, R)$ and $\perp(\omega, R)$ is a sub-end of $\top(\omega, S)$. Let $\mathcal{R}_{\perp(S)}$ be the subset of \mathcal{R}_S of rays, which belong to $\perp(\omega, S)$. Let $\mathcal{R}_{\perp(R)}$ be defined accordingly. After replacing rays with tails all rays in $\mathcal{R} := \mathcal{R}_{\perp(S)} \cup \mathcal{R}_{\perp(R)} \cup \{R\} \cup \{S\}$ are pairwise disjoint. More over, R and S both lie on the central path of $RG(\mathcal{R})$ and are both not endpoints of this central path. Thus either $S \in \perp(\omega, R)$ or $R \in \perp(\omega, S)$ contradicting Lemma 6.7.15.

For the transitivity, let us suppose that R, S, T are rays, such that $R \leq_\omega S$ and $S \leq_\omega T$. We may assume that R and S , and S and T are disjoint. As \leq_ω is anti-symmetric, it is $T \not\leq_\omega S$, hence $T \in \perp(\omega, S)$. Thus, R and T belong to different ends of $G - S$, thus we may assume that they are also disjoint. As S therefore separates R from T , by Lemma 6.7.10, T does not separate S from R . Thus, R and S belong to the same end of $G - T$. Hence $R \in \top(\omega, T)$. \square

REMARK 6.7.17. *Let $R, S \in \text{core}(\omega)$ and let \mathcal{R} be a finite family of disjoint ω -rays.*

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

LEMMA 6.7.18. *Let $R, S \in \text{core}(\omega)$. Let $Z \subseteq V(G)$ be a finite set such that $\top(\omega, S)$ and $\perp(\omega, S)$ are separated by Z in $G - V(S)$. Let $H \subseteq G - Z$ be a connected subgraph which is disjoint to S and contains R , and let $T \subseteq H$ be some core ω -ray. Then S is in the same relative \leq_ω -order to T as to R .*

PROOF. Assume $S \leq_\omega R$ and hence $R \in \top(\omega, S)$. Since H is connected, we obtain that $T \in \top(\omega, S)$ as well and hence $S \leq_\omega T$. The other case is analogous. \square

LEMMA AND DEFINITION 6.7.19. *Let \mathcal{R} be a finite family of disjoint core ω -rays. Then there exists a family \mathcal{R}' of disjoint ω -rays such that $RG(\mathcal{R})$ is precisely the inner vertices of the central path of $RG(\overline{\mathcal{R}})$. Even though such a family is not unique, we denote by $\overline{\mathcal{R}}$ an arbitrary such family.*

DEFINITION 6.7.20. If \mathcal{P} is a linkage from \mathcal{R} to \mathcal{S} then a *sub-linkage* of \mathcal{P} is just a subset of \mathcal{P} , considered as a linkage from the corresponding subset of \mathcal{R} to \mathcal{S} .

REMARK 6.7.21. *A sub-linkage of a transitional linkage is transitional.*

PROOF. By Remark 6.7.17(2) the rays in \mathcal{R} are linearly ordered by \leq_ω . Let R denote the \leq_ω -smallest and S denote the \leq_ω -greatest element of \mathcal{R} . As in the proof of Lemma 6.7.16, consider the sets $\mathcal{R}_{\perp(R)}$ and $\mathcal{R}_{\top(S)}$, which are without loss of generality minimal with respect to their defining property. Now $\mathcal{R}_{\perp(R)} \subseteq \perp(\omega, R)$ and $R' \in \top(\omega, R)$ for every $R' \in \mathcal{R} \setminus \{R\}$ and hence tails of $\mathcal{R}_{\perp(R)}$ are disjoint to $\bigcup \mathcal{R}$. Analogously, $\mathcal{R}_{\top(S)} \subseteq \top(\omega, S)$ and $R' \in \perp(\omega, S)$ for every $R' \in \mathcal{R} \setminus \{S\}$ and hence tails of $\mathcal{R}_{\top(S)}$ are disjoint to $\bigcup \mathcal{R}$. Finally, $\mathcal{R}_{\top(S)} \subseteq \top(\omega, R)$ and $\mathcal{R}_{\perp(R)} \subseteq \perp(\omega, S)$ by Lemma 6.7.14, yielding that tails of $\mathcal{R}_{\top(S)}$ are necessarily disjoint from tails in $\mathcal{R}_{\perp(R)}$. Their the union of those tails with \mathcal{R} yields a set $\overline{\mathcal{R}}$ as desired. \square

DEFINITION 6.7.22. Let \mathcal{R}, \mathcal{S} be finite families of disjoint ω -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 .

The following remarks are a direct consequence of the definitions and Corollary 6.7.7.

REMARK 6.7.23. *Let $\mathcal{R}, \mathcal{S}, \mathcal{T}$ be finite families of disjoint ω -rays, let $\mathcal{R}' \subseteq \mathcal{R}$ be a subfamily of core rays, and let $\mathcal{P}_1, \mathcal{P}_2$ be a linkages from \mathcal{R} to \mathcal{S} and from $(\mathcal{R} \circ_{\mathcal{P}_1} \mathcal{S})$ to \mathcal{T} respectively.*

- (1) *If \mathcal{P}_1 is transitional and \mathcal{R}' is on the central path of \mathcal{R} , then it is preserving on \mathcal{R}' .*
- (2) *If \mathcal{P}_1 is preserving on \mathcal{R}' , then the sub-linkage of \mathcal{P}_1 from \mathcal{R}' to the respective subfamily of \mathcal{S} is transitional.*
- (3) *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}' .*
- (4) *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 6.7.24. *Let \mathcal{R} and \mathcal{S} be finite families of disjoint core rays of ω , 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 $\mathcal{T} := (\overline{\mathcal{S}} \setminus \mathcal{S}) \cup \mathcal{S}' \subseteq \overline{\mathcal{S}}$. Take a transitional linkage from $\overline{\mathcal{R}}$ to \mathcal{T} . This linkage can be viewed as a linkage from $\overline{\mathcal{R}}$ to $\overline{\mathcal{S}}$, is preserving on \mathcal{R} by Remark 6.7.23(1),

and hence the sub-linkage from \mathcal{R} to \mathcal{S}' is also preserving on \mathcal{R} by Remark 6.7.23(3) as well as transitional by Remark 6.7.21. \square

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

To show that a given graph G is \preceq -ubiquitous, we shall assume that $nG \preceq \Gamma$ 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 is established in [31] and [32].

DEFINITION 6.8.1 (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} is *connected* if every member H of \mathcal{F} is connected. Note that this is the case precisely if G is connected.
- A G -tribe \mathcal{F}' in Γ is a G -subtribe⁵ of a G -tribe \mathcal{F} in Γ , denoted by $\mathcal{F}' \preceq \mathcal{F}$, if there is an injection $\Psi: \mathcal{F}' \rightarrow \mathcal{F}$ such that for each $F' \in \mathcal{F}'$ there is an injection $\varphi_{F'}: F' \rightarrow \Psi(F')$ with $V(H') \subseteq V(\varphi_{F'}(H'))$ for every $H' \in F'$. The G -subtribe \mathcal{F}' is called *flat*, denoted by $\mathcal{F}' \subseteq \mathcal{F}$, if there is such an injection Ψ satisfying $F' \subseteq \Psi(F')$.
- A thick G -tribe \mathcal{F} in Γ is *concentrated at an end* ϵ of Γ , if for every finite vertex set X of Γ , the G -tribe $\mathcal{F}_X = \{F_X: F \in \mathcal{F}\}$ consisting of the layers $F_X = \{H \in F: H \not\subseteq C(X, \epsilon)\} \subseteq F$ is a thin subtribe of \mathcal{F} .

We note that, if G is connected, every thick G -tribe \mathcal{F} contains a thick subtribe \mathcal{F}' such that every $H \in \bigcup \mathcal{F}$ is a tidy IG . We will use the following lemmas from [31].

LEMMA 6.8.2 (Removing a thin subtribe, [31, 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 6.8.3 (Pigeon hole principle for thick G -tribes, [31, 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} .*

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

LEMMA 6.8.4 ([31, 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 6.8.5 ([31, 5.5]). *Let G be a connected graph and Γ a graph containing a thick connected G -tribe \mathcal{F} concentrated at an end ϵ of Γ . Then the following assertions hold:*

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

The following lemma from [32] shows that we can restrict ourself to thick G -tribes which are concentrated at thick ends.

LEMMA 6.8.6 ([32, 6.7]). *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 trees $T_0 \subseteq T_1 \subseteq \dots$ such that $\bigcup T_i = T$ and to construct inductively a family of finitely many $IG[T_{k+1}]$ s which extend the $IG[T_k]$ s built previously (cf. Definition 6.4.6). The extensiveness of the tree-decomposition ensures 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 displays self-similarity.

In order to extend our $IG[T_k]$ at each step, we will want to assume that the IG s in \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 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 6.3.13) which might either be ϵ , or another end of Γ . We would like that our G -tribe \mathcal{F} makes a consistent choice 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 with [32, Lemma 9.1] we may assume that $RG_H(H^\downarrow(\mathcal{R}_e))$ is a path for each H in the G -tribe \mathcal{F} .

By moving to a thick subtribe, we may assume that every ray in every $H \in \bigcup \mathcal{F}$ is core, in which case \leq_ϵ imposes a linear order on every family of rays $H^\downarrow(\mathcal{R}_e)$, which induces one of the two distinct orientations of the path $RG_H(H^\downarrow(\mathcal{R}_e))$ (reference to make this clear/precise). 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 6.8.7. Let G be a connected locally finite graph with a extensive tree-decomposition (T, \mathcal{V}) , S be an initial subtree of T . Let $H \subseteq \Gamma$ be an 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 for every ray $R \in \omega$ we have $H^\downarrow(R) \in \epsilon$. 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 for every ray $R \in \omega$ we have $H^\downarrow(R) \notin \epsilon$. 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 $e \in E$, \mathcal{H} determines whether ω_e converges to ϵ . If \mathcal{H} weakly agrees about $\partial(S)$ we let

$$\partial_\epsilon(S) := \{e \in \partial(S) : \omega_e \text{ converges to } \epsilon \text{ according to } \mathcal{H}\},$$

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

and write

$S^{-\epsilon}$ for the component of the forest $T - \partial_\epsilon(S)$ containing the root of T ,

S^ϵ for the component of the forest $T - \partial_{-\epsilon}(S)$ containing the root of T .

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 $\text{core}(\epsilon)$; and
- for every $e \in \partial_\epsilon(S)$ there is a linear order $\leq_{\mathcal{F},e}$ on $S(e)$ such that the order induced on $H^\downarrow(\mathcal{R}_e)$ by $\leq_{\mathcal{F},e}$ 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.

REMARK 6.8.8. *We note that 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 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 6.8.9 ([32]). Let ω be an end of a graph G . We say ω is *linear* if $RG(\mathcal{R})$ is a path for every finite family \mathcal{R} of disjoint ω -rays.

LEMMA 6.8.10 ([32, 8.1]). *Let ϵ be a non-pebbly end of Γ and let \mathcal{F} be a 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}' such that ω_H is linear for every $H \in \bigcup \mathcal{F}'$.*

COROLLARY 6.8.11. *Let G be a connected locally finite graph with an extensive tree-decomposition (T, \mathcal{V}) , S be 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 S . Then ω_e is linear for every $e \in \partial_\epsilon(S)$.*

PROOF. For any $e \in \partial_\epsilon(S)$ apply Lemma 6.8.10 to \mathcal{F} with $\omega_H = \omega_e$ for each $H \in \bigcup \mathcal{F}$. \square

LEMMA 6.8.12. *Let G be a connected locally-finite graph with a tree-decomposition (T, \mathcal{V}) . Let \mathcal{F} be a thick G -tribe in Γ concentrated at ϵ which weakly agrees about some finite $\partial(S) \subseteq E(T)$. Then \mathcal{F} has a flat thick subtribe \mathcal{F}' so that \mathcal{F}' strongly agrees about $\partial(S)$.*

PROOF. Clear. \square

LEMMA 6.8.13. *Let G be a connected locally-finite graph with an extensive tree-decomposition (T, \mathcal{V}) . Let $H \subseteq \Gamma$ be an IG and ϵ an end of Γ . Let e be an edge of T , such that $H(\omega_e) \neq \epsilon$. There is a finite set $X \subseteq V(G)$ such that for every finite $X' \supseteq X$ there exists a push-out H_e of H along e so that $C_\Gamma(X', G(\omega_e)) \neq C_\Gamma(X', \epsilon)$ and*

- (1) $H_e(G[B(e)]) \subseteq C_\Gamma(X', G(\omega_e))$,
- (2) $H_e(G[B(e)]) \setminus X \subseteq H(G[B(e')])$ for an edge e' on R_e , and
- (3) $H_e(G[A(e)])$ extends $H(G[A(e)])$ fixing $A(e) \setminus S(e)$.

PROOF. Let $X_1 \subseteq V(\Gamma)$ be a finite vertex set such that $C_\Gamma(X, G(\omega_e)) \neq C_\Gamma(X, \epsilon)$, then given any finite $X' \supseteq X_1$, surely $C_\Gamma(X', G(\omega_e)) \neq C_\Gamma(X', \epsilon)$. Since X_1 is finite, there are only finitely many $v \in G$ whose branch sets $H(v)$ meet X_1 . By extensiveness, every vertex of G is contained in only finitely many parts of the tree-decomposition, and so there exists an edge e_1 on R_e with

$$H(G[B(e_1)]) \cap X_1 = \emptyset.$$

For each $s \in S(e)$ let P_s be the initial segment of $R_{e,s}$ up to the first time it meets $S(e_1)$. Let

$$X = X_1 \cup \bigcup_{v \in V(P_s), s \in S(e)} H(v).$$

Then, given any $X' \supseteq X$, as before there is an edge e' on R_e such that

$$H(G[B(e')]) \cap X' = \emptyset.$$

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 $\max\{\text{dist}(e^-, e_1^-), \text{dist}(e^-, e'^-)\} := n$. Then by Definition 6.4.11 and Lemma 6.4.12 there is a push-out H_e of H along e to depth n .

By Definition 6.4.11 $V(H_e(G[B(e)])) \subseteq V(H_e(W)) \cup X$ and hence (1) and (2) hold, and also $H_e([A(e)])$ extends $H(G[A(e)])$ fixing $A(e) \setminus S(e)$. \square

LEMMA 6.8.14. *Let G be a connected locally finite graph with an extensive tree-decomposition (T, \mathcal{V}) with root r . 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 sub-tribe \mathcal{F}' of \mathcal{F} such that*

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

PROOF. Since $d(r)$ is finite, by choosing a thick flat subtribe of \mathcal{F} , we may assume that \mathcal{F} weakly agrees about $\partial(\{r\})$. Moreover, by Lemma 6.8.12, we may even assume that \mathcal{F} strongly agrees about $\partial(\{r\})$.

For every member H of \mathcal{F} , and for every $e \in \partial_{-\epsilon}(\{r\})$ there exists by Lemma 6.8.13 a finite set X_e such that for every finite $X' \supseteq X_e$ there is a push-out H_e of H along e so that $C_\Gamma(X', G(\omega_e)) \neq C_\Gamma(X', \epsilon)$ and

- (1) $H_e(G[B(e)]) \subseteq C_\Gamma(X', G(\omega_e))$,
- (2) $H_e(G[B(e)]) \setminus X_e \subseteq H(G[B(e')])$ for an edge e' on R_e , and
- (3) $H_e(G[A(e)])$ extends $H(G[A(e)])$ fixing $A(e) \setminus S(e)$.

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 6.4.10 and by (3) 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 6.4.11, since these vertices come from $H(\mathcal{R}_e)$ 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, we claim that H strongly agrees with H' about $\partial(\{r\})$ for every member H of \mathcal{F} . Indeed, by construction for each $e \in \partial_{-\epsilon}(\{r\})$, $H'(G[B(e)]) \subseteq C_\Gamma(X', G(\omega_e))$, 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 $\text{core}(\epsilon)$.

Then, since \mathcal{F} is strongly concentrated at ϵ and strongly agrees about $\partial(\{r\})$ it follows that (1) and (2) hold for \mathcal{F}' . It remains to show that \mathcal{F}' is well-separated from ϵ at $\{r\}$.

However, we claim that for each member H of \mathcal{F} the set X defined above separates $H'(\{r\}^{-\epsilon})$ from ϵ in Γ . Indeed,

$$H'(\{r\}^{-\epsilon}) = H'(\{r\}) \cup \bigcup_{e \in \partial_{-\epsilon}(\{r\})} H'(G[B(e)]),$$

and so $H'(\{r\}^{-\epsilon}) \cap C_\Gamma(X, \epsilon) = \emptyset$. It follows that \mathcal{F}' satisfies the conclusion of the lemma. \square

LEMMA 6.8.15 (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 an IG and ϵ an end of Γ . Let S be a finite 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 6.4.11) such that $\{H'\}$ is well-separated from ϵ at $\tilde{S} = S \cup \{f\}$.*

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

By applying Lemma 6.8.13 to every $e \in \partial_{-\epsilon}(\tilde{S}) \setminus \partial(S)$, we obtain a finite set $X \supseteq X'$ and a family $(H_e : e \in \partial_{-\epsilon}(\tilde{S}) \setminus \partial(S))$ where each H_e is a push out of H along e such that

- (1) $H_e(G[B(e)]) \subseteq C_\Gamma(X, H(\omega_e))$,
- (2) $H_e(G[B(e)]) \subseteq H(G[B(e')])$ for some edge e' on R_e , and
- (3) $H_e(G[A(e)])$ extends $H(G[A(e)])$ fixing $A(e) \setminus S(e)$.

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 H is well-separated from ϵ at S there is a finite set X such that $H(S^{-\epsilon}) \cap C_\Gamma(X, \epsilon) = \emptyset$. Let

$$\bar{X} = X \cup \bigcup_{e \in \partial_{-\epsilon}(\tilde{S}) \setminus \partial(S)} V(H_e(S(e))),$$

note that \bar{X} is finite.

It is sufficient to show that \bar{X} separates $H'(G[B(e)])$ from ϵ in Γ for each $e \in \partial_{-\epsilon}(\tilde{S})$, since then \bar{X} together with $H'(S)$ separates $H'(S^{-\epsilon})$ from ϵ in Γ . Given an edge $e \in \partial_{-\epsilon}(\tilde{S})$

either $e \in \partial_{-\epsilon}(S)$ or $e \in \partial_{-\epsilon}(\tilde{S}) \setminus \partial(S)$. In the first case, since

$$H'(G[B(e)]) \subseteq \bigcup_{e' \in \partial_{-\epsilon}(\tilde{S}) \setminus \partial(S)} H_{e'}(G[B(e)]) \subseteq H(G[B(e)]) \cup \bigcup_{e' \in \partial_{-\epsilon}(\tilde{S}) \setminus \partial(S)} H_{e'}(S(e')),$$

by (3), it follows that $H'(G[B(e)]) \cap C_\Gamma(\overline{X}, \epsilon) = \emptyset$.

In the second case $e \in \partial_{-\epsilon}(\tilde{S}) \setminus \partial(S)$, and so again it follows from (3) that

$$H'(G[B(e)]) \subseteq H_e(G[B(e)]) \cup \bigcup_{e \neq e' \in \partial_{-\epsilon}(\tilde{S}) \setminus \partial(S)} H_{e'}(S(e)).$$

Hence, $H'(G[B(e)]) \cap C_\Gamma(\overline{X}, \epsilon) = \emptyset$.

□

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 G -tribe \mathcal{F}_n which agrees about $\partial(T_n)$, which will allow us to extend our $IG[T_n]$ s to $IG[T_{n+1}]$ s. In order to perform the next stage of our construction we will need to ‘refine’ \mathcal{F}_n to a G -tribe \mathcal{F}_{n+1} which agrees about the boundary of T_{n+1} .

This would be a relatively simple application of the pigeon hole principle for G -tribes, Lemma 6.8.3, 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 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 to link.

LEMMA 6.8.16 (*G -tribe refinement lemma*). *Let G be a connected locally finite graph with an extensive tree-decomposition (T, \mathcal{V}) , let S be a 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 $\tilde{S} = S \cup \{f\}$. 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 $F' \in \mathcal{F}'$ such that for all $H \in F$ there is an $H' \in F'$ which is a push-out of H to depth m along f .

PROOF. For every member H of \mathcal{F} consider a sequence $(H^{(i)} : i \in \mathbb{N})$ where $H^{(i)}$ is a push-out of H along f to depth at least i . After choosing a subsequence of $(H^{(i)} : i \in \mathbb{N})$ and relabelling (monotonically), we may assume that for each H , the set $\{H^{(i)} : i \in \mathbb{N}\}$ weakly agrees on $\partial(\tilde{S})$, i.e. for every $e \in \partial(\tilde{S})$ either $H^{(i)}(R) \in \epsilon$ for every $R \in \omega_e$ and all i or $H^{(i)}(R) \notin \epsilon$ 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 6.8.3 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 6.8.11 for every $e \in \partial_\epsilon(\tilde{S})$ (w.r.t. \mathcal{F}_1) the ray graph $RG_G(\mathcal{R}_e)$ is a path. We pick an arbitrary orientation of this path and denote by \leq_e the corresponding linear order on \mathcal{R}_e .

Again for every member $H \in \bigcup \mathcal{F}_1$ define

$$d_H : \{H^{(i)} : i \in \mathbb{N}\} \rightarrow \{-1, 0, 1\}^{\partial_\epsilon(\tilde{S})}$$

where

$$d_H(H^{(i)})_e = \begin{cases} 0 & \text{if } H^{(i)}(\mathcal{R}_e) \text{ are not all core rays,} \\ +1 & \text{if } H^{(i)}(\mathcal{R}_e) \text{ are all core rays and } \leq_e \text{ agrees with } \leq_e, \\ -1 & \text{if } H^{(i)}(\mathcal{R}_e) \text{ are all core rays and } \leq_e \text{ agrees with } \geq_e. \end{cases}$$

Since d_H has finite range we may assume as above, 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 $d : \bigcup \mathcal{F}_1 \rightarrow \{-1, 0, 1\}^{\partial_\epsilon(\tilde{S})}$ with $d(H) = d_H(H^{(1)})$ ($= d_H(H^{(i)})$ for all i). Again, we may choose a thick flat subtribe $\mathcal{F}_2 \subseteq \mathcal{F}_1$ such that d is constant on \mathcal{F}_2 .

Note that no coordinate of d takes the value 0. Indeed, for $e \in \partial_\epsilon(\tilde{S})$ and every layer $F \in \mathcal{F}_2$ the rays in $(H^{(1)}(\mathcal{R}_e) : H \in F)$ are disjoint, and for large enough F it cannot be the case that there is a non-core ray in every $H^{(1)}(\mathcal{R}_e)$.

We can now apply Lemma 6.8.15 to each $H^{(i)}$ to obtain $H'^{(i)}$, the collection of which is well-separated from ϵ at \tilde{S} . Note that $H'^{(i)}$ is still a push-out of H along f to depth 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 6.8.5. 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).

(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) & (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

6.9. The inductive argument

In this section we prove Theorem 6.4.9. Given a connected, locally finite 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 6.8.4 we may assume that \mathcal{F} is concentrated at some end ϵ of Γ , which by Lemma 6.8.6 we may assume is a thick end, and by Lemma 6.6.5 we may assume is not pebbly. Hence, by the work of Section 6.7 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 . However, the two cases are very similar, with the grid-like case being significantly simple. Therefore we will first prove Theorem 6.4.9 in the case where ϵ is half-grid-like, and then in Section 6.9.2 we will briefly sketch the differences for the grid-like case.

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

- \mathcal{F} is concentrated at ϵ ;
- ϵ is thick;
- ϵ is not pebbly;
- ϵ is half-grid-like.

6.9.1. The half-grid-like case. As explained in Section 6.2, our strategy will be to take some sequence of subtrees $S_1 \subseteq S_2 \subseteq S_3 \dots$ of T , such that $\bigcup_i S_i = T$, 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 $T_n^{-\epsilon}$ with respect to some tribe \mathcal{F}_n which 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 $S_n = T_n^{-\epsilon}$ 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 6.8.12 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* which separates all 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’ve built so far.

Furthermore, in order to extend these copies in the next set 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 6.9.1 (Bounder, extender). Let \mathcal{F} be a thick G -tribe which is concentrated at ϵ and strongly agrees about $\partial(S)$ for some 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 $\text{init}(E_{e,s,i}) \in Q_i(s)$.
- Given an extender \mathcal{E} , an edge $e \in \partial_\epsilon(S)$ and $i \in [k]$ we let

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

Recall that, since ϵ is half-grid like, there is a partial order \leq_ϵ defined on the core rays of ϵ , see Lemma 6.7.16. Furthermore, if \mathcal{F} strongly agrees about $\partial(S)$ then, as in Definition 6.8.7, for each $e \in \partial_\epsilon(S)$ there is a linear order $\leq_{\mathcal{F},e}$ on $S(e)$.

DEFINITION 6.9.2 (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 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 $n \in \mathbb{N}$, with $r := \lfloor n/2 \rfloor$ and $s := \lceil n/2 \rceil$ we have:

- (I1) a thick G -tribe \mathcal{F}_r in Γ which
- is concentrated at ϵ ;
 - strongly agrees about $\partial(T_r)$;
 - is well-separated from ϵ at T_r ; and
 - whenever $l < k \leq r$, $T_k^{-\epsilon}$ with respect to \mathcal{F}_k is a sub-tree of $T_l^{-\epsilon}$ with respect to \mathcal{F}_l .
- (I2) a family $\mathcal{Q}_n = (Q_i^n : i \in [s])$ of s pairwise disjoint inflated copies of $G(T_r^{-\epsilon})$ (where $T_r^{-\epsilon}$ is considered with respect to \mathcal{F}_r) in Γ ;
- if $n \geq 1$, we additionally require that Q_i^n extends Q_i^{n-1} for all $i \leq s - 1$;
- (I3) an extension scheme (X_n, \mathcal{E}_n) for \mathcal{Q}_n ;
- (I4) if n is even and $\partial_\epsilon(T_r) \neq \emptyset$, we require that there is a set \mathcal{J}_r of disjoint core ϵ -rays disjoint to \mathcal{E}_n with $|\mathcal{J}_r| \geq (|\partial_\epsilon(T_r)| + 1) \cdot |\mathcal{E}_n|$.

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 sub-tree of $T_l^{-\epsilon}$ with respect to \mathcal{F}_l for all $k < l$, we have $\bigcup_{n \in \mathbb{N}} T_n^{-\epsilon} = T$ (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 for the case $n = 0$ we first note that since $T_0 = t_0$, by Lemma 6.8.14 there is a thick subtribe \mathcal{F}_0 of \mathcal{F} which satisfies (I1). Let us further take

- $\mathcal{Q}_0 = \mathcal{E}_0 = X_0 = \emptyset$;
- \mathcal{J}_0 be any suitably large set of disjoint core rays of ϵ .

The following notation will be useful throughout the construction. Given $e \in E(T)$ and some inflated copy H of G , recall that $H^\downarrow(\mathcal{R}_e)$ denotes the family $(H^\downarrow(R_{e,s}) : s \in S(e))$. Given a G -tribe \mathcal{F} , a layer $F \in \mathcal{F}$ and a family of 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 $Q_{k+1}^{n+1} = H(T_k^{-\epsilon})$ gives us a further inflated copy of $G(T_k^{-\epsilon})$ disjoint from all the previous ones. We set $Q_i^{n+1} = Q_i^n$ for all $i \in [k]$ and $\mathcal{Q}_{n+1} = (Q_i^{n+1} : i \in [k+1])$. Using that \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 6.9.1)

Consider the family $\mathcal{R}^- := \bigcup \{\mathcal{R}_e^- : e \in \partial_\epsilon(T_k)\}$. Moreover, set $\mathcal{C} := \mathcal{E}_n^- \cup \mathcal{J}_k$ and consider $\bar{\mathcal{C}}$ as in Definition 6.7.19. Let $Y \subseteq C(X_n, \epsilon)$ be a finite set which is a transition box between $\bar{\mathcal{E}}_n^-$ and $\bar{\mathcal{C}}$ as in Lemma 6.3.17. Let \mathcal{F}' be a flat thick G -subtribe of \mathcal{F}_k such that each member of \mathcal{F}' is contained in $C(X_n \cup Y, \epsilon)$, which exists by Lemma 6.8.5 since both X_n and Y are finite.

Let R be an arbitrary element of \mathcal{R} . Let $F \in \mathcal{F}'$ be large enough such that we may apply Lemma 6.3.16 to find a transitional linkage $\mathcal{P} \subseteq C(X_n \cup Y, \epsilon)$ from $\bar{\mathcal{C}}$ to $F^\downarrow(\mathcal{R}^-)$ after $X_n \cup Y$ avoiding some member $H \in F$. Note that, since X_n is a bounder and $\mathcal{P} \subseteq C(X_n \cup Y, \epsilon)$, \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 $X_{n+1} \subseteq \Gamma$ containing $X_n \cup Y$ 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, Remark 6.7.23(1) implies that the linkage is preserving on \mathcal{C} . Since all rays in $F^\downarrow(\mathcal{R}^-)$ are core rays, \leq 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) 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 such an interval on $F^\downarrow(\mathcal{R}^-)$ that

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

By Lemma 6.7.24 and Remark 6.7.23(1 and 3), and Lemma 6.3.17 there is a linkage \mathcal{P}'' from \mathcal{E}_n^- to \mathcal{A} contained in $\Gamma[Y]$ which is preserving on \mathcal{E}_n^- .

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

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

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

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

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

We claim that $(X_{n+1}, \mathcal{E}_{n+1})$ is an extension scheme for \mathcal{Q}_{n+1} and hence property (I3) is satisfied. Since every ray in \mathcal{E}_{n+1} has a tail which is also a tail of a ray in $F^\downarrow(\mathcal{R}^-)$, property (ES2) is satisfied by Remark 6.7.171. Since \mathcal{P}' is preserving on \mathcal{A}' and \mathcal{P}'' is preserving on \mathcal{E}_n^- , Remark 6.7.23(4) 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 $s \in S(e)$, 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 $\mathcal{P}'' + \mathcal{P}'$ is preserving on \mathcal{E}_n^- 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 6.3. 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 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 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 6.9.1)

By (I1) we can apply Lemma 6.8.16 to \mathcal{F}_{k-1} in order to find a thick G -tribe \mathcal{F}_k and a thick flat sub-tribe \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 6.7.24, there is, for every subset \mathcal{S} of $F^\downarrow(\mathcal{R}_f)$ of size $|\mathcal{E}_f^-|$ a transitional linkage $\mathcal{P} \subseteq C(X_n, \epsilon)$ from $\overline{\mathcal{E}_n^-}$ to $\overline{F^\downarrow(\mathcal{R}_f)}$ after $X_n \cup \text{init}(\mathcal{E}_n)$ such that \mathcal{P} links \mathcal{E}_f to \mathcal{S} , if we view it as a linkage from $\overline{\mathcal{E}_n^-}$ to $\overline{F^\downarrow(\mathcal{R}_f)}$. Since all rays in \mathcal{E}_f and in $F^\downarrow(\mathcal{R}_f)$ are core rays, any such linkage is preserving on \mathcal{E}_f .

Let us choose $H_1, H_2, \dots, H_k \in F$ and let $\mathcal{S} = \left(H_i^\downarrow(R_{f,s}) : i \in [k], s \in S(f) \right)$. 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 $H^\downarrow(\mathcal{R}_f)$ forms an interval in $F^\downarrow(\mathcal{R}_f)$ it follows that, after perhaps relabelling the H_i , for every $i \in [k]$ and $s \in S(f)$, \mathcal{P} links $E_{f,s,i}^n$ to $H_i^\downarrow(R_{f,s})$.

Let $Z \subseteq V(\Gamma)$ be a finite set such that $\top(\omega, R)$ and $\perp(\omega, R)$ are separated by Z in $\Gamma - V(R)$ for all $R \in F^\downarrow(\mathcal{R}_f)$ (cf. Lemma 6.7.18).

Since $|F|$ is finite and (T, \mathcal{V}) is an extensive tree-decomposition there exists an $m \in \mathbb{N}$ such that if $e \in R_f$ with $\text{dist}(f^-, e^-) = m$ then $H(B(e)) \cap (X_n \cup Z \cup V(\bigcup \mathcal{P})) = \emptyset$. Let $\vec{F} \in \mathcal{F}_k$ be as in Lemma 6.8.16(vi) for F with such an m .

Hence, by definition, for each $H_i \in F$ there is some subgraph $W_i \subseteq H(B(e))$ which is an $\overline{IG[B(f)]}$ such that for each $s \in S(f)$, $W_i(s)$ contains the first point 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

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

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 point 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

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

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 6.4.10) 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 \vec{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

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

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

We claim that $(X_{n+1}, \mathcal{E}_{n+1})$ is an extension scheme. Property (ES1) is apparent. Since the G -tribes \mathcal{F}_k and \mathcal{F}^* both strongly agree about $\partial(T_k)$, 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 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 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 Remark 6.7.23(1) 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 $(\vec{H}_i^\downarrow(R_{e,s}) : s \in S(e))$. Since $\vec{H}_i \in \vec{F} \in \mathcal{F}_k$ 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 6.3.20 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 \vec{H}_i to which the tails of the rest of \mathcal{E}_{n+1}^- are disjoint. Furthermore, by choice of Z and Lemma 6.7.18 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 $H \in F \in \mathcal{F}^*$. Since there is at least one core ϵ -ray in each $H \in F \in \mathcal{F}^*$, we can find family of at least $|F| - |\mathcal{E}_{n+1}|$ such rays. 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. □

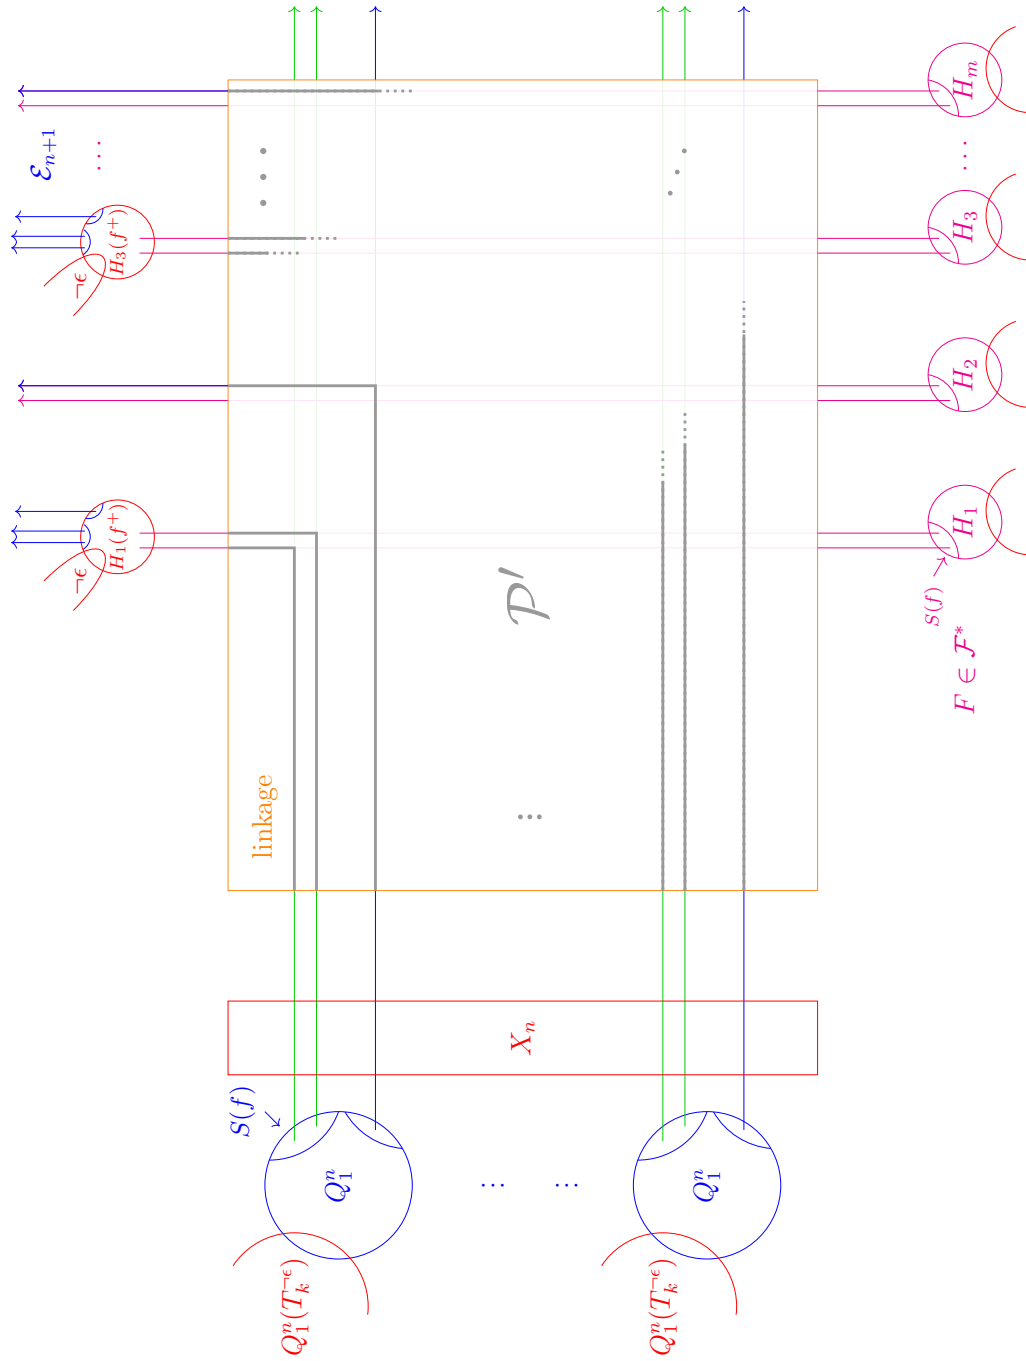


FIGURE 6.4. Extending the copies when $n = 2k - 1$ is odd.

6.9.2. 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 6.8.11, $RG_G(\mathcal{R}_e)$ 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 $G(e^+)$ 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 is such that for each $i \in [k]$ there is some $H \in F$ such that ray $E_{e,s,i}$ is linked to $H^\downarrow(R_{e,s})$ for each $s \in S(e)$. However, since any transitional linkage we construct between \mathcal{E} 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 $RG_G(\mathcal{R}_e)$ 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 a number such that no family of disjoint ϵ -rays has a ray graph which is N -pebble win. 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 $i \in [n]$ ⁶.

Every large enough ray graph is a cycle, which has a correct orientation by Lemma 6.7.2 and we would like to say that this orientation is induced by a global ‘partial cyclic order’ defined on the core rays of ϵ .

By a similar argument as in Section 6.7.3 one can show the following:

⁶We note that it is possible to show that, if ϵ is grid-like, then in fact $N = 3$.

LEMMA 6.9.3. *Let R and R' be disjoint core rays of ϵ . Then in $G - (V(R) \cup V(R'))$ the end ϵ splits into precisely two different ends.*

DEFINITION 6.9.4. Let R and R' be a core ray 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 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.

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 $C \subseteq X^3$ 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 6.9.5. *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 disjoint family of at least $N + 3$ ϵ -rays $(R_i: i \in [n])$ the cyclic order induced on $(R_i: i \in [n])$ by C_ϵ agrees with the correct orientation.

Again by a similar argument as in Section 6.7.3 one can show that this relation is in fact a partial cyclic order and that it always agrees with the correct orientation of large enough ray graphs. Furthermore, by Lemma 6.7.3, 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 disjoint family of ω -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 $R_i < R_j < R_k$.

Recall that, 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 $i \leq j$.

We say \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 $\text{core}(\epsilon)$; and
- for every $e \in \partial_\epsilon(T_n)$ there is a linear order $\leq_{\mathcal{F}, e}$ on $S(e)$ such that $H^\downarrow(\mathcal{R}_e)(\leq_{\mathcal{F}, e})$ agrees with C_ϵ on $H^\downarrow(\mathcal{R}_e)$ for all $H \in \bigcup \mathcal{F}$.

Using this definition the G -tribe refinement lemma (Lemma 6.8.16) can also be shown to hold in the case where ω 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 6.9.6 (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 $i \in [k]$.

The 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 don't 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 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 $H_1, \dots, H_N \in F$ and to take

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

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.

6.10. 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's recall how well-quasi-ordering was used in the proofs of Lemmas 6.5.6 and 6.5.12. Lemma 6.5.6 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 6.5.12, 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 [125] 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 6.1.1 and 6.1.2:

CONJECTURE 6.10.1. Any locally finite connected graph in which all blocks are finite is ubiquitous.

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

Thomas has conjectured 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, something which we hope to do in the next paper in this series [\[34\]](#).

Minimal obstructions for normal spanning trees

Diestel and Leader have characterised connected graphs that admit a normal spanning tree via two classes of forbidden minors. One class are Halin's (\aleph_0, \aleph_1) -graphs: bipartite graphs with bipartition (A, B) such that $|A| = \aleph_0$, $|B| = \aleph_1$ and every vertex of B has infinite degree.

Our main result is that under Martin's Axiom and the failure of the Continuum Hypothesis, the class of forbidden (\aleph_0, \aleph_1) -graphs in Diestel and Leader's result can be replaced by one single instance of such a graph.

Under CH, however, the class of (\aleph_0, \aleph_1) -graphs contains minor-incomparable elements, namely graphs of binary type, and \mathcal{U} -indivisible graphs. Assuming CH, Diestel and Leader asked whether every (\aleph_0, \aleph_1) -graph has an (\aleph_0, \aleph_1) -minor that is either indivisible or of binary type, and whether any two \mathcal{U} -indivisible graphs are necessarily minors of each other. For both questions, we construct examples showing that the answer is in the negative.

7.1. Introduction

A (*graph theoretic*) *tree* is a connected, acyclic graph. A subgraph H of a graph G is called *spanning* if H has the same vertex set as G . Thus, a *spanning tree* T of a connected graph G is a connected, acyclic subgraph containing every vertex of G . A tree is *rooted* if it has one designated vertex, called the *root*. Fixing a root of a graph-theoretic tree T induces a natural tree order on its vertex set $V(T)$ with the root as unique minimal element.

A rooted spanning tree T of a graph G is called *normal* if the end-vertices of any edge of G are comparable in the natural tree order of T , see e.g [54, §1.5]. Intuitively, all the edges of G run 'parallel' to branches of T , but never 'across'. Every countable connected graph has a normal spanning tree, but uncountable graphs might not, as demonstrated by complete graphs on uncountably many vertices [54, 8.2.3].

Halin [89, 7.2] observed that as a consequence of a theorem of Jung, the property of having a normal spanning tree is minor-closed, i.e. preserved under taking (connected) minors. Here, a graph H is a *minor* of another graph G , written $H \preceq G$, if to every vertex $x \in H$ we can assign a (possibly infinite) connected set $V_x \subseteq V(G)$, called the *branch set* of x , so that these sets V_x are disjoint for different x and G contains a $V_x - V_y$ edge whenever xy is an edge of H .

Halin's observation opens up the possibility of a forbidden minor characterisation for the property of admitting normal spanning trees. In the universe of finite graphs, the famous Seymour-Robertson Theorem asserts that any minor-closed property of finite graphs can be characterised by *finitely* many forbidden minors, see e.g. [54, §12.7]. Whilst for infinite graphs, we generally need an infinite list of forbidden minors, Diestel and Leader have shown that for the property of having a normal spanning tree, the forbidden minors come in two structural types.

Following Halin, a bipartite graph with bipartition (A, B) is called an (\aleph_0, \aleph_1) -graph if $|A| = \aleph_0$, $|B| = \aleph_1$, and every vertex in B has infinite degree.

THEOREM (Diestel and Leader, [58]). *A connected graph admits a normal spanning tree if and only if it does not contain an (\aleph_0, \aleph_1) -graph or an AT-graph (a certain kind of graph whose vertex set is an order-theoretic Aronszajn tree) as a minor.*

In the same paper, they ask how one might further describe the minor-minimal graphs within the class of (\aleph_0, \aleph_1) -graphs.

One family of possibly minimal (\aleph_0, \aleph_1) -graphs suggested by Diestel and Leader are the *binary trees with tops*, also called (\aleph_0, \aleph_1) -graphs of *binary type*: Let A be a binary tree of countable height, and let B index \aleph_1 -many branches of A . We form an (\aleph_0, \aleph_1) -graph with bipartition (A, B) by connecting every vertex $b \in B$ to infinitely many points on its branch. Details on these graphs can be found in Section 7.2. We can now state our main result as follows.

THEOREM 7.1.1. *Let T be an arbitrary binary tree with tops. Under Martin's Axiom and the failure of the Continuum Hypothesis, the graph T embeds into any other (\aleph_0, \aleph_1) -graph as a subgraph.*

Answering a question by Diestel and Leader, it follows that it is consistent with the usual axioms of set theory ZFC that there is a minor-minimal graph without a normal spanning tree. As a second consequence, we can extend Diestel and Leader's result as follows.

THEOREM 7.1.2. *Let T be an arbitrary binary tree with tops. Under Martin's Axiom and the failure of the Continuum Hypothesis, a graph has a normal spanning tree if and only if it does not contain T , or an AT-graph as a minor.*

However, under the Continuum Hypothesis (CH) the situation is different. Now, there exist *indivisible* (\aleph_0, \aleph_1) -graphs, i.e. graphs (\mathbb{N}, B) where for every partition $\mathbb{N} = A_1 \dot{\cup} A_2$, only one of the induced graphs (A_1, B) and (A_2, B) contains an (\aleph_0, \aleph_1) -subgraph. Note that for every indivisible graph (\mathbb{N}, B) there is a corresponding (non-principal) ultrafilter \mathcal{U} consisting of all subsets $A \subseteq \mathbb{N}$ such that (A, B) contains an (\aleph_0, \aleph_1) -subgraph. Indivisible graphs with associated ultrafilter \mathcal{U} are also called *\mathcal{U} -indivisible*.

In [58, 8.1], Diestel and Leader proved that binary trees with tops and indivisible graphs form two minor-incomparable classes of (\aleph_0, \aleph_1) -graphs. Further, they mention the following two problems involving indivisible graphs:

QUESTION 1 (Diestel and Leader). Assuming CH, does every (\aleph_0, \aleph_1) -graph have an (\aleph_0, \aleph_1) -minor that is either indivisible or of binary type?

QUESTION 2 (Diestel and Leader). Assuming CH, are any two \mathcal{U} -indivisible (\aleph_0, \aleph_1) -graphs necessarily minors of each other?

One particular property of (\aleph_0, \aleph_1) -graphs of binary type is that they are *almost disjoint* (AD): neighbourhoods of any two distinct B -vertices intersect only finitely (see Section 7.2 for further details). Of course, not every (\aleph_0, \aleph_1) -graph has this property, as complete bipartite graphs show. However, our first result in this paper is that we can always restrict our attention to almost disjoint (\aleph_0, \aleph_1) -graphs: In Theorem 7.3.3 below, we show that every (\aleph_0, \aleph_1) -graph has an AD- (\aleph_0, \aleph_1) -subgraph.

Once we have made this reduction, we turn towards Questions 1 and 2. In Theorem 7.5.1, we show that Question 1 has a negative answer. Our construction refines a strategy developed by Roitman and Soukup for the combinatorial analysis of almost disjoint families. We then construct in Theorem 7.6.2 two \mathcal{U} -indivisible graphs that are not minor-equivalent, answering Question 2 in the negative.

7.2. Collections of infinite subsets of \mathbb{N} , and (\aleph_0, \aleph_1) -graphs

The following connection between collections of infinite subsets of \mathbb{N} and (\aleph_0, \aleph_1) -graphs will be used frequently in this paper. Let G be an (\aleph_0, \aleph_1) -graph with bipartition (A, B) , and enumeration $B = \{b_\alpha : \alpha < \omega_1\}$. Identifying A with the integers \mathbb{N} , we can encode G as (multi-)set $\langle N(b_\alpha) : \alpha < \omega_1 \rangle$ of infinite subsets of \mathbb{N} . Conversely, given any multiset $\langle N_\alpha : \alpha < \omega_1 \rangle$ of infinite subsets of \mathbb{N} , we can form an (\aleph_0, \aleph_1) -graph with bipartition (\mathbb{N}, B) by setting $N(b_\alpha) := N_\alpha$.

This correspondence allows us to translate graph-theoretic problems about (\aleph_0, \aleph_1) -graphs to the realm of infinite combinatorics. Let A and B be subsets of \mathbb{N} . If $A \setminus B$ is finite, we say that A is *almost contained* in B , or A is contained in B *mod finite*, and write $A \subseteq^* B$. Consequently, A and B are *almost equal*, $A =^* B$, if $A \subseteq^* B$ and $B \subseteq^* A$ (which means their symmetric difference is finite).

Given any collection \mathcal{P} of infinite subsets of \mathbb{N} , we say that an infinite set $A \subseteq \mathbb{N}$ is a *pseudo-intersection* for \mathcal{P} if $A \subseteq^* P$ for all $P \in \mathcal{P}$. Every countable \mathcal{P} that is directed by \subseteq^* has a pseudo-intersection.

A collection \mathcal{A} of infinite subsets of \mathbb{N} is an *almost disjoint family* (AD-family) if $A \cap A' =^* \emptyset$ for all A, A' in \mathcal{A} (in other words, if the pairwise intersection of elements of \mathcal{A} is always finite). By a diagonalisation argument, every infinite AD-family can be extended to an uncountable AD-family.

The simplest example of an (\aleph_0, \aleph_1) -graph is the complete bipartite graph K_{\aleph_0, \aleph_1} . Binary trees with tops as introduced above are strictly smaller (with respect to the minor relation \preceq) examples of (\aleph_0, \aleph_1) -graphs, as they have the property that $|N(b) \cap N(b')| < \infty$ for all $b \neq b' \in B$. Changing our perspective, we see that in this case, the collection $\langle N(b_\alpha) : \alpha < \omega_1 \rangle$ forms an almost disjoint family on \mathbb{N} . Let us call any (\aleph_0, \aleph_1) -graph with this last property an *almost disjoint (\aleph_0, \aleph_1) -graph*, or for short an AD- (\aleph_0, \aleph_1) -graph.

A *tree* $\mathcal{T} = (T, <)$ in the order-theoretic sense is a partially ordered set T with a smallest element such that all predecessor sets $t^\downarrow = \{s \in T : s < t\}$ are well-ordered by $<$. The order type of t^\downarrow is called the *height* of t , and denoted by $\text{ht}(t)$. The set of all elements of \mathcal{T} of height α is denoted by $\mathcal{T}(\alpha)$, and called the α^{th} level of \mathcal{T} . A subset $S \subseteq T$ of a tree $\mathcal{T} = (T, <)$ is an *initial subtree* if $t^\downarrow \subseteq S$ for all $t \in S$. By $\mathcal{T}(\leq \alpha) = \bigcup_{\beta \leq \alpha} \mathcal{T}(\beta)$ we mean the initial subtree of \mathcal{T} consisting of all elements of \mathcal{T} of height at most α .

A linearly ordered subset of \mathcal{T} is also called a *chain*. A *branch* of a tree \mathcal{T} is an inclusion-maximal chain. The collection of branches is also denoted by $\mathcal{B}(\mathcal{T})$. For b a branch and α an ordinal, $b \upharpoonright \alpha$ denotes the unique element of $b \cap \mathcal{T}(\alpha)$. An *Aronszajn tree* is an uncountable tree such that all levels and all branches are countable. The *binary tree of countable height* is the tree $2^{<\omega}$, the set of all finite binary sequences, ordered by extension. Similarly, a *binary tree of finite height* is a tree isomorphic to $2^{<\omega}(\leq n)$ for some $n \in \mathbb{N}$.

In the following, we list some special types of (\aleph_0, \aleph_1) -graphs (suggested by Diestel and Leader [58]), and some well-known types of almost disjoint families (studied by Roitman and Soukup [138]), all of which will play a role in this paper.

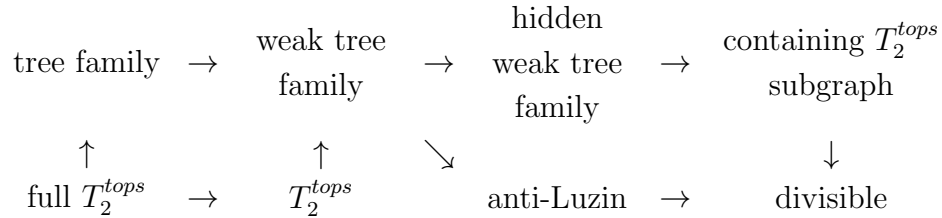
Graph-theoretic perspective (Diestel & Leader).

- T_2^{tops} : Let $A = 2^{<\omega}$ be a binary tree of height ω , and B be a set of \aleph_1 -many branches of A . Any graph isomorphic to some (\aleph_0, \aleph_1) -graph formed on the vertex set $A \dot{\cup} B$ by connecting every vertex $b \in B$ to infinitely many points on its branch is called a T_2^{tops} , or an (\aleph_0, \aleph_1) -graph of binary type.
- *full T_2^{tops}* : As above, but now connect every vertex $b \in B$ to all points on its branch.
- *divisible*: An (\aleph_0, \aleph_1) -graph with bipartition (A, B) is divisible if there are partitions $A = A_1 \dot{\cup} A_2$ and $B = B_1 \dot{\cup} B_2$ such that both (A_1, B_1) and (A_2, B_2) contain (\aleph_0, \aleph_1) -subgraphs.
- *\mathcal{U} -indivisible*: For a non-principal ultrafilter \mathcal{U} , an (\aleph_0, \aleph_1) -graph with bipartition (\mathbb{N}, B) is called \mathcal{U} -indivisible if for all $A \in \mathcal{U}$ we have $N(b) \subseteq^* A$ for all but countably many $b \in B$.

Set-theoretic perspective (Roitman & Soukup).

- *tree-family*: An uncountable AD-family \mathcal{A} on \mathbb{N} is a tree-family if there is a tree-ordering \mathcal{T} of countable height on \mathbb{N} so that for every $A \in \mathcal{A}$ there is a branch of \mathcal{T} which almost equals A .
- *weak tree-family*: As above, but now it is only required that there is an injective assignment from \mathcal{A} to branches of \mathcal{T} such that every $A \in \mathcal{A}$ is almost contained in its assigned branch.
- *hidden (weak) tree-family*: \mathcal{A} is a hidden (weak) tree family if for some countable tree T , $\{T \cap a : a \in \mathcal{A}\}$ a (weak) tree family.
- *anti-Luzin*: An AD-family \mathcal{A} is anti-Luzin if for all uncountable $\mathcal{B} \subseteq \mathcal{A}$ there are uncountable $\mathcal{C}, \mathcal{D} \subseteq \mathcal{B}$ such that $\bigcup \mathcal{C} \cap \bigcup \mathcal{D}$ is finite.

Comparing the different notions. There are striking similarities between the graph-theoretic and the set-theoretic perspective. We gather dependencies between the above concepts in the following diagram. All these implications are straightforward from the definitions.



A little less straightforward is the fact that none of the arrows in the above diagram can generally be reversed. This is witnessed by the following examples.

OBSERVATION 7.2.1. *Under CH, there is a binary tree with tops which is not a tree family.*

CONSTRUCTION SKETCH. Consider a binary tree order \mathcal{T} on \mathbb{N} and, using CH, enumerate its branches $\mathcal{B}(\mathcal{T}) = \{b_\alpha : \alpha < \omega_1\}$. In order to diagonalize against all possible tree families, enumerate all tree orders of countable height on \mathbb{N} as $\{\mathcal{T}_\alpha : \alpha < \omega_1\}$. Now if $|b_\alpha \cap b| = \infty$ for some branch b of \mathcal{T}_α , then choose $N_\alpha \subseteq b_\alpha$ such that $N_\alpha \subsetneq^* b$. Otherwise, put $N_\alpha = b_\alpha$. Then $(N_\alpha : \alpha < \omega_1)$ is as desired. \square

Hence, the implications ‘tree family \rightarrow weak tree family’ and ‘full $T_2^{\text{tops}} \rightarrow T_2^{\text{tops}}$ ’ cannot be reversed.

Next, if in a full T_2^{tops} one additionally makes all tops adjacent to one special node of the tree, one obtains a tree family which cannot be a T_2^{tops} , because in a T_2^{tops} without isolated points on the countable side, only the root of the tree can be simultaneously adjacent to all tops. In particular, the implications ‘ $T_2^{\text{tops}} \rightarrow$ weak tree family’ and ‘full $T_2^{\text{tops}} \rightarrow$ tree family’ cannot be reversed.

Hidden weak tree families need not be anti-Luzin, see [138, p.58]. In particular, the implications ‘weak tree family \rightarrow hidden weak tree family’ and ‘anti-Luzin \rightarrow divisible’ cannot be reversed. In Theorem 7.5.1 below, we construct under CH an anti-Luzin family which contains no T_2^{tops} subgraph, so the implications ‘weak tree family \rightarrow anti-Luzin’ and ‘containing T_2^{tops} subgraph \rightarrow divisible’ cannot be reversed. Finally, the implication ‘hidden weak tree family \rightarrow containing a T_2^{tops} subgraph’ cannot be reversed:

OBSERVATION 7.2.2. *Under CH, there is an AD-family $(N_\alpha: \alpha < \omega_1)$ containing a T_2^{tops} subgraph but which is not a hidden weak tree family.*

CONSTRUCTION SKETCH. Consider a binary tree order \mathcal{T} on \mathbb{N} and enumerate its branches $\mathcal{B}(\mathcal{T}) = \{b_\alpha: \alpha < \omega_1\}$. Enumerate all tree orders of countable height with groundset some infinite subset of \mathbb{N} as $\{\mathcal{T}_\alpha: \alpha < \omega_1\}$. Every N_α will be the union of at most two $b_{\beta_1(\alpha)}$ and $b_{\beta_2(\alpha)}$. At step $\alpha < \omega_1$, we have $\beta = \sup \{b_{\beta_1(\gamma)}, b_{\beta_2(\gamma)}: \gamma < \alpha\} < \omega_1$. If there is b_δ with $\delta > \beta$ such that b_δ is not almost contained in a single branch of \mathcal{T}_α , put $N_\alpha = b_\delta$. If all b_δ with $\delta > \beta$ are almost contained in the same branch of \mathcal{T}_α , put $N_\alpha = b_{\beta+1}$. Otherwise, there are $\beta_1(\alpha) > \beta$ and $\beta_2(\alpha) > \beta$ such that $b_{\beta_1(\alpha)}$ and $b_{\beta_2(\alpha)}$ are almost contained in different branches of \mathcal{T}_α . Put $N_\alpha = b_{\beta_1(\alpha)} \cup b_{\beta_2(\alpha)}$. Then it is easily checked that $(N_\alpha: \alpha < \omega_1)$ is as desired. \square

However, under $\text{MA} + \neg\text{CH}$, every $< \mathfrak{c}$ -sized AD family is a hidden weak tree family [138, 4.4], so the last construction cannot be done in ZFC alone.

7.3. Finding almost disjoint (\aleph_0, \aleph_1) -subgraphs

Almost disjoint (\aleph_0, \aleph_1) -graphs are natural candidates for smaller obstruction sets in Diestel and Leader’s result. In this section, we prove that indeed, every (\aleph_0, \aleph_1) -graph contains an almost disjoint (\aleph_0, \aleph_1) -subgraph.

We say that a collection \mathcal{F} of infinite subsets of some countably infinite set has an *almost disjoint refinement* if there is a choice of infinite subsets $A_F \subseteq F$ such that $\mathcal{A} = \{A_F: F \in \mathcal{F}\}$ is an almost disjoint family.

THEOREM 7.3.1 (Baumgartner, Hajnal and Mate; Hechler). *Every $< \mathfrak{c}$ -sized collection of infinite subsets of \mathbb{N} has an almost disjoint refinement.*

The theorem is due to Baumgartner, Hajnal and Mate [18, 2.1], and independently due to Hechler [94, 2.1]. For convenience, we will indicate the proof below.

COROLLARY 7.3.2. *Assume $\neg\text{CH}$. Every (\aleph_0, \aleph_1) -graph has a spanning AD- (\aleph_0, \aleph_1) -subgraph.*

PROOF. An almost disjoint refinement corresponds, in the graph-theoretic perspective, to a subgraph obtained by deleting, at every vertex on the B -side, co-infinitely many

incident edges. Since we did not remove any vertices, we obtain indeed a spanning AD- (\aleph_0, \aleph_1) -subgraph. \square

Theorem 7.3.1 does not hold for families of size \mathfrak{c} (consider the collection of all infinite subsets of \mathbb{N}). Still, we can prove that the corresponding result for subgraphs is true nonetheless (but we can no longer guarantee spanning subgraphs).

THEOREM 7.3.3. *Every (\aleph_0, \aleph_1) -graph has an AD- (\aleph_0, \aleph_1) -subgraph.*

First, a piece of notation. Let \mathcal{F} be a collection of infinite subsets of \mathbb{N} , and \mathcal{A} be an almost disjoint family. Following Hechler, [94], we say that \mathcal{A} *covers* \mathcal{F} if for every $F \in \mathcal{F}$, the collection $\{A \in \mathcal{A} : |F \cap A| = \infty\}$ is of size $|\mathcal{A}|$.

Hechler showed that a collection \mathcal{F} of infinite subsets of \mathbb{N} has an almost disjoint refinement if and only if there is an almost disjoint family of size $|\mathcal{F}|$ covering \mathcal{F} [94, 2.3]. We shall only make use of the backwards implication, the proof of which is nicely illustrated in the claim below.

PROOF OF THEOREM 7.3.3. Suppose we are given an (\aleph_0, \aleph_1) -graph G with bipartition (\mathbb{N}, B) , an enumeration $B = \{b_\alpha : \alpha < \omega_1\}$ and neighbourhoods $N_\alpha = N(b_\alpha)$.

CLAIM. If $\{N_\alpha : \alpha < \omega_1\}$ forms an uncountable decreasing chain mod finite (i.e. $N_\beta \subseteq^* N_\alpha$ for all $\alpha < \beta$), then G has an AD- (\aleph_0, \aleph_1) -subgraph.

For the claim, consider two alternatives. Either, $\mathcal{N} = \{N_\alpha : \alpha < \omega_1\}$ has an infinite pseudo-intersection A , in which case any uncountable AD-family $\mathcal{A} = \{A_\alpha : \alpha < \omega_1\}$ on A covers $\{N_\alpha : \alpha < \omega_1\}$. Picking $N'_\alpha = N_\alpha \cap A_\alpha$ readily provides an almost disjoint refinement of \mathcal{N} . And if \mathcal{N} does not have an infinite pseudo-intersection, then moving to a subgraph, we may assume that $C_\alpha = N_\alpha \setminus N_{\alpha+1}$ is infinite for all $\alpha < \omega_1$. Now if $\alpha < \beta$ then $C_\alpha \cap C_\beta \subseteq N_\alpha \setminus N_{\alpha+1} \cap N_\beta$ is finite, as $N_\beta \setminus N_{\alpha+1}$ is finite by assumption. So $\{C_\alpha : \alpha < \omega_1\}$ gives rise to an AD- (\aleph_0, \aleph_1) -subgraph of G , establishing the claim.

Now suppose there exists an infinite set $A \subseteq \mathbb{N}$ with the property that for every infinite $C \subseteq A$ there is an uncountable set $K_C = \{\beta < \omega_1 : |N_\beta \cap C| = \infty\}$. Let us construct, by recursion,

- (1) a faithfully indexed set $\{N_{\mu_\alpha} : \alpha < \omega_1\} \subseteq \mathcal{N}$, and
- (2) infinite subsets $C_\alpha \subseteq N_{\mu_\alpha} \cap A$ such that $C_\alpha \subseteq^* C_\beta$ for all $\alpha > \beta$.

First, let $\mu_0 = \min K_A$ and put $C_0 = A \cap N_{\mu_0}$, an infinite subset of A . Next, let $\alpha < \omega_1$ and suppose μ_β and C_β have been defined according to (1) and (2) for all $\beta < \alpha$. Let \tilde{C}_α be an infinite pseudo-intersection of the countable collection $\{C_\beta : \beta < \alpha\}$. We may assume that $\tilde{C}_\alpha \subseteq A$ and let $\mu_\alpha = \min(K_{\tilde{C}_\alpha} \setminus \{\mu_\beta : \beta < \alpha\})$. Then $C_\alpha = \tilde{C}_\alpha \cap N_{\mu_\alpha}$ is as required.

Once the recursion is completed, we can move to the subgraph on $(A, \{\mu_\alpha : \alpha < \omega_1\})$ with neighbourhoods $N(\mu_\alpha)$ given by C_α . By property (2), the claim applies and we obtain an AD- (\aleph_0, \aleph_1) -subgraph.

Thus, we can assume that every infinite subset of \mathbb{N} , and in particular every N_α contains an infinite subset C_α such that K_{C_α} is countable. Recursively, pick an increasing transfinite subsequence $\{\nu_\alpha : \alpha < \omega_1\}$ of ω_1 , defined recursively by $\nu_0 = 0$ and

$$\nu_\alpha = \sup \left(\{\nu_\beta : \beta < \alpha\} \cup \bigcup_{\beta < \alpha} K_{C_{\nu_\beta}} \right) + 1 < \omega_1.$$

We claim that $\{C_{\nu_\alpha} : \alpha < \omega_1\}$ gives rise to an AD- (\aleph_0, \aleph_1) -subgraph of G . It is a subgraph, since by construction, we have $C_{\nu_\alpha} \subseteq N(\nu_\alpha)$. And it is almost disjoint, since given two arbitrary neighbourhoods C_{ν_α} and C_{ν_β} with say $\nu_\alpha < \nu_\beta$, we have $C_{\nu_\alpha} \cap C_{\nu_\beta} \subseteq C_{\nu_\alpha} \cap N_{\nu_\beta}$, which is finite since $\nu_\beta \notin K_{\nu_\alpha}$ by construction. \square

For completeness, we provide the proof of Theorem 7.3.1.

PROOF OF THEOREM 7.3.1. Let $\mathcal{F} = \{F_\alpha : \alpha < \kappa\}$ be a $\kappa < \mathfrak{c}$ sized family of infinite subsets of \mathbb{N} . We want to find an almost disjoint family $\mathcal{B} = \{B_\alpha : \alpha < \kappa\}$ such that $B_\alpha \subseteq F_\alpha$ for all $\alpha < \kappa$.

Step 1: Split each F_α into an almost disjoint family $\mathcal{S}_\alpha = \{S_\xi^\alpha : \xi < \kappa^+\}$, i.e. all S_ξ^α are infinite subsets of F_α , and $S_\xi^\alpha \cap S_\zeta^\alpha$ is finite whenever $\xi \neq \zeta < \kappa^+$. As $\kappa^+ \leq \mathfrak{c}$, this is always possible. Note that κ^+ is a regular cardinal.

Step 2: From our definition of ‘covering’ after Theorem 7.3.3, it follows that a κ^+ -sized AD-family \mathcal{S}_α covers $\{F_\beta\}$ iff $\{S_\xi^\alpha \cap F_\beta : |S_\xi^\alpha \cap F_\beta| = \infty\}$ is a κ^+ -sized AD-family on F_β . For all $\alpha < \kappa$ we use

$$Y_\alpha = \{\beta < \kappa : \mathcal{S}_\alpha \text{ covers } \{F_\beta\}\}$$

to build a partition of κ into (possibly empty) sets $\{X_\alpha : \alpha < \kappa\}$, defined by $X_0 = Y_0$ and $X_\alpha = Y_\alpha \setminus \bigcup_{\beta < \alpha} Y_\beta$.

Step 3: For all $\alpha \notin Y_\beta$ there is $\kappa(\alpha, \beta) < \kappa^+$ such that $|F_\alpha \cap S_\xi^\beta| < \infty$ for all $\xi \geq \kappa(\alpha, \beta)$. Define

$$\eta = \sup \{\kappa(\alpha, \beta) : \beta < \kappa, \alpha \notin Y_\beta\} < \kappa^+.$$

Step 4: Here, we pick the almost disjoint refinement. For all β there is $\alpha(\beta)$ such that $\beta \in X_{\alpha(\beta)}$. For all $\beta \in X_\alpha$ we choose different $\xi(\beta) > \eta$ and define $B_\beta = S_{\xi(\beta)}^{\alpha(\beta)} \cap F_\beta$. Since the X_α form a partition of κ , this is a well-defined assignment. Now consider $\beta < \gamma$. We need to show that $B_\beta \cap B_\gamma$ is finite.

- If $\alpha(\beta) = \alpha = \alpha(\gamma)$ then $B_\beta \cap B_\gamma \subseteq S_{\xi(\beta)}^\alpha \cap S_{\xi(\gamma)}^\alpha$ which is finite, since both sets are elements of the same AD-family \mathcal{S}_α .
- Otherwise, if say $\alpha(\beta) < \alpha(\gamma)$, then $\gamma \notin Y_{\alpha(\beta)}$, so $B_\beta \cap B_\gamma \subseteq S_{\xi(\beta)}^{\alpha(\beta)} \cap F_\gamma$ is finite since $\xi(\beta) > \eta \geq \kappa(\gamma, \alpha(\beta))$. \square

7.4. The situation under Martin's Axiom

In this section we prove that under $MA + \neg CH$, any binary tree with tops serves as a one-element obstruction set for the class of (\aleph_0, \aleph_1) -graphs. For background on Martin's Axiom, see [106, III.3]. We begin with a sequence of lemmas.

LEMMA 7.4.1. *Under $MA + \neg CH$, every (\aleph_0, \aleph_1) -graph contains a spanning subgraph isomorphic to a binary tree with tops.*

PROOF. Let (A, B) be an (\aleph_0, \aleph_1) -graph. We want to find an infinite set $T \subseteq A$ plus a tree order \prec on T such that $\mathcal{T} = (T, \prec)$ is isomorphic to $2^{<\omega}$, and an injective map $h: B \rightarrow \mathcal{B}(\mathcal{T})$ (assigning to each element $b \in B$ a unique branch of \mathcal{T}) such that $N(b) \cap h(b)$ is infinite for all $b \in B$. Once we have achieved this, we delete for every $b \in B$ all edges from b to $A \setminus h(b)$ to obtain a binary tree with tops with bipartition (T, B) . The remaining vertices in $A \setminus T$ can be easily interweaved with \mathcal{T} as isolated vertices to obtain a spanning such subgraph.

To build this tree \mathcal{T} , we consider finite approximations (T_p, \prec_p) to \mathcal{T} (which will be finite initial segments of \mathcal{T}), and then use Martin's Axiom to find a consistent way to build the desired full binary tree. Formally, consider the partial order (\mathbb{P}, \leq) consisting of tuples $p = (T_p, \prec_p, B_p, h_p)$ such that

- $T_p \subseteq A$ finite, and \prec_p a tree-order on T_p such that (T_p, \prec_p) is a binary tree of some finite height,
- $B_p \subseteq B$ finite, and
- $h_p: B_p \rightarrow \mathcal{B}((T_p, \prec_p))$ an injective assignment of branches,

and $p \leq q$ if

- (T_q, \prec_q) is an initial subtree of (T_p, \prec_p) ,
- $B_q \subseteq B_p$, and
- h_p extends h_q in the sense $h_p(b) \supseteq h_q(b)$ for all $b \in B_q$.

To see that (\mathbb{P}, \leq) is ccc, consider an uncountable collection

$$\{p_\alpha = (T_\alpha, \prec_\alpha, B_\alpha, h_\alpha) : \alpha < \omega_1\} \subseteq \mathbb{P}.$$

By the Δ -System Lemma [106, III.2.6], there is a finite root $R \subseteq B$ and an uncountable $K \subseteq \omega_1$ such that $B_\alpha \cap B_\beta = R$ for all $\alpha \neq \beta \in K$. And since there are only countably many finite subsets of A , each with only finitely many possible tree-orders and branch-assignments for R , there is an uncountable $K' \subseteq K$ such that $(T_\alpha, \prec_\alpha) = (T_\beta, \prec_\beta)$ and $h_\alpha \upharpoonright R = h_\beta \upharpoonright R$ for all $\alpha \neq \beta \in K'$. But then for any $\alpha \neq \beta \in K'$, $q = (T_\alpha, \prec_\alpha, B_\alpha \cup B_\beta, h_\alpha \cup h_\beta)$ is a condition below p_α and p_β (where we possibly have to increase T_α by one level so a suitable extension of $h_\alpha \cup h_\beta$ can be injective).

Next we claim that for all $b \in B$ and $n \in \omega$, the set

$$D_{b,n} = \{p \in \mathbb{P} : b \in B_p \text{ and } |h_p(b) \cap N(b)| \geq n\}$$

is dense. To see this, consider any condition $q \in \mathbb{P}$ and suppose (T_q, \prec_q) has height k . Choose any subset of $F_b \subseteq N(b) \setminus T_q$ of size n , and extend T_q to a full binary tree T_p of height $k + n$, making sure that $F_b \subseteq h_p(b)$.

Finally, by Martin's Axiom there is a filter \mathcal{G} meeting each of our $\aleph_1 < \mathfrak{c}$ many dense sets in $\mathcal{D} = \{D_{b,n} : b \in B, n \in \omega\}$. Then

$$\mathcal{T} = (T, \prec) = \left(\bigcup_{p \in \mathcal{G}} T_p, \bigcup_{p \in \mathcal{G}} \prec_p \right)$$

is a countable binary tree, and

$$h: B \rightarrow \mathcal{B}(\mathcal{T}), b \mapsto \bigcup_{p \in \mathcal{G}} h_p(b)$$

is an injective function witnessing that $N(b) \cap h(b)$ is infinite, for our dense sets make sure it has cardinality at least n for all $n \in \mathbb{N}$. \square

We remark that it has been shown in either of [144, Thm. 6], [160, 2.3] or [138, 4.4] (in historical order) that under $\text{MA} + \neg \text{CH}$, every almost disjoint family of size $< \mathfrak{c}$ contains a hidden tree family, which together with our Theorem 7.3.3 and the observations in Section 7.2 implies the result of Lemma 7.4.1.

However, we will now strengthen the claim of Lemma 7.4.1 to hold for *full binary trees with tops*. Clearly, binary trees with tops have fewer edges, and are therefore easier to find as subgraphs than full binary trees with tops. But under Martin's Axiom, it turns out that the additional leeway is not needed. Note though that in the previous theorem, we could find a *spanning* binary tree with tops. In the next theorem, we can obtain full binary trees with tops as subgraphs, but can no longer guarantee that they are spanning.

LEMMA 7.4.2. *Under $\text{MA} + \neg \text{CH}$, every (\aleph_0, \aleph_1) -graph contains a full binary tree with tops as a subgraph.*

PROOF. Let (A, B) be an (\aleph_0, \aleph_1) -graph. We want to find an infinite set $T \subseteq A$ plus a tree order \prec on T such that $\mathcal{T} = (T, \prec)$ is isomorphic to $2^{<\omega}$, and an uncountable $B_T \subseteq B$ plus an injective map $h: B_T \rightarrow \mathcal{B}(\mathcal{T})$ (assigning to each element $b \in B_T$ a unique branch of \mathcal{T}) such that $h(b) \subseteq N(b)$ for all $b \in B_T$. Once we have achieved this, we delete for every $b \in B_T$ all edges from b to $T \setminus h(b)$ to obtain the desired full binary tree (T, B_T) with tops.

To find this tree \mathcal{T} , we build countably many such trees in parallel, which together take care of all $b \in B$. Consider the partial order (\mathbb{P}, \leq) consisting of tuples $p = (T_p, \prec_p, B_p, h_p)$ such that

- $T_p \subseteq A$ finite, and \prec_p a tree-order on T_p such that (T_p, \prec_p) is a binary tree of some finite height,
- $B_p \subseteq B$ finite,

- $h_p: B_p \rightarrow \mathcal{B}((T_p, \prec_p))$ an injective assignment of branches, and
- $h_p(b) \subseteq N(b)$ for all $b \in B_p$

and $p \leq q$ if

- (T_q, \prec_q) is an initial subtree of (T_p, \prec_p) ,
- $B_q \subseteq B_p$, and
- h_p extends h_q in the sense $h_p(b) \supseteq h_q(b)$ for all $b \in B_q$.

As in the proof of Lemma 7.4.1, this partial order is ccc, and hence so is the finite support product

$$\prod_{n < \omega}^{\text{fin}} \mathbb{P} := \{\vec{p} \in \mathbb{P}^\omega : |\{n : \vec{p}_n \neq \mathbb{1}\}| < \infty\}$$

by [106, III.3.43].

We claim that for all $b \in B$, the set $D_b = \{\vec{p} : \exists n \in \omega \text{ s.t. } b \in B_{\vec{p}_n}\}$ is dense in $\prod_{n < \omega}^{\text{fin}} \mathbb{P}$. And indeed, to any condition \vec{p} which does not yet mention b we can simply add b to a free coordinate, even using the empty tree.

So by Martin's Axiom, there is a filter \mathcal{G} meeting every one of our $\aleph_1 < \mathfrak{c}$ many dense sets in $\mathcal{D} = \{D_b : b \in B\}$. It follows that

$$\left\{ (T_n, B_n) = \left(\bigcup_{\vec{p} \in \mathcal{G}} T_{\vec{p}_n}, \bigcup_{\vec{p} \in \mathcal{G}} B_{\vec{p}_n} \right) : n \in \mathbb{N} \right\}$$

is a countable collection of binary trees with tops, such that $B = \bigcup_{n \in \mathbb{N}} B_n$. Thus, at least one of them, say B_n , is uncountable. It follows that in $(T, B_T) = (T_n, B_n)$ we have found our full binary tree with tops embedded as a subgraph as desired. \square

We now proceed to showing that under MA, any two binary trees with tops embed into each other. Consider the binary tree $T = 2^{<\omega}$. A subset $B \subseteq \mathcal{B}(T)$ of branches is called dense (or \aleph_1 -dense) if for every $t \in T$ the set $B(t) = \{b \in B : t \in b\}$ has size at least \aleph_0 (or \aleph_1 respectively).

It is well known that the Cantor set 2^ω is countable dense homogeneous, i.e. for every two countable dense subsets $A, B \subseteq 2^\omega$ there is a self-homeomorphism f of 2^ω such that $f(A) = B$. It is also known that under $\text{MA} + \neg \text{CH}$, this assertion can be strengthened to \aleph_1 -dense subsets of 2^ω , see for example [17, 3.2] and [145]. In the following, we shall see that a mild refinement of this approach, namely adding condition (d) to the partial order below, also works for (\aleph_0, \aleph_1) -graphs of binary type. In this condition (d) below, a level $\mathcal{T}(\alpha)$ of a tree \mathcal{T} is said to *separate* a collection of branches $B \subseteq \mathcal{B}(\mathcal{T})$ if $B(t) = \{b \in B : t \in b\}$ has size at most one for all $t \in \mathcal{T}(\alpha)$.

LEMMA 7.4.3. *Under $\text{MA} + \neg \text{CH}$, any two full \aleph_1 -dense binary trees with tops are isomorphic.*

PROOF. Suppose $G = (T_A, A)$ and $H = (T_B, B)$ are two full \aleph_1 -dense binary trees with tops. For convenience, we treat $a \in A$ as branch of the tree T_A . Recall that $a \upharpoonright n$ denotes the unique node of the branch a of height n .

It is clear that A and B can be partitioned into \aleph_1 many disjoint countable dense sets $\{A_\alpha : \alpha < \omega_1\}$ and $\{B_\alpha : \alpha < \omega_1\}$ respectively. Consider the partial order (\mathbb{P}, \leq) consisting of tuples $p = (f_p, g_p)$ such that

- (a) f_p is a finite injection with $\text{dom}(f_p) \subseteq A$ and $\text{ran}(f_p) \subseteq B$,
- (b) if $x \in A_\alpha$ then $f_p(x) \in B_\alpha$,
- (c) g_p is an order isomorphism between $T_A(\leq n_p)$ and $T_B(\leq n_p)$ for some $n_p \in \mathbb{N}$,
- (d) $T_A(n_p)$ separates $\text{dom}(f_p)$ and $T_B(n_p)$ separates $\text{ran}(f_p)$,
- (e) for all $a \in \text{dom}(f_p)$ we have $g_p(a \upharpoonright n_p) = f_p(a) \upharpoonright n_p$,

and define $p \leq q$ if

- $f_p \supseteq f_q$, and
- $g_p \supseteq g_q$.

To see that (\mathbb{P}, \leq) is ccc, consider an uncountable collection

$$\{p_\alpha = (f_\alpha, g_\alpha) : \alpha < \omega_1\} \subseteq \mathbb{P}.$$

Applying the Δ -System Lemma to all sets of the form $I_\alpha = \{\gamma : A_\gamma \cap \text{dom}(f_\alpha) \neq \emptyset\}$ (for $\alpha < \omega_1$), we obtain a finite root R and an uncountable $K \subseteq \omega_1$ such that $I_\alpha \cap I_\beta = R$ for all $\alpha \neq \beta \in K$.

Since there are only countably many different finite subsets of $A' = \bigcup_{\alpha \in R} A_\alpha$, we may assume that $\text{dom}(f_\alpha) \cap A' = S = \text{dom}(f_\beta) \cap A'$ for all $\alpha \neq \beta \in K$. And since (b) implies that there are only countably many choices for $f_\alpha \upharpoonright S$, we may assume that $f_\alpha \upharpoonright S = f_\beta \upharpoonright S$ for all $\alpha \neq \beta \in K$. Finally, since there are only countably many different g_α , we may assume that all $g_\alpha : T_A(\leq n) \rightarrow T_B(\leq n)$ agree.

But now any two conditions in $\{p_\alpha : \alpha \in K\}$ are compatible. By (b) and the definition of R , the map $f = f_\alpha \cup f_\beta$ is a well-defined injective partial map. Extend g_α to an order isomorphism $g : T_A(\leq m) \rightarrow T_B(\leq m)$ for some sufficiently large $m \geq n$, making sure that (d) and (e) are satisfied. Then (f, g) is a condition below f_α and f_β , so (\mathbb{P}, \leq) is ccc.

As our dense sets, we will consider

- (1) $D_n = \{p \in \mathbb{P} : T_A(\leq n) \subseteq \text{dom}(g_p)\}$, for $n \in \mathbb{N}$,
- (2) $D_a = \{p \in \mathbb{P} : a \in \text{dom}(f_p)\}$ for $a \in A$, and
- (3) $D_b = \{p \in \mathbb{P} : b \in \text{ran}(f_p)\}$ for $b \in B$.

To see that sets in (1) are dense, consider any condition $q = (f_q, g_q) \in \mathbb{P}$ and assume that $\text{dom}(g_q) = T_A(\leq m)$ for some $m < n$. Since for every $t \in T(m)$ there is at most one $a \in \text{dom}(f_q)$ such that $t \in a$ by (d), it is clear that we can extend g_q to a function g_p defined on $T_A(\leq n)$ by mapping the upset t^\uparrow in $T_A(\leq n)$ to the corresponding upset of

$g_q(t)^\uparrow$ of $T_B(\leq n)$ such that the branch $a \upharpoonright t^\uparrow$ is mapped to $f_q(a) \upharpoonright g_q(t)^\uparrow$. For $f_p = f_q$ we have $p = (f_p, g_p)$ is a condition in D_n below q .

To see that sets in (2) are dense, consider any condition $q \in \mathbb{P}$ and assume that $a \notin \text{dom}(f_q)$. Say $\text{dom}(g_q) = T_A(\leq n)$ for a given $n \in \mathbb{N}$. By (1) we may assume that $T_A(n)$ separates $\text{dom}(f_q) \cup \{a\}$. Find $t \in T_A(n)$ such that $t \in a$. Note that $a \in A_\alpha$ for some $\alpha < \omega_1$. By density of B_α , we may pick $b \in B_\alpha$ extending $g_q(t)$. Then $f_p = f_q \cup \langle a, b \rangle$ and $g_p = g_q$ gives a condition in D_a below q . The argument for (3) is similar.

Finally, Martin's Axiom gives us a filter \mathcal{G} meeting all specified dense sets. But then (2) and (3) force that $f = \bigcup_{p \in \mathcal{G}} f_p: A \rightarrow B$ is a bijection, and (1) forces that $g = \bigcup_{p \in \mathcal{G}} g_p: T_A \rightarrow T_B$ is an isomorphism of trees. In combination with property (e), we have $g[a] = f(a)$ for all $a \in A$, and this means, since G and H were full binary trees with tops, that $f \cup g: G \rightarrow H$ is an isomorphism of graphs. \square

THEOREM 7.4.4. *Under $MA + \neg CH$, any binary tree with tops embeds into all other (\aleph_0, \aleph_1) -graphs as a subgraph.*

PROOF. Suppose $G = (T_A, A)$ is a binary tree with tops, and H an arbitrary (\aleph_0, \aleph_1) -graph. Our task is to embed G into H as a subgraph. By Lemma 7.4.2, we may assume that $H = (T_B, B)$ is a full binary tree with tops.

Our plan is (a) to extend G to a full \aleph_1 -dense binary tree with tops G' , and (b) to find in H a full \aleph_1 -dense binary tree with tops H' as a subgraph. Then Lemma 7.4.3 implies that

$$G \hookrightarrow G' \cong H' \hookrightarrow H,$$

establishing the theorem.

Only item (b) requires proof. For this, we observe that every uncountable set of branches X of a binary tree T contains at least one *complete accumulation point*, i.e. a branch $x \in X$ such that for every $t \in x$, the set $B(t) = \{y \in X : t \in y\}$ is uncountable. Indeed, otherwise for every $x \in X$ there is t_x such that $B(t_x)$ is countable, and hence $X \subseteq \bigcup_{t_x \in T} B(t_x)$ is countable, a contradiction.

It follows that in fact all but at most countably many points of X are complete accumulation points, so without loss of generality, we may assume that every point of B is a complete accumulation point. Consider $T'_B = \bigcup_{b \in B} b \subseteq T_B$. Then T'_B is a (subdivided) binary tree, so after deleting all non-splitting nodes from T'_B , we obtain a full \aleph_1 -dense binary tree with tops H' as desired. The proof is complete. \square

7.5. A third type of (\aleph_0, \aleph_1) -graph

In this section we present a counterexample to the main open question from [58, §8], which is our Question 1 from the beginning.

THEOREM 7.5.1. *Under CH , there is an almost disjoint (\aleph_0, \aleph_1) -graph which contains no (\aleph_0, \aleph_1) -minor that is indivisible or of binary type.*

Our proof is inspired by the proof strategy of the following result due to Roitman & Soukup: *Under CH plus the existence of a Suslin tree, there is an uncountable anti-Luzin AD-family containing no uncountable hidden weak tree families* [138, 4.6]. Note though, that not containing a binary (\aleph_0, \aleph_1) -graph as a minor or just as a subgraph are stronger assertions than not containing an uncountable hidden weak tree family.

We shall make use of the following lemma.

LEMMA 7.5.2. *Whenever \mathcal{T}^* is Aronszajn, and B an uncountable set of branches of \mathcal{T}^* such that no two elements of B have the same order type, there are incompatible elements $s, t \in \mathcal{T}^*$ both contained in uncountably many branches of B .*

PROOF. The proof follows [138, 4.7]. Consider an Aronszajn tree \mathcal{T}^* , and let B be an uncountable set of branches of \mathcal{T}^* such that no two elements of B have the same order type.

Suppose for a contradiction that whenever s and t are incompatible, then either $B(s) = \{b \in B : s \in b\}$ is countable or $B(t) = \{b \in B : t \in b\}$ is countable. Then $S = \{s : B(s) \text{ is uncountable}\}$ forms a chain, hence is countable. So there is $\alpha < \omega_1$ with $\mathcal{T}^*(\alpha) \cap S = \emptyset$. But now all but countably many elements of B are contained in the countable set $\bigcup_{s \in \mathcal{T}^*(\alpha)} B(s)$, a contradiction. \square

PROOF OF THEOREM 7.5.1. Consider an Aronszajn tree \mathcal{T}^* , and let B be an uncountable set of branches of \mathcal{T}^* such that no two elements of B have the same order type.

Using CH, let $\{\mathcal{T}_\alpha = (T_\alpha, <_\alpha) : \alpha < \omega_1\}$ enumerate all trees of countable height whose underlying set is an infinite family of non-empty disjoint subsets of \mathbb{N} . For a subset $C \subseteq \mathbb{N}$ we define $C(\mathcal{T}_\alpha) = \{t \in T_\alpha : C \cap t \neq \emptyset\}$.

Let us construct, by recursion on $\alpha < \omega_1$,

- families $\{C_t : t \in \mathcal{T}^*(\alpha)\}$ of infinite subsets of \mathbb{N} , and
- countable families B_α of branches of \mathcal{T}_α ,

such that

- (a) for all $s, t \in \mathcal{T}^*$ we have $C_t \subseteq^* C_s$ if $s < t$, and $C_s \cap C_t =^* \emptyset$ if s and t are incomparable,
- (b) for all $s \neq t \in \mathcal{T}^*(\alpha)$, we have $C_s(\mathcal{T}_\alpha) \cap C_t(\mathcal{T}_\alpha) =^* \emptyset$, and
- (c) for all $t \in \mathcal{T}^*(\alpha)$, if $C_t(\mathcal{T}_\alpha)$ contains an infinite chain in \mathcal{T}_α , then there is $b \in B_\alpha$ such that $C_t(\mathcal{T}_\alpha) \subseteq^* b$.

For the construction, suppose for some $\alpha < \omega_1$ that we have already constructed infinite sets $C_t \subseteq \mathbb{N}$ for all $t \in \mathcal{T}^*$ of height strictly less than α . By (a), we may pick for every $t \in \mathcal{T}^*(\alpha)$ an infinite pseudo-intersection D_t of the family $\{C_s : s < t\}$. Using that every level $\mathcal{T}^*(\alpha)$ of our Aronszajn tree \mathcal{T}^* is countable, find an almost disjoint refinement $\{D'_t : t \in \mathcal{T}^*(\alpha)\}$ of $\{D_t : t \in \mathcal{T}^*(\alpha)\}$. This can be done either by hand, or by invoking

Theorem 7.3.1. Similarly, we can find a further refinement $\{D_t'' : t \in \mathcal{T}^*(\alpha)\}$ such that $D_s''(\mathcal{T}_\alpha) \cap D_t''(\mathcal{T}_\alpha) =^* \emptyset$ for all $s \neq t \in \mathcal{T}^*(\alpha)$. This takes care of property (b).

For (c), we use the Aronszajn property to enumerate $\mathcal{T}^*(\alpha) = \{t_n : n \in \mathbb{N}\}$. For $n \in \mathbb{N}$, if $D_{t_n}''(\mathcal{T}_\alpha)$ has infinite intersection with some branch of \mathcal{T}_α , we pick one such branch b_n and pick an infinite subset $C_{t_n} \subseteq D_{t_n}''$ such that $C_{t_n}(\mathcal{T}_\alpha) \subseteq b_n$. Otherwise, we simply put $C_{t_n} = D_{t_n}''$ (and let b_n be an arbitrary branch). This final refinement preserves (a) and (b), and after putting $B_\alpha = \{b_n : n \in \mathbb{N}\}$, we see that also (c) is satisfied.

Having completed the construction, we may pick, by (a), for every branch $b \in B$ an infinite pseudo-intersection $N(b)$ along the branch b , i.e. $N(b) \subseteq^* C_t$ for all $t \in b$. It follows from (a) that $\{N(b) : b \in B\}$ is an almost disjoint family of size ω_1 .

Let G be the almost disjoint (\aleph_0, \aleph_1) -graph with bipartition (\mathbb{N}, B) where the neighbourhood of $b \in B$ is $N(b)$.

CLAIM. Property (c) implies that no (\aleph_0, \aleph_1) -minor of G is of binary type.

To see the claim, suppose that $H = (\mathcal{T}, X)$ is an (\aleph_0, \aleph_1) -minor of G of binary type. Since any non-trivial branch set of the bipartite graph G must contain a vertex from \mathbb{N} , we may assume, without loss of generality, that $X \subseteq B$, and that every branch set $X_t \subseteq V(G)$ corresponding to a vertex of $t \in \mathcal{T}$ intersects \mathbb{N} . Further, there is an injective function $h : X \rightarrow Br(\mathcal{T})$ mapping points in X to branches of \mathcal{T} such that $N_G(x)(\mathcal{T}) \cap h(x)$ is infinite for all $x \in X$.

However, the tree $\mathcal{T} = \mathcal{T}_\alpha$ appears in our enumeration. Without loss of generality, $X \subseteq \{b \in B : ht(b) > \alpha\}$. But then (c) implies that $\text{ran}(h) \subseteq B_\alpha$, which is countable, contradicting that X is uncountable and h injective.

CLAIM. Property (b) implies that every (\aleph_0, \aleph_1) -minor of G is divisible.

Suppose that $H = (A, X)$ is an (\aleph_0, \aleph_1) -minor of G . As before, we may assume that $X \subseteq B$ and that the branch sets $X_a \subseteq V(G)$ for $a \in A$ intersect \mathbb{N} . Note that $\mathcal{X} = \{X_a \cap \mathbb{N} : a \in A\}$ is the underlying set of uncountably many of our trees T_α .

Now by Lemma 7.5.2, there are incomparable $s, t \in \mathcal{T}^*$ each contained in uncountably many branches of X . Find $\alpha \geq ht(s), ht(t)$ such that $\mathcal{X} = T_\alpha$, and find $s', t' \in \mathcal{T}^*(\alpha)$ extending s and t respectively such that $\mathcal{C} = \{b \in X : s' \in b\}$ and $\mathcal{D} = \{b \in X : t' \in b\}$ are both uncountable.

But then (b) implies that $C_{s'}(\mathcal{T}_\alpha)$ and its complement witness that H is divisible. Indeed, each $b \in \mathcal{C}$ has co-finitely many of its neighbours in $C_{s'}(\mathcal{T}_\alpha)$, since $N(b) \subseteq^* C_{s'}$ for all $b \in \mathcal{C}$, and similarly, each $b \in \mathcal{D}$ has co-finitely many of its neighbours in $C_{t'}(\mathcal{T}_\alpha)$, as $N(b) \subseteq^* C_{t'}$ for all $b \in \mathcal{D}$. \square

Since every AD-family built in the above way satisfying (a) is anti-Luzin [138, 4.10], we obtain the following corollary.

COROLLARY 7.5.3. *Under CH, there is an uncountable anti-Luzin AD-family which contains no uncountable hidden weak tree families.*

This improves the corresponding result from [138, 4.6], where it was proved under the additional assumption of the existence of a Suslin tree.

7.6. More on indivisible graphs

In this final section, we investigate indivisible graphs in more detail. Our aim is to construct a counterexample to Question 2 from the introduction. First however, we consider the question of when precisely indivisible graphs exist.

We recall two cardinal invariants in infinite combinatorics. The ultrafilter number \mathfrak{u} is the least cardinal of a collection \mathcal{U} of infinite subsets of \mathbb{N} that form a base of some non-principal ultrafilter on \mathbb{N} . In formulas,

$$\mathfrak{u} = \min \{ |\mathcal{U}| : \mathcal{U} \subseteq [\mathbb{N}]^\omega \text{ is a base for a non-principal ultrafilter on } \mathbb{N} \}.$$

(Recall that \mathcal{U} is a base for an ultrafilter \mathcal{V} if $\mathcal{U} \subseteq \mathcal{V}$ and for all $V \in \mathcal{V}$ there is $U \in \mathcal{U}$ such that $U \subseteq V$.) We call $\mathcal{R} \subseteq [\mathbb{N}]^\omega$ a *reaping family* if for all $A \in [\mathbb{N}]^\omega$ there is $R \in \mathcal{R}$ such that either $|A \cap R|$ or $|R \setminus A|$ is finite. The *reaping number* \mathfrak{r} is the least size of a reaping family. In formulas,

$$\mathfrak{r} = \min \{ |\mathcal{R}| : \mathcal{R} \subseteq [\mathbb{N}]^\omega \text{ and } \forall A \in [\mathbb{N}]^\omega \exists R \in \mathcal{R} (A \cap R =^* \emptyset \vee R \setminus A =^* \emptyset) \}.$$

THEOREM 7.6.1. *The equality $\mathfrak{u} = \omega_1$ implies that indivisible (\aleph_0, \aleph_1) -graphs exist, whereas $\mathfrak{r} > \omega_1$ implies they do not exist.*

PROOF. Let \mathcal{V} be a non-principal ultrafilter and let $\{U_\alpha : \alpha < \omega_1\}$ be a base for \mathcal{V} . We will build an indivisible (\aleph_0, \aleph_1) -graph with bipartition (\mathbb{N}, B) as follows. Let $B = \{b_\alpha : \alpha < \omega_1\}$. For every b_α we let $N(b_\alpha)$ be an infinite pseudo-intersection of the family $(U_\beta)_{\beta < \alpha}$. It is easy to check that this yields a graph as desired.

Conversely, if (\mathbb{N}, B) is indivisible, then for every $A \subseteq \mathbb{N}$, all but countably many elements of $\{N(b) : b \in B\}$ are almost contained in A or almost disjoint from A . It follows that $\{N(b) : b \in B\}$ is a reaping family and therefore $\mathfrak{r} = \omega_1$. \square

In particular, it is well-known (see [159]) that we have $\omega_1 \leq \mathfrak{r} = \pi\mathfrak{u} \leq \mathfrak{u} \leq \mathfrak{c}$, where $\pi\mathfrak{u}$ is the least cardinal of a local π -base of some non-principal ultrafilter on \mathbb{N} . Since it is consistent that $\omega_1 = \mathfrak{u} < \mathfrak{c}$, it follows that CH is independent of the existence of indivisible (\aleph_0, \aleph_1) -graphs. However, we do not know whether indivisible graphs exist in the Bell-Kunen model where $\omega_1 = \pi\mathfrak{u} < \mathfrak{u}$, [19].

Lastly, we observe the following connection between indivisible graphs and π -bases: The neighbourhoods $N(b_\alpha)$ of an \mathcal{U} -indivisible (\aleph_0, \aleph_1) -graph form a π -base for \mathcal{U} . And conversely, if a family $\{N_\alpha : \alpha < \omega_1\}$ of infinite subsets of \mathbb{N} forms a π -base for a *unique* ultrafilter \mathcal{U} , then the corresponding (\aleph_0, \aleph_1) -graph is indivisible.

We are now ready to answer Question 2 in the negative.

THEOREM 7.6.2. *Assume CH. Let \mathcal{U} be a non-principal ultrafilter on the natural numbers. For every \mathcal{U} -indivisible (\aleph_0, \aleph_1) -graph G there exists an \mathcal{U} -indivisible (\aleph_0, \aleph_1) -graph H such that $G \not\leq H$.*

PROOF. Using CH, let $\{U_\alpha : \alpha < \omega_1\}$ be an enumeration of the elements of \mathcal{U} , and let $\{\mathcal{X}_\alpha : \alpha < \omega_1\}$ be an enumeration of all infinite sequences of non-empty disjoint subsets of \mathbb{N} . For $\alpha < \omega_1$ write $\mathcal{X}_\alpha = (X_n^\alpha : n \in \mathbb{N}) \in \mathcal{P}(\mathbb{N})^{\mathbb{N}}$.

Suppose G is a \mathcal{U} -indivisible (\aleph_0, \aleph_1) -graph with bipartition (\mathbb{N}, B) . We write $B = \{b_\alpha : \alpha < \omega_1\}$. Our graph H will be an (\aleph_0, \aleph_1) -graph with bipartition (\mathbb{N}, C) where $C = \{c_\alpha : \alpha < \omega_1\}$. Our task is to define suitable neighbourhoods $N(c_\alpha)$ for all $\alpha < \omega_1$. We will do this as follows. At step $\alpha < \omega_1$, choose a neighbourhood $N(c_\alpha) \subseteq \mathbb{N}$ such that

- (1) $N(c_\alpha) \subseteq^* U_\beta$ for all $\beta \leq \alpha$, and
- (2) for any $\gamma, \delta \leq \alpha$ there is $n \in N(b_\gamma)$ such that $N(c_\alpha) \cap X_n^\delta = \emptyset$.

To build the neighbourhood $N(c_\alpha) = \{m_k : k \in \mathbb{N}\}$ recursively, enumerate the set $\{U_\beta : \beta \leq \alpha\}$ as $\{U^n : n \in \mathbb{N}\}$ and $\{(\beta, \gamma) : \beta, \gamma \leq \alpha\}$ as $\{(\beta_n, \gamma_n) : n \in \mathbb{N}\}$.

To choose m_k , note that since the collection $\{X_n^{\gamma_k} : n \in N(b_{\beta_k})\}$ is infinite and disjoint, there is an index $n_k \in N(b_{\beta_k})$ such that $X_{n_k}^{\gamma_k} \notin \mathcal{U}$ and $X_{n_k}^{\gamma_k} \cap \{m_l : l < k\} = \emptyset$. Now pick

$$m_k \in \bigcap_{l \leq k} (U^l \setminus X_{n_l}^{\gamma_l}) \in \mathcal{U}.$$

This choice of $N(c_\alpha) = \{m_k : k \in \mathbb{N}\}$ clearly satisfies (1). To see that it satisfies (2), note that $X_{n_k}^{\gamma_k} \cap \{m_l : l < k\} = \emptyset$ by our choice of n_k , and $X_{n_k}^{\gamma_k} \cap \{m_l : l \geq k\} = \emptyset$ by our choice of the m_l for $l \geq k$. This completes the recursive construction of the graph H .

CLAIM. H is \mathcal{U} -indivisible.

This is immediate from (1).

CLAIM. G is not a minor of H .

Suppose for contradiction that it is. Without loss of generality, we may assume that every vertex on the \mathbb{N} -side of G has uncountable degree. Write $V_n, W_\alpha \subseteq V(H)$ ($n \in \mathbb{N}$, $\alpha < \omega_1$) for the branch sets of the vertices in \mathbb{N} and B respectively. By our assumption on the degrees of the vertices on the \mathbb{N} -side of G , it follows that $V_n \cap \mathbb{N} \neq \emptyset$ for all $n \in \mathbb{N}$. Thus, $(V_n \cap \mathbb{N} : n \in \mathbb{N}) = \mathcal{X}_\gamma$ for some $\gamma < \omega_1$.

Also, since only countably many branch sets can intersect \mathbb{N} , there is some $\delta < \omega_1$ such that $W_\alpha = \{c_{\beta(\alpha)}\}$ for all $\alpha > \delta$. Also, since branch sets must be disjoint, the function $\beta : \alpha \mapsto \beta(\alpha)$ is injective.

Let $\eta = \max\{\gamma, \delta\}$. We claim that for all $\alpha > \eta$, we have $\beta(\alpha) < \alpha$. Indeed, W_α needs to have an edge to all V_n for $n \in N(b_\alpha)$, which requires that $c_{\beta(\alpha)}$ has an edge to

X_n^γ for all $n \in N(b_\alpha)$. However, if $\alpha \leq \beta(\alpha)$, then this is impossible, as (2) implies that $N(c_{\beta(\alpha)}) \cap X_{n_0}^\gamma = \emptyset$ for at least one $n_0 \in N(b_\alpha)$.

Thus, we have $\beta(\alpha) < \alpha$ for all $\alpha > \eta$. By Fodor's Lemma [106, III.6.14], however, this implies that the map β is constant on an uncountable subset of ω_1 , contradicting its injectivity. \square

QUESTION 3. Assume CH. Is it true that for every \mathcal{U} -indivisible (\aleph_0, \aleph_1) -graph G there exists a \mathcal{U} -indivisible (\aleph_0, \aleph_1) -graph H such that both $G \not\preceq H$ and $H \not\preceq G$?

Part 2

Topological infinite graph theory

Hamilton decompositions of one-ended Cayley graphs

We prove that any one-ended, locally finite Cayley graph $G(\Gamma, S)$, where Γ is an abelian group and S is a finite generating set of non-torsion elements, admits a decomposition into edge-disjoint Hamiltonian (i.e. spanning) double-rays. In particular, the n -dimensional grid \mathbb{Z}^n admits a decomposition into n edge-disjoint Hamiltonian double-rays for all $n \in \mathbb{N}$.

8.1. Introduction

A *Hamiltonian cycle* of a finite graph is a cycle which includes every vertex of the graph. A finite graph $G = (V, E)$ is said to have a *Hamilton decomposition* if its edge set can be partitioned into disjoint sets $E = E_1 \dot{\cup} E_2 \dot{\cup} \dots \dot{\cup} E_r$ such that each E_i is a Hamiltonian cycle in G .

The starting point for the theory of Hamilton decompositions is an old result by Walecki from 1890 according to which every finite complete graph of odd order has a Hamilton decomposition (see [4] for a description of his construction). Since then, this result has been extended in various different ways, and we refer the reader to the survey of Alspach, Bermond and Sotteau [5] for more information.

Hamiltonicity problems have also been considered for infinite graphs, see for example the survey by Gallian and Witte [167]. While it is sometimes not obvious which objects should be considered the correct generalisations of a Hamiltonian cycle in the setting of infinite graphs, for one-ended graphs the undisputed solution is to consider *double-rays*, i.e. infinite, connected, 2-regular subgraphs. Thus, for us a *Hamiltonian double-ray* is then a double-ray which includes every vertex of the graph, and we say that an infinite graph $G = (V, E)$ has a *Hamilton decomposition* if we can partition its edge set into edge-disjoint Hamiltonian double-rays.

In this paper we will consider infinite variants of two long-standing conjectures on the existence of Hamilton decompositions for finite graphs. The first conjecture concerns Cayley graphs: Given a finitely generated abelian group $(\Gamma, +)$ and a finite generating set S of Γ , the *Cayley graph* $G(\Gamma, S)$ is the multi-graph with vertex set Γ and edge multi-set

$$\{(x, x + g) : x \in \Gamma, g \in S\}.$$

CONJECTURE 8.1.1 (Alspach [2, 3]). *If Γ is an abelian group and S generates G , then the simplification of $G(\Gamma, S)$ has a Hamilton decomposition, provided that it is $2k$ -regular for some k .*

Note that if $S \cap -S = \emptyset$, then $G(\Gamma, S)$ is automatically a $2|S|$ -regular simple graph. If $G(\Gamma, S)$ is finite and 2-regular, then the conjecture is trivially true. Bermond, Favaron and Maheo [22] showed that the conjecture holds in the case $k = 2$. Liu [113] proved certain cases of the conjecture for finite 6-regular Cayley graphs, and his result was further extended by Westlund [164]. Liu [114, 115] also gave some sufficient conditions on the generating set S for such a decomposition to exist.

Our main theorem in this paper is the following affirmative result towards the corresponding infinite analogue of Conjecture 8.1.1:

THEOREM 8.1.2. *Let Γ be an infinite, finitely generated abelian group, and let S be a generating set such that every element of S has infinite order. If the Cayley graph $G = G(\Gamma, S)$ is one-ended, then it has a Hamilton decomposition.*

We remark that under the assumption that elements of S are non-torsion, the simplification of $G(\Gamma, S)$ is always isomorphic to a Cayley graph $G(\Gamma, S')$ with $S' \subseteq S$ and $S' \cap -S' = \emptyset$, and so our theorem implies the corresponding version of Conjecture 8.1.1 for non-torsion generators, in particular for Cayley graphs of \mathbb{Z}^n with arbitrary generators.

In the case when $G = G(\Gamma, S)$ is two-ended, there are additional technical difficulties when trying to construct a decomposition into Hamiltonian double-rays. In particular, since each Hamiltonian double-ray must meet every finite edge cut an odd number of times, there can be parity reasons why no decomposition exists. One particular two-ended case, namely where $\Gamma \cong \mathbb{Z}$, has been considered by Bryant, Herke, Maenhaut and Webb [40], who showed that when $G(\mathbb{Z}, S)$ is 4-regular, then G has a Hamilton decomposition unless there is an odd cut separating the two ends.

The second conjecture about Hamiltonicity that we consider concerns Cartesian products of graphs: Given two graphs G and H the *Cartesian product* (or product) $G \square H$ is the graph with vertex set $V(G) \times V(H)$ in which two vertices (g, h) and (g', h') are adjacent if and only if either

- $g = g'$ and h is adjacent to h' in H , or
- $h = h'$ and g is adjacent to g' in G .

Kotzig [103] showed that the Cartesian product of two cycles has a Hamilton decomposition, and conjectured that this should be true for the product of three cycles. Bermond extended this conjecture to the following:

CONJECTURE 8.1.3 (Bermond [21]). *If G_1 and G_2 are finite graphs which both have Hamilton decompositions, then so does $G_1 \square G_2$.*

Alspach and Godsil [6] showed that the product of any finite number of cycles has a Hamilton decomposition, and Stong [147] proved certain cases of Conjecture 8.1.3 under additional assumptions on the number of Hamilton cycles in the decomposition of G_1 and G_2 respectively.

Applying techniques we developed to prove Theorem 8.1.2, we show as our second main result of this paper that Conjecture 8.1.3 holds for countably infinite multi-graphs.

THEOREM 8.1.4. *If G and H are countable multi-graphs which both have Hamilton decompositions, then so does their product $G \square H$.*

Note that the restriction to countable multi-graphs, i.e multi-graphs with countably many vertices and edges, is necessary. Indeed the existence of a spanning double ray implies that G and H have countable vertex sets. But then if G contains a countable edge cut, then so does $G \square H$. However, if H has uncountably many edges, then any Hamilton decomposition of $G \square H$ must consist of uncountably many edge-disjoint double-rays, contradicting the existence of a countable edge cut.

The paper is structured as follows: In Section 8.2 we mention some group theoretic results and definitions we will need. In Section 8.3 we state our main lemma, the *Covering Lemma*, and show that it implies Theorem 8.1.2. The proof of the Covering Lemma will be the content of Section 8.4. In Section 8.5 we apply our techniques to prove Theorem 8.1.4. Finally, in Section 8.6 we list open problems and possible directions for further work.

8.2. Notation and preliminaries

If $G = (V, E)$ is a graph, and $A, B \subseteq V$, we denote by $E(A, B)$ the set of edges between A and B , i.e. $E(A, B) = \{(x, y) \in E : x \in A, y \in B\}$. For $A \subseteq V$ or $F \subseteq E$ we write $G[A]$ and $G[F]$ for the subgraph of G induced by A and F respectively.

For $A, B \subseteq \Gamma$ subsets of an abelian group Γ we write $-A := \{-a : a \in A\}$ and $A + B := \{a + b : a \in A, b \in B\} \subseteq \Gamma$. If Δ is a subgroup of Γ , and $A \subseteq \Gamma$ a subset, then $A^\Delta = \{a + \Delta : a \in A\}$ denotes the family of corresponding cosets. If $g \in \Gamma$ we say that the *order* of g is the smallest $k \in \mathbb{N}$ such that $k \cdot g = 0$. If such a k exists, then g is a *torsion element*. Otherwise, we say the order of g is infinite and g is a *non-torsion element*. For $k \in \mathbb{N}$ we write $[k] = \{1, 2, \dots, k\}$.

The following terminology will be used throughout.

DEFINITION 8.2.1. Given a graph G , an edge-colouring $c : E(G) \rightarrow [s]$ and a colour $i \in [s]$, the *i -subgraph* is the subgraph of G induced by the edge set $c^{-1}(i)$, and the *i -components* are the components of the i -subgraph.

DEFINITION 8.2.2 (Standard and almost-standard colourings of Cayley graphs). Let Γ be an infinite abelian group, $S = \{g_1, g_2, \dots, g_s\}$ a finite generating set for Γ such that every $g_i \in S$ has infinite order, and let G be the Cayley graph $G(\Gamma, S)$.

- The *standard colouring* of G is the edge colouring $c_{\text{std}}: E(G) \rightarrow [s]$ such that $c_{\text{std}}((x, x + g_i)) = i$ for each $x \in \Gamma, g_i \in S$.
- Given a subset $X \subseteq V(G)$ we say that a colouring c is *standard on X* if c agrees with c_{std} on $G[X]$. Similarly if $F \subseteq E(G)$ we say that c is *standard on F* if c agrees with c_{std} on F .
- A colouring $c: E(G) \rightarrow [s]$ is *almost-standard* if the following are satisfied:
 - there is a finite subset $F \subseteq E(G)$ such that c is standard on $E(G) \setminus F$;
 - for each $i \in [s]$ the i -subgraph is spanning, and each i -component is a double-ray.

DEFINITION 8.2.3 (Standard squares and double-rays). Let Γ and S be as above. Given $x \in \Gamma$ and $g_i \neq g_j \in S$, we call

$$\blacksquare(x, g_i, g_j) := \{(x, x + g_i), (x, x + g_j), (x + g_i, x + g_i + g_j), (x + g_j, x + g_i + g_j)\}$$

an (i, j) -square with base point x , and

$$\longleftrightarrow(x, g_i) := \{(x + ng_i, x + (n + 1)g_i) : n \in \mathbb{Z}\}$$

an i -double-ray with base point x .

Moreover, given a colouring $c: E(G(\Gamma, S)) \rightarrow [s]$ we call $\blacksquare(x, g_i, g_j)$ and $\longleftrightarrow(x, g_i)$ an (i, j) -standard square and i -standard double-ray if c is standard on $\blacksquare(x, g_i, g_j)$ and $\longleftrightarrow(x, g_i)$ respectively.

Since Γ is an abelian group, every $\blacksquare(x, g_i, g_j)$ is a 4-cycle in $G(\Gamma, S)$ (provided $g_i \neq -g_j$), and since S contains no torsion elements of Γ , $\longleftrightarrow(x, g_k)$ really is a double-ray in the Cayley graph $G(\Gamma, S)$.

Let Γ be a finitely generated abelian group. By the Classification Theorem for finitely generated abelian groups (see e.g. [75]), there are integers n, q_1, \dots, q_r such that $\Gamma \cong \mathbb{Z}^n \oplus \bigoplus_{i=1}^r \mathbb{Z}_{q_i}$, where \mathbb{Z}_q is the additive group of the integers modulo q . In particular, for each Γ there is an integer n and a finite abelian group Γ_{fin} such that $\Gamma \cong \mathbb{Z}^n \oplus \Gamma_{\text{fin}}$.

The following structural theorem for the ends of finitely generated abelian groups is well-known:

THEOREM 8.2.4. *For a finitely generated group $\Gamma \cong \mathbb{Z}^n \oplus \Gamma_{\text{fin}}$, the following are equivalent:*

- $n \geq 2$,
- there exists a finite generating set S such that $G(\Gamma, S)$ is one-ended, and
- for all finite generating sets S , the Cayley graph $G(\Gamma, S)$ is one-ended.

PROOF. See e.g. [141, Proposition 5.2] for the fact the number of ends of $G(\Gamma, S)$ is independent of the choice of the generating set S , and [141, Theorem 5.12] for the equivalence with the first item. \square

A group Γ satisfying one of the conditions from Theorem 8.2.4 is called *one-ended*.

COROLLARY 8.2.5. *Let Γ be an abelian group, $S = \{g_1, \dots, g_s\}$ be a finite generating set such that the Cayley graph $G(\Gamma, S)$ is one-ended. Then, for every $g_i \in S$ of infinite order, there is some $g_j \in S$ such that $\langle g_i, g_j \rangle \cong (\mathbb{Z}^2, +)$.*

PROOF. Suppose not. It follows that in $\Gamma/\langle g_i \rangle$ every element has finite order, and since it is also finitely generated, it is some finite group Γ_f such that $\Gamma \cong \mathbb{Z} \oplus \Gamma_f$. Thus, by Theorem 8.2.4, G is not one-ended, a contradiction. \square

8.3. The covering lemma and a high-level proof of the main theorem

Every Cayley graph $G(\Gamma, S)$ comes with a natural edge colouring c_{std} , where we colour an edge $(x, x + g_i)$ with $x \in \Gamma$ and $g_i \in S$ with the index i of the corresponding generating element g_i . If every element of S has infinite order, then every i -subgraph of $G(\Gamma, S)$ consists of a spanning collection of edge-disjoint double-rays, see Definitions 8.2.1 and 8.2.2. So, it is perhaps a natural strategy to try to build a Hamiltonian decomposition by combining each of these monochromatic collections of double-rays into a single monochromatic spanning double-ray.

Rather than trying to do this directly, we shall do it in a series of steps: given any colour $i \in [s]$ where $s = |S|$ and any finite set $X \subseteq V(G)$, we will show that one can change the standard colouring at finitely many edges, in particular only edges outside of X , so that there is a single double-ray in the colour i which covers X . Moreover, we can ensure that the resulting colouring maintains enough of the structure of the standard colouring that we can repeat this process inductively: it should remain *almost-standard*, i.e. all monochromatic components are still double-rays, see Definition 8.2.2. By taking an appropriate sequence of sets $X_1 \subseteq X_2 \subseteq \dots$ exhausting the vertex set of G , and varying which colour i we consider, we can ensure that in the limit, each colour class consists of a single spanning double-ray, giving us the desired Hamilton decomposition.

In this section, we formulate our key lemma, namely the Covering Lemma 8.3.1, which allows us to do each of these steps. We will then show how Theorem 8.1.2 follows from the Covering Lemma. The proof of the Covering Lemma is given in Section 8.4.

LEMMA 8.3.1 (Covering lemma). *Let Γ be an infinite, one-ended abelian group, $S = \{g_1, g_2, \dots, g_s\}$ a finite generating set such that every $g_i \in S$ has infinite order, and $G = G(\Gamma, S)$ the corresponding Cayley graph.*

Then for every almost-standard colouring c of G , every colour i and every finite subset $X \subseteq V(G)$, there exists an almost-standard colouring \hat{c} of G such that

- $\hat{c} = c$ on $E(G[X])$, and
- some i -component in \hat{c} covers X .

PROOF OF THEOREM 8.1.2 GIVEN LEMMA 8.3.1. Fix an enumeration $V(G) = \{v_n : n \in \mathbb{N}\}$. Let $X_0 = D'_0 = D'_{-1} = \dots = D'_{-s+1} = \{v_0\}$ and $c_0 = c_{\text{std}}$. For each $n \geq 1$ we will recursively construct almost-standard colourings $c_n : E(G) \rightarrow [s]$, finite subsets $X_n \subseteq V(G)$, ($n \bmod s$)-components D_n of c_n and finite paths $D'_n \subseteq D_n$ such that for every $n \in \mathbb{N}$

- (1) $X_{n-1} \cup \{v_n\} \subseteq X_n$,
- (2) $V(D'_{n-1}) \subseteq X_n$,
- (3) $X_n \subseteq V(D'_n)$,
- (4) D'_n properly extends the path D'_{n-s} (the ‘previous’ path of colour $n \bmod s$) in both endpoints of D'_{n-s} , and
- (5) c_n agrees with c_{n-1} on $E(G[X_n])$.

Suppose inductively for some $n \in \mathbb{N}$ that c_n , X_n , D_n and D'_n have already been defined. Choose some $X_{n+1} \supseteq X_n \cup \{v_n\}$ large enough such that (1) and (2) are satisfied. Applying Lemma 8.3.1 with input c_n and X_{n+1} provides us with a colouring c_{n+1} such that (5) is satisfied and some $(n+1 \bmod s)$ -component D_{n+1} covers X_{n+1} . Since c_{n+1} is almost-standard, D_{n+1} is a double-ray. Furthermore, since c_{n+1} agrees with c_n on $E(G[X_{n+1}])$, by the inductive hypothesis it agrees with c_k on $E(G[X_{k+1}])$ for each $k \leq n$.

Therefore, since $D'_{n+1-s} \subseteq X_{n-s+2}$ is a path of colour $(n+1 \bmod s)$ in c_{n+1-s} , it follows that $D'_{n+1-s} \subseteq D_{n+1}$ and so we can extend D'_{n+1-s} to a sufficiently long finite path $D'_{n+1} \subseteq D_{n+1}$ such that (3) and (4) are satisfied at stage $n+1$.

Once the construction is complete, we define $T_1, \dots, T_s \subseteq G$ by

$$T_i = \bigcup_{n \equiv i \pmod{s}} D'_n$$

and claim that they form a decomposition of G into edge-disjoint Hamiltonian double-rays. Indeed, by (4), each T_i is a double-ray. That they are edge-disjoint can be seen as follows: Suppose for a contradiction that $e \in E(T_i) \cap E(T_j)$. Choose $n(i)$ and $n(j)$ minimal such that $e \in E(D'_{n(i)}) \subseteq E(T_i)$ and $e \in E(D'_{n(j)}) \subseteq E(T_j)$. We may assume that $n(i) < n(j)$, and so $e \in E(G[X_{n(i)+1}])$ by (2). Furthermore, by (5) it follows that $c_{n(j)}$ agrees with $c_{n(i)}$ on $E(G[X_{n(i)+1}])$. However by construction $c_{n(j)}(e) = j \neq i = c_{n(i)}(e)$ contradicting the previous line.

Finally, to see that each T_i is spanning, consider some $v_n \in V(G)$. By (1), $v_n \in X_n$. Pick $n' \geq n$ with $n' \equiv i \pmod{s}$. Then by (3), $D'_{n'} \subseteq T_i$ covers $X_{n'}$ which in turn contains v_n , as $v_n \in X_n \subseteq X_{n'}$ by (1). \square

8.4. Proof of the Covering Lemma

8.4.1. Blanket assumption. Throughout this section, let us now fix

- a one-ended infinite abelian group Γ with finite generating set $S = \{g_1, \dots, g_s\}$ such that every element of S has infinite order,
- an almost-standard colouring c of the Cayley graph $G = G(\Gamma, S)$,

- a finite subset $X \subseteq \Gamma$ such that c is standard on $V(G) \setminus X$,
- a colour i , say $i = 1$, and corresponding generator $g_1 \in S$, for which we want to show Lemma 8.3.1, and finally
- a second generator in S , say g_2 , such that $\Delta := \langle g_1, g_2 \rangle \cong (\mathbb{Z}^2, +)$, see Corollary 8.2.5.

8.4.2. Overview of proof. We want to show Lemma 8.3.1 for the Cayley graph G , colouring c , generator g_1 and finite set X . The cosets of $\langle g_1, g_2 \rangle$ in Γ cover $V(G)$, and in the standard colouring the edges of colour 1 and 2 form a grid on $\langle g_1, g_2 \rangle$. So, since c is almost-standard, on each of these cosets the edges of colour 1 and 2 will look like a grid, apart from some finite set.

Our aim is to use the structure in these grids to change the colouring c to one satisfying the conclusions of Lemma 8.3.1. It will be more convenient to work with large finite grids, which we require, for technical reasons, to have an even number of rows. This is the reason for the slight asymmetry in the definition below.

NOTATION 8.4.1. Let $g_i, g_j \in \Gamma$. For $N, M \in \mathbb{N}$ we write

$$\langle g_i, g_j \rangle_{N,M} := \{ng_i + mg_j : n, m \in \mathbb{Z}, -N \leq n \leq N, -M < m \leq M\} \subseteq \langle g_i, g_j \rangle \subseteq \Gamma.$$

The structure of our proof can be summarised as follows. First, in Section 8.4.3, we will show that there is some N_0 and some ‘nice’ finite set P of representatives of cosets of $\langle g_1, g_2 \rangle$ such that $P + \langle g_1, g_2 \rangle_{N_0, N_0}$ covers X . We will then, in Section 8.4.4 pick sufficiently large numbers $N_0 < N_1 < N_2 < N_3$ and consider the grids $P + \langle g_1, g_2 \rangle_{N_3, N_1}$. Using the structure of the grids we will make local changes to the colouring inside $P + (\langle g_1, g_2 \rangle_{N_3, N_1} \setminus \langle g_1, g_2 \rangle_{N_0, N_0})$ to construct our new colouring \hat{c} . This new colouring \hat{c} will then agree with c on the subgraph induced by $P + \langle g_1, g_2 \rangle_{N_0, N_0} \supseteq X$, and be standard on $V(G) \setminus (P + \langle g_1, g_2 \rangle_{N_3, N_1})$, and hence, as long as we ensure all the colour components are double-rays, almost-standard.

These local changes will happen in three steps. First, in Step 1, we will make local changes inside $x_\ell + (\langle g_1, g_2 \rangle_{N_3, N_1} \setminus \langle g_1, g_2 \rangle_{N_2, N_1})$ for each $x_\ell \in P$, in order to make every i -component meeting $P + \langle g_1, g_2 \rangle_{N_2, N_1}$ a finite cycle.

Next, in Step 2, we will make local changes inside $x_\ell + (\langle g_1, g_2 \rangle_{N_2, N_1} \setminus \langle g_1, g_2 \rangle_{N_1, N_1})$ for each $x_\ell \in P$, in order to combine the cycles meeting this translate of the grid into a single cycle.

Finally, in Step 3, we will make local changes inside $P + (\langle g_1, g_2 \rangle_{N_1, N_1} \setminus \langle g_1, g_2 \rangle_{N_0, N_0})$, in order to join the cycles for different x_ℓ into a single cycle covering $P + \langle g_1, g_2 \rangle_{N_0, N_0}$. We then make one final local change to turn this finite cycle into a double-ray.

8.4.3. Identifying the relevant cosets.

LEMMA 8.4.2. *There exist $N_0 \in \mathbb{N}$ and a finite set $P = \{x_0, \dots, x_t\} \subseteq \Gamma$ such that*

- $P^\Delta = \{x_0 + \Delta, \dots, x_t + \Delta\}$ is a path in $G(\Gamma/\Delta, (S \setminus \{g_1, g_2\})^\Delta)$, and
- $X \subseteq P + \langle g_1, g_2 \rangle_{N_0, N_0}$.

PROOF. Since X is finite, there is a finite set $Y = \{y_1, \dots, y_k\} \subseteq \Gamma$ such that the cosets in $Y^\Delta = \{y_1 + \Delta, \dots, y_k + \Delta\}$ are all distinct and cover X . Moreover, since every $(y_\ell + \Delta) \cap X$ is finite, there exists $N_0 \in \mathbb{N}$ such that

$$(y_\ell + \langle g_1, g_2 \rangle) \cap X = (y_\ell + \langle g_1, g_2 \rangle_{N_0, N_0}) \cap X$$

for all $1 \leq \ell \leq k$. Then $X \subseteq Y + \langle g_1, g_2 \rangle_{N_0, N_0}$.

Next, by a result of Nash-Williams [122], every Cayley graph of a countably infinite abelian group has a Hamilton double-ray, and it is a folklore result (see [167]) that every Cayley graph of a finite abelian group has a Hamilton cycle. So in particular, the Cayley graph of $(\Gamma/\Delta, (S \setminus \{g_1, g_2\})^\Delta)$, has a Hamilton cycle or double-ray, say H . Let $P \supseteq Y$ be a finite set of representatives of the cosets of Δ such that P^Δ is the set of vertices of a subpath of H . It is clear that P is as required. \square

- For the rest of this section let us fix $N_0 \in \mathbb{N}$ and $P = \{x_0, \dots, x_t\} \subseteq \Gamma$ to be as given by Lemma 8.4.2.

8.4.4. Picking sufficiently large grids. In order to choose our grids large enough to be able to make all the necessary changes to our colouring, we will first need the following lemma, which guarantees that we can find, for each $k \neq 1, 2$ and $x \in \Gamma$, many distinct standard k -double-rays which go between the cosets $x + \Delta$ and $(x + g_k) + \Delta$.

LEMMA 8.4.3. *For any $g_k \in S \setminus \{g_1, g_2\}$ and any pair of distinct cosets $x + \Delta$ and $(g_k + x) + \Delta$, there are infinitely many distinct standard k -double-rays R for the colouring c with $E(R) \cap E(x + \Delta, (g_k + x) + \Delta) \neq \emptyset$.*

PROOF. It clearly suffices to prove the assertion for $c = c_{\text{std}}$. We claim that either

$$\mathcal{R}_1 = \{\leftrightarrow(x + mg_1, g_k) : m \in \mathbb{Z}\} \text{ or } \mathcal{R}_2 = \{\leftrightarrow(x + mg_2, g_k) : m \in \mathbb{Z}\}$$

is such a collection of disjoint standard k -double-rays.

Suppose that \mathcal{R}_1 is not a collection of disjoint double-rays. Then there are $m \neq m' \in \mathbb{Z}$ and $n, n' \in \mathbb{Z}$ such that

$$mg_1 + ng_k = m'g_1 + n'g_k.$$

Since g_1 has infinite order, it follows that $n \neq n'$, too, and so we can conclude that there are $\ell, \ell' \in \mathbb{Z} \setminus \{0\}$ such that $\ell g_1 = \ell' g_k$. Similarly, if \mathcal{R}_2 is not a collection of disjoint double-rays, then we can find $q, q' \in \mathbb{Z} \setminus \{0\}$ such that $qg_2 = q'g_k$. However, it now follows that

$$q'\ell g_1 = q'(\ell' g_k) = \ell'(q'g_k) = \ell'qg_2,$$

contradicting the fact that $\langle g_1, g_2 \rangle \cong (\mathbb{Z}^2, +)$. This establishes the claim.

Finally, observe that if say \mathcal{R}_1 is a disjoint collection, then for every $R_m = \curvearrowright(x + mg_1, g_k) \in \mathcal{R}_1$ we have $(x + mg_1, x + mg_1 + g_k) \in E(R_m) \cap E(x + \Delta, (g_k + x) + \Delta)$ as desired. \square

We are now ready to define our numbers $N_0 < N_1 < N_2 < N_3$. Recall that N_0 and $P = \{x_0, \dots, x_t\}$ are given by Lemma 8.4.2. For each $\ell \in [t]$, let $g_{n(\ell)}$ be some generator in $S \setminus \{g_1, g_2\}$ that induces the edge between $x_{\ell-1} + \Delta$ and $x_\ell + \Delta$ on the path P^Δ . Note that $n(\ell) \in [s] \setminus \{1, 2\}$ for all ℓ .

By Lemma 8.4.3, we may find t^2 many disjoint standard double-rays

$$\mathcal{R} = \{R_\ell^k : 1 \leq k, \ell \leq t\}$$

such that for every ℓ , the double-rays in $\{R_\ell^k = \curvearrowright(y_\ell^k, g_{n(\ell)}) : k \in [t]\}$ are standard $n(\ell)$ -double-rays containing an edge

$$e_\ell^k = (y_\ell^k, y_\ell^k + g_{n(\ell)}) \in E(R_\ell^k) \cap E(x_{\ell-1} + \Delta, x_\ell + \Delta)$$

so that all $T_\ell^k = \blacksquare(y_\ell^k, g_1, g_{n(\ell)})$ are $(1, n(\ell))$ -standard squares for c which have empty intersection with $\{x_{\ell-1}, x_\ell\} + \langle g_1, g_2 \rangle_{N_0, N_0}$. Furthermore we may assume that these standard squares are all edge-disjoint. Then

- let $N_1 > N_0$ be sufficiently large such that the subgraph induced by $P + \langle g_1, g_2 \rangle_{N_1-3, N_1-3}$ contains all standard squares T_ℓ^k mentioned above.
- Let N_2 be arbitrary with $N_2 \geq 5N_1$.
- Let N_3 be arbitrary with $N_3 \geq N_2 + 2N_1$.

8.4.5. The cap-off step. Our main tool for locally modifying our colouring is the following notion of ‘colour switchings’, which is also used in [113]. Informally, given a four cycle on which the edge colouring alternates between two colours, to perform a colour switching on this square we exchange the colours of the edges.

DEFINITION 8.4.4 (Colour switching of standard squares). Given an edge colouring $c: E(G(\Gamma, S)) \rightarrow [s]$ and an (i, j) -standard square $\blacksquare(x, g_i, g_j)$, a *colour switching* on $\blacksquare(x, g_i, g_j)$ changes the colouring c to the colouring c' such that

- $c' = c$ on $E \setminus \blacksquare(x, g_i, g_j)$,
- $c'((x, x + g_i)) = c'((x + g_j, x + g_i + g_j)) = j$,
- $c'((x, x + g_j)) = c'((x + g_i, x + g_i + g_j)) = i$.

It would be convenient if colour switchings maintained the property that a colouring is almost-standard. Indeed, if c is standard on $E(G) \setminus F$ then c' is standard on $E(G) \setminus (F \cup \blacksquare(x, g_i, g_j))$. Also, it is a simple check that if the i and j -subgraphs of G for c are 2-regular and spanning, then the same is true for c' . However, some i or j -components may change from double-rays to finite cycles, and vice versa.

STEP 1 (Cap-off step). There is a colouring c' obtained from c by colour switchings of finitely many $(1, 2)$ -standard squares such that

- $c' = c$ on $E(G[X])$;
- every 1-component in c' meeting $P + \langle g_1, g_2 \rangle_{N_2, N_1}$ is a finite cycle intersecting both $P + (\langle g_1, g_2 \rangle_{N_3, N_1} \setminus \langle g_1, g_2 \rangle_{N_2, N_1})$ and $P + \langle g_1, g_2 \rangle_{N_1, N_1}$;
- every other 1-component, and all other components of all other colour classes of c' are double-rays;
- c' is standard outside of $P + \langle g_1, g_2 \rangle_{N_3, N_1}$ and inside of $P + (\langle g_1, g_2 \rangle_{N_2, N_1} \setminus \langle g_1, g_2 \rangle_{N_0, N_0})$;
- for each $x_\ell \in P$, the set of vertices

$$\{x_\ell + ng_1 + mg_2 : N_1 \leq |n| \leq N_2, m \in \{N_1, N_1 - 1\}\}$$

is contained in a single 1-component of c' .

PROOF. For $\ell \in [t]$ and $q \in [N_1]$ let $R_q^\ell = \blacksquare(v_q^\ell, g_1, g_2)$ and $L_q^\ell = \blacksquare(w_q^\ell, g_1, g_2)$ be the $(1, 2)$ -squares with base point $v_q^\ell = x_\ell + (N_3 + 1 - 2q) \cdot g_1 + (N_1 + 1 - 2q) \cdot g_2$ and $w_q^\ell = x_\ell - (N_3 + 2 - 2q) \cdot g_1 + (N_1 + 1 - 2q) \cdot g_2$ respectively. The square L_q^ℓ is the mirror image of R_q^ℓ with respect to the y -axis of the grid $x_\ell + \langle g_1, g_2 \rangle$, however the base points are not mirror images, accounting for the slight asymmetry in the definitions.

Since $N_3 \geq N_2 + 2N_1$, it follows that

$$R_q^\ell \cup L_q^\ell \subseteq E(x_\ell + (\langle g_1, g_2 \rangle_{N_3, N_1} \setminus \langle g_1, g_2 \rangle_{N_2, N_1}))$$

for all $q \in [N_1]$, and so by assumption on c , all R_q^ℓ and L_q^ℓ are indeed standard $(1, 2)$ -squares. We perform colour switchings on R_q^ℓ and L_q^ℓ for all $\ell \in [t]$ and $q \in [N_1]$, and call the resulting edge colouring c' . It is clear that $c' = c$ on $E(G[X])$ and that c' is standard outside of $P + \langle g_1, g_2 \rangle_{N_3, N_1}$ and inside of $P + (\langle g_1, g_2 \rangle_{N_2, N_1} \setminus \langle g_1, g_2 \rangle_{N_0, N_0})$. Let $C \subseteq G$




FIGURE 8.1. Performing colour switchings of standard squares at positions indicated by ‘x’ in a copy $x_\ell + \langle g_1, g_2 \rangle_{N_3, N_1}$ of a finite grid.

denote the region consisting of all vertices that lie in $x_\ell + \langle g_1, g_2 \rangle_{N_3, N_1}$ for some ℓ between a pair L_q^ℓ and R_q^ℓ for some q , i.e.

$$C = \bigcup_{\ell=1}^t \bigcup_{q=1}^{N_1} \bigcup_{m=1}^2 \{x_\ell + ng_1 + (N_1 + m - 2q)g_2 : |n| \leq N_3 + 1 - 2q\}.$$

Then $P + \langle g_1, g_2 \rangle_{N_2, N_1} \subseteq C$. By construction, there are no edges of colour 1 in c' leaving C , that is, $E(C, V(G) \setminus C) \cap c'^{-1}(1) = \emptyset$. In particular, since the 1-subgraph of G under c' remains 2-regular and spanning, as remarked above, all 1-components under c' inside C are finite cycles, whose union covers C .

Also, since each 1-component of c is a double-ray, it must leave the finite set $P + \langle g_1, g_2 \rangle_{N_3, N_1}$ and hence meets some R_q^ℓ or L_q^ℓ . Therefore, by construction each 1-component of c' inside C meets some R_q^ℓ or L_q^ℓ and so, since c' is standard outside of $P + \langle g_1, g_2 \rangle_{N_0, N_0}$ except at the squares R_q^ℓ or L_q^ℓ , each such 1-component meets both $P + (\langle g_1, g_2 \rangle_{N_3, N_1} \setminus \langle g_1, g_2 \rangle_{N_2, N_1})$ and $P + \langle g_1, g_2 \rangle_{N_1, N_1}$.

Moreover, all other colour components remain double-rays. This is clear for all k -components of G if $k \neq 1, 2$ (as the colour switchings of $(1, 2)$ -standard squares did not affect these other colours). However, it is also clear for the 1-coloured double-rays outside of C and also for all 2-coloured components, as we chose our standard squares R_q^ℓ and L_q^ℓ ‘staggered’, so as not to create any finite monochromatic cycles, see Figure 8.1 (recall that every $x_\ell + \Delta$ is isomorphic to the grid).

Finally, since $N_1 > N_0$, the edge set

$$\begin{aligned} & \{(x_\ell + ng_1 + N_1g_2, x_\ell + (n+1)g_1 + N_1g_2) : -N_3 \leq |n| < N_3 - 1\} \\ & \cup \{(v_1^\ell, v_1^\ell + g_2), ((w_1^\ell + g_1, w_1^\ell + g_1 + g_2))\} \\ & \cup \{(x_\ell + ng_1 + (N_1 - 1)g_2, x_\ell + (n+1)g_1 + (N_1 - 1)g_2) : -N_3 \leq n < -N_1\} \\ & \cup \{(x_\ell + ng_1 + (N_1 - 1)g_2, x_\ell + (n+1)g_1 + (N_1 - 1)g_2) : N_1 \leq n < N_3\} \end{aligned}$$

meets only R_1^ℓ and L_1^ℓ and therefore is easily seen to be part of the same 1-component of c' . In Figure 8.1, these edges correspond to the red line at the top, and the two lines below it on either side of $x_\ell + \langle g_1, g_2 \rangle_{N_1, N_1}$. \square

8.4.6. Combining cycles inside each coset of Δ . In the previous step we chose the $(1, 2)$ -standard squares at which we performed colour switchings in a staggered manner in the grids $x_l + \langle g_1, g_2 \rangle_{N_3, N_1}$, so that we could guarantee that all the 2-components were still double-rays afterwards. In later steps we will no longer be able to be as explicit about which standard squares we perform colour switchings at, and so we will require the following definitions to be able to say when it is ‘safe’ to perform a colour switching at a standard square.

DEFINITION 8.4.5 (Crossing edges). Suppose $R = \{(v_i, v_{i+1}) : i \in \mathbb{Z}\}$ is a double-ray and $e_1 = (v_{j_1}, v_{j_2})$ and $e_2 = (v_{k_1}, v_{k_2})$ are edges with $j_1 < j_2$ and $k_1 < k_2$. We say that e_1 and e_2 *cross* on R if either $j_1 < k_1 < j_2 < k_2$ or $k_1 < j_1 < k_2 < j_2$.

LEMMA 8.4.6. *For an edge-colouring $c : E(G(\Gamma, S)) \rightarrow [s]$, suppose that $\blacksquare(x, g_i, g_k)$ is an (i, k) -standard square with $g_i \neq -g_k$, and further that the two k -coloured edges*

$(x, x + g_k)$ and $(x + g_i, x + g_i + g_k)$ of $\blacksquare(x, g_i, g_k)$ lie on the same standard k -double-ray $R = \leftrightarrow(x, g_k)$. Then the two i -coloured edges of $\blacksquare(x, g_i, g_k)$ cross on R .

PROOF. Write $e_1 = (x, x + g_i)$ and $e_2 = (x + g_k, x + g_k + g_i)$ for the two i -coloured edges of $\blacksquare(x, g_i, g_k)$. The assumption that $(x, x + g_k)$ and $(x + g_i, x + g_i + g_k)$ both lie on $\leftrightarrow(x, g_k)$ implies that $g_i = rg_k$ for some $r \in \mathbb{Z} \setminus \{-1, 0, 1\}$. If $r > 1$, we have $x < x + g_k < x + g_i < x + g_k + g_i$ (where $<$ denotes the natural linear order on the vertex set of the double-ray), and if $r < -1$, we have $x + g_i < x + g_k + g_i < x < x + g_k$, and so the edges e_1 and e_2 indeed cross on R . \square

DEFINITION 8.4.7 (Safe standard square). Given an edge colouring $c: E(G(\Gamma, S)) \rightarrow [s]$ we say an (i, k) -standard square $T = \blacksquare(x, g_i, g_k)$ is *safe* if $g_i \neq -g_k$ and either

- the k -components for c meeting T are distinct double-rays, or
- there is a unique k -component for c meeting T , which is a double-ray on which $(x, x + g_i)$ and $(x + g_k, x + g_i + g_k)$ cross.

The following lemma tells us, amongst other things, that if we perform a colour switching at a safe $(1, k)$ -standard square then the k -components in the resulting colouring meeting that square will still be double-rays.

LEMMA 8.4.8. Let $c: E(G(\Gamma, S)) \rightarrow [s]$ be an edge colouring, $T = \blacksquare(x, g_i, g_k)$ be an (i, k) -standard square with $g_i \neq -g_k$, and c' be the colouring obtained by performing a colour switching on T . Then the following statements are true:

- If there are two distinct i -components C_1 and C_2 for c meeting T which are both 2-regular, at least one of which is a finite cycle, then there is a single i -component for c' meeting T which is 2-regular and whose vertex set is $V(C_1) \cup V(C_2)$;
- If the k -components for c meeting T are distinct double-rays then the k -components for c' meeting T are distinct double-rays;
- If there is a unique k -component for c meeting T , which is a double-ray on which $(x, x + g_i)$ and $(x + g_k, x + g_i + g_k)$ cross, then there is unique k -component for c' meeting T , which is a double-ray.

PROOF. Let us write $e_i = (x, x + g_i)$, $e_k = (x, x + g_k)$, $e'_i = (x + g_k, x + g_i + g_k)$ and $e'_k = (x + g_i, x + g_i + g_k)$, so that $\blacksquare(x, g_i, g_j) = \{e_i, e_k, e'_i, e'_k\}$.

For the first item, let the i -components for c be $e_i \in C_1$ and $e'_i \in C_2$, where without loss of generality C_2 is a finite cycle. Then $C_2 - e'_i$ is a finite path, and $C_1 - e_i$ has at most 2 components, one containing x and one containing $x + g_i$. Hence, the i -component for c' meeting T , $(C_1 \cup C_2) - \{e_i, e'_i\} + \{e_k, e'_k\}$, is connected and has vertex set $V(C_1) \cup V(C_2)$.

For the second item, let the k -components for c be $e_k \in D_1$ and $e'_k \in D_2$. Then $D_1 - e_k$ has two components, a ray starting at x and a ray starting at $x + g_k$. Similarly, $D_2 - e'_k$ has two components, a ray starting at $x + g_i$ and a ray starting at $x + g_i + g_k$. Hence, the




FIGURE 8.2. The two situations in Lemma 8.4.8 with i in red and k in blue.

k -components for c' meeting T , which are the components of $(D_1 \cup D_2) - \{e_k, e'_k\} + \{e_i, e'_i\}$, are distinct double-rays.

Finally, if there is a single k -component D for c meeting T such that D is a double-ray, then $D - \{e_k, e'_k\}$ consist of three components. Since e_i and e'_i cross on D there are two cases as to what these components are. Either the components consist of two rays, starting at x and $x + g_i + g_k$ and a finite path from $x + g_k$ to $x + g_i$, or the components consist of two rays, starting at $x + g_i$ and $x + g_k$, and a finite path from $x + g_i + g_k$ to x . In either case, the k -component for c' meeting T , namely $D - \{e_k, e'_k\} + \{e_i, e'_i\}$, is a double-ray. \square

Lemma 8.4.8 is also useful as the first item allows us to use $(1, k)$ colour switchings to combine two 1-components into a single 1-component which covers the same vertex set.

STEP 2 (Combining cycles step). We can change c' from Step 1 via colour switchings of finitely many $(1, 2)$ -standard squares to a colouring c'' satisfying

- $c'' = c' = c$ on $E(G[X])$;
- every 1-component in c'' meeting $P + \langle g_1, g_2 \rangle_{N_2, N_1}$ is a finite cycle intersecting both $P + (\langle g_1, g_2 \rangle_{N_3, N_1} \setminus \langle g_1, g_2 \rangle_{N_2, N_1})$ and $P + \langle g_1, g_2 \rangle_{N_1, N_1}$;
- every other 1-component, and all other components of all other colour classes of c'' are double-rays;
- every 1-component in c'' meeting some $x_k + (\langle g_1, g_2 \rangle_{N_2, N_1} \setminus \langle g_1, g_2 \rangle_{N_0, N_0})$ covers $x_k + (\langle g_1, g_2 \rangle_{N_2, N_1} \setminus \langle g_1, g_2 \rangle_{N_0, N_0})$;
- c'' is standard outside of $P + \langle g_1, g_2 \rangle_{N_3, N_1}$ and inside of $P + (\langle g_1, g_2 \rangle_{N_1, N_1} \setminus \langle g_1, g_2 \rangle_{N_0, N_0})$.

PROOF. Our plan will be to go through the ‘grids’ $x_k + \langle g_1, g_2 \rangle_{N_2, N_1}$ in order, from $k = 0$ to t , and use colour switchings to combine all the 1-components which meet $x_k + (\langle g_1, g_2 \rangle_{N_2, N_1} \setminus \langle g_1, g_2 \rangle_{N_0, N_0})$ into a single 1-component. We note that, since c' is not standard on X , it may be the case that these 1-components also meet $x_{k'} + \langle g_1, g_2 \rangle_{N_2, N_1}$ for $k' \neq k$.

We claim inductively that there exists a sequence of colourings $c' = c_0, c_1, \dots, c_t = c''$ such that for each $0 \leq \ell \leq t$:

- $c_\ell = c' = c$ on $E(G[X])$;

- every 1-component in c_ℓ meeting $P + \langle g_1, g_2 \rangle_{N_2, N_1}$ is a finite cycle intersecting both $P + (\langle g_1, g_2 \rangle_{N_3, N_1} \setminus \langle g_1, g_2 \rangle_{N_2, N_1})$ and $P + \langle g_1, g_2 \rangle_{N_1, N_1}$;
- for every $k \leq \ell$, every 1-component in c_ℓ meeting $x_k + (\langle g_1, g_2 \rangle_{N_2, N_1} \setminus \langle g_1, g_2 \rangle_{N_0, N_0})$ covers $x_k + (\langle g_1, g_2 \rangle_{N_2, N_1} \setminus \langle g_1, g_2 \rangle_{N_0, N_0})$;
- for every $k > \ell$, $c_\ell = c'$ on $x_k + \langle g_1, g_2 \rangle_{N_2, N_1}$
- every other 1-component, and all other components of all other colour classes of c_ℓ are double-rays;
- c_ℓ is standard outside of $P + \langle g_1, g_2 \rangle_{N_3, N_1}$ and inside of $P + (\langle g_1, g_2 \rangle_{N_1, N_1} \setminus \langle g_1, g_2 \rangle_{N_0, N_0})$.

In Step 1 we constructed $c_0 = c'$ such that this holds. Suppose that $0 < \ell \leq t$, and that we have already constructed c_k for $k < \ell$.

For $q \in [4N_1 - 2]$ we define $T_q = \blacksquare(v_q, g_1, g_2)$ to be the $(1, 2)$ -square with base point

$$v_q = \begin{cases} x_\ell + (N_2 + 2 - 2q)g_1 + (N_1 - q)g_2 & \text{if } q \leq 2N_1 - 1, \text{ and} \\ x_\ell - (N_2 + 3 - 2q')g_1 + (N_1 - q')g_2 & \text{if } q' = q - (2N_1 - 1) \geq 1. \end{cases}$$

With these definitions, T_{2N_1-1+q} is the mirror image of T_q for all $q \in [2N_1 - 1]$ along the y -axis. Moreover, since $N_2 \geq 5N_1$, each T_q is contained within $x_k + (\langle g_1, g_2 \rangle_{N_2, N_1} \setminus \langle g_1, g_2 \rangle_{N_1, N_1})$, see Figure 8.3.

We will combine the 1-components in $c_{\ell-1}$ which meet $x_\ell + (\langle g_1, g_2 \rangle_{N_2, N_1} \setminus \langle g_1, g_2 \rangle_{N_0, N_0})$ into a single component by performing colour switchings at some of the $(1, 2)$ -squares T_q . Let us show first that most of the induction hypotheses are maintained regardless of the subset of the T_q we make switchings at.




FIGURE 8.3. The standard squares T_q , with a colour switching performed at T_2 .

We note that, since $c_{\ell-1}$ is standard inside of $x_\ell + (\langle g_1, g_2 \rangle_{N_2, N_1} \setminus \langle g_1, g_2 \rangle_{N_0, N_0})$ and outside of $P + \langle g_1, g_2 \rangle_{N_3, N_1}$, and $g_1 \neq -g_2$, each T_q is a safe $(1, 2)$ -standard square for $c_{\ell-1}$. Furthermore, by construction, even if we perform colour switchings at any subset of the T_q , the remaining squares remain standard and safe.

Hence, by Lemma 8.4.8 and the induction assumption, after performing colour switchings at any subset of the standard squares T_q all 2-components of the resulting colouring will be double-rays. Secondly, these colour switchings will not change the colouring outside of $P + \langle g_1, g_2 \rangle_{N_2, N_1}$ and inside of $P + \langle g_1, g_2 \rangle_{N_1, N_1}$, or in any $x_k + \langle g_1, g_2 \rangle_{N_2, N_1}$ with

$k \neq \ell$. In particular, every 1-component not meeting $P + \langle g_1, g_2 \rangle_{N_2, N_1}$ will still be a double-ray. Finally, again by Lemma 8.4.8, every 1-component of the resulting colouring meeting $P + \langle g_1, g_2 \rangle_{N_2, N_1}$ will be a finite cycle which covers the vertex set of some union of 1-components in $c_{\ell-1}$, and hence will intersect both $P + (\langle g_1, g_2 \rangle_{N_3, N_1} \setminus \langle g_1, g_2 \rangle_{N_2, N_1})$ and $P + \langle g_1, g_2 \rangle_{N_1, N_1}$.

Let us write $e_q = (v_q, v_q + g_1)$ for each $q \in [4N_1 - 2]$. Since $c_{\ell-1} = c'$ on $x_\ell + \langle g_1, g_2 \rangle_{N_2, N_1}$, and by Step 1 c' is standard on $x_\ell + (\langle g_1, g_2 \rangle_{N_2, N_1} \setminus \langle g_1, g_2 \rangle_{N_0, N_0})$, each 1-component of $c_{\ell-1}$ that meets $x_\ell + (\langle g_1, g_2 \rangle_{N_2, N_1} \setminus \langle g_1, g_2 \rangle_{N_0, N_0})$ contains at least one e_q . Also, e_1 and e_{2N_1} belong to the same 1-component by the last claim in Step 1. Let us write \mathcal{C} for the collection of such cycles, and consider the map

$$\alpha: \mathcal{C} \rightarrow \{1, \dots, 4N_1 - 1\}, \quad C \mapsto \min \{q: e_q \in E(C)\},$$

which maps each cycle to the first e_q that it contains. Since \mathcal{C} is a disjoint collection of cycles, the map α is injective. Now let c_ℓ be the colouring obtained from $c_{\ell-1}$ by switching all standard squares in

$$\mathcal{T} = \{T_q: q \in \text{im}(\alpha)\} \setminus \{T_1\}.$$

We claim that c_ℓ satisfies our induction hypothesis for ℓ . By the previous comments it will be sufficient to show

CLAIM 10. Every 1-component in c_ℓ meeting $x_\ell + (\langle g_1, g_2 \rangle_{N_2, N_1} \setminus \langle g_1, g_2 \rangle_{N_0, N_0})$ covers $x_\ell + (\langle g_1, g_2 \rangle_{N_2, N_1} \setminus \langle g_1, g_2 \rangle_{N_0, N_0})$.

To see this, we index $\mathcal{C} = \{C_1, \dots, C_r\}$ such that $u < v$ implies $\alpha(C_u) < \alpha(C_v)$, and consider the sequence of colourings $\{c^z: z \in [r]\}$ where $c^1 = c_{\ell-1}$ and each c^z is obtained from c^{z-1} by switching the standard square $T_{\alpha(C_z)}$.

Let us show by induction that for every $z \in [r]$ there is a 1-component of c^z which covers $\bigcup_{y \leq z} C_y$. For $z = 1$ the claim is clearly true. So, suppose $z > 1$. Since $\alpha(C_z)$ is minimal in $\{\alpha(C_y): y \geq z\}$ it follows that $e_q \in \bigcup_{y < z} C_y$ for every $q < \alpha(C_z)$. Note that, since $c_{\ell-1} = c'$ on $x_\ell + \langle g_1, g_2 \rangle_{N_2, N_1}$, it follows from the final claim in the Cap-off step that C_1 contains both e_1 and e_{2N_1} , and so $\alpha(C_z) \neq 2N_1$.

Consider the standard square $T_{\alpha(C_z)}$. Since $c_{\ell-1} = c'$ on $x_\ell + \langle g_1, g_2 \rangle_{N_2, N_1}$, by construction the edge ‘opposite’ to $e_{\alpha(C_z)}$ in $T_{\alpha(C_z)}$, that is, $e_{\alpha(C_z)} + g_j$, is in the same 1-component in $c_{\ell-1}$ as $e_{\alpha(C_z)-1}$, and hence is contained in $\bigcup_{y < z} C_y$.

Therefore, by Lemma 8.4.8, after performing an $(1, 2)$ -colour switching at $T_{\alpha(C_z)}$, the 1-component in c^z contains $\bigcup_{y \leq z} C_y$.

Hence, there is a 1-component of $c_\ell = c^r$ which covers $\bigcup_{y \leq r} C_y$, and so there is a unique 1-component of c_ℓ meeting $x_\ell + (\langle g_1, g_2 \rangle_{N_2, N_1} \setminus \langle g_1, g_2 \rangle_{N_0, N_0})$ which covers it, establishing the claim. \square

8.4.7. Combining cycles across different cosets of Δ . In the third and final step we join the finite cycles covering each $x_\ell + (\langle g_1, g_2 \rangle_{N_1, N_1} \setminus \langle g_1, g_2 \rangle_{N_0, N_0})$ into a single finite cycle, and then make one final switch to absorb this cycle into a double-ray. The resulting colouring will then satisfy the conditions of Lemma 8.3.1.

STEP 3 (Combining cosets step). We can change c'' from the previous lemma to an almost-standard colouring \hat{c} such that

- $\hat{c} = c'' = c' = c$ on $E(G[X])$;
- Some component in colour 1 covers $P + \langle g_1, g_2 \rangle_{N_1, N_1}$.

PROOF. Recall that $P = \{x_0, \dots, x_t\}$ is such that $P^\Delta = \{x_0 + \Delta, \dots, x_t + \Delta\}$ is a finite, graph-theoretic path in the Cayley graph of the quotient Γ/Δ with generating set $(S \setminus \{g_1, g_2\})^\Delta$. Moreover, recall from Section 8.4.4 that $N_1 > N_0$ was chosen so that for the initial colouring c there were t^2 many disjoint standard double-rays

$$\mathcal{R} = \{R_\ell^k : 1 \leq k, \ell \leq t\}$$

such that for every ℓ , the double-rays in $\{R_\ell^k = \rightsquigarrow(y_\ell^k, g_{n(\ell)}) : k \in [t]\}$ are standard $n(\ell)$ -double-rays containing an edge

$$e_\ell^k = (y_\ell^k, y_\ell^k + g_{n(\ell)}) \in E(R_\ell^k) \cap E(x_{\ell-1} + \Delta, x_\ell + \Delta)$$

so that all $T_\ell^k = \blacksquare(y_\ell^k, g_1, g_{n(\ell)})$ are edge-disjoint $(1, n(\ell))$ -standard squares for the colouring c contained in the subgraph induced by $P + \langle g_1, g_2 \rangle_{N_1-3, N_1-3}$ which have empty intersection with $\{x_{\ell-1}, x_\ell\} + \langle g_1, g_2 \rangle_{N_0, N_0}$. However, since we only altered the $(1, 2)$ -subgraphs of G in Step 1 and 2, it is clear that all these standard double-rays and standard squares for c remain standard also for the colourings c' and in particular c'' .

We claim that there exists a function $k: [t] \rightarrow [t] \cup \{\perp\}$ such that iteratively switching $T_\ell^{k(\ell)}$ (or not doing anything at all if $k(\ell) = \perp$) results in a sequence of colourings $c'' = c_0, c_1, \dots, c_t$ such that for each $0 \leq \ell \leq t$,

- (1) a single finite 1-component in c_ℓ covers $\{x_0, \dots, x_\ell\} + (\langle g_1, g_2 \rangle_{N_1, N_1} \setminus \langle g_1, g_2 \rangle_{N_0, N_0})$,
- (2) for every k , every 1-component in c_ℓ meeting $x_k + (\langle g_1, g_2 \rangle_{N_1, N_1} \setminus \langle g_1, g_2 \rangle_{N_0, N_0})$ is a finite cycle covering $x_k + (\langle g_1, g_2 \rangle_{N_1, N_1} \setminus \langle g_1, g_2 \rangle_{N_0, N_0})$, and
- (3) every other 1-component, and all other components of all other colour classes in c_ℓ are double-rays.

In Step 2 we constructed a colouring $c_0 = c''$ for which properties (1)–(3) are satisfied. Now suppose that $\ell \geq 1$, and that the colouring $c_{\ell-1}$ obtained by switching the standard squares $\{T_{\ell'}^{k(\ell')} : \ell' \in [\ell-1], k(\ell') \neq \perp\}$ satisfies (1)–(3). By construction, each such standard square $T_{\ell'}^{k(\ell')}$ has an edge in common with the ray $R_{\ell'}^{k(\ell')}$ and potentially one further $n(\ell')$ -component. But since we had reserved more than $\ell-1$ different rays $R_\ell^1, \dots, R_\ell^t$, it follows that some ray R_ℓ^K remains a standard $n(\ell)$ -coloured component for $c_{\ell-1}$.




FIGURE 8.4. Using $(1, n(\ell))$ -standard squares to join up different cosets. For this picture, we assume wlog that $x_{\ell+1} = x_\ell + g_{n(\ell+1)}$.

Both edges $(y_\ell^K, y_\ell^K + g_i)$ and $(y_\ell^K + g_{n(\ell)}, y_\ell^K + g_{n(\ell)} + g_i)$ of T_ℓ^K are contained in $\{x_{\ell-1}, x_\ell\} + (\langle g_1, g_2 \rangle_{N_1, N_1} \setminus \langle g_1, g_2 \rangle_{N_0, N_0})$, and hence are, by assumption (2), covered by finite 1-cycles in $c_{\ell-1}$. If both edges lie in the same finite 1-cycle, there is nothing to do and we set $k(\ell) := \perp$, so that $c_\ell = c_{\ell-1}$. However, if they lie on different finite cycles, set $k(\ell) := K$. Then, in our procedure we perform a colour switching on the standard square $T_\ell^{k(\ell)}$ and claim that the resulting c_ℓ is as required. By Lemma 8.4.8, the two finite 1-components merge into a single finite cycle, and so (1) and (2) are certainly satisfied for c_ℓ .

To see (3), we need to verify that $T_\ell^{k(\ell)}$ is, when we perform the switching, safe. However, $T_\ell^{k(\ell)}$ was chosen so that the edge $(y_\ell^{k(\ell)}, y_\ell^{k(\ell)} + g_{n(\ell)}) \in T_\ell^{k(\ell)}$ lies on a standard double-ray $R = R_\ell^{k(\ell)}$ of $c_{\ell-1}$. Also, by the inductive assumption (3), the second $n(\ell)$ -coloured edge $(y_\ell^{k(\ell)} + g_i, y_\ell^{k(\ell)} + g_i + g_{n(\ell)}) \in T_\ell^{k(\ell)}$ lies on an $n(\ell)$ -coloured double-ray R' in $c_{\ell-1}$. If R and R' are distinct, then $T_\ell^{k(\ell)}$ is safe, and if $R = R'$ then, since R is a standard $n(\ell)$ -double-ray, Lemma 8.4.6 implies that $T_\ell^{k(\ell)}$ is safe. Hence c_ℓ satisfies (3). This completes the induction step.

Thus, by (1) and (3), we obtain an edge-colouring c_t for G such that a single finite 1-component covers $P + (\langle g_1, g_2 \rangle_{N_1, N_1} \setminus \langle g_1, g_2 \rangle_{N_0, N_0})$, and all other 1-components and all other components of other colour classes in c_t are double-rays. Furthermore, since every 1-component which meets $P + \langle g_1, g_2 \rangle_{N_0, N_0}$ must meet $P + (\langle g_1, g_2 \rangle_{N_1, N_1} \setminus \langle g_1, g_2 \rangle_{N_0, N_0})$, it follows that the 1-component in fact covers $P + \langle g_1, g_2 \rangle_{N_0, N_0}$. Moreover, since $T_\ell^{k(\ell)} \subseteq P + \langle g_1, g_2 \rangle_{N_1-3, N_1-3}$ for all $\ell \in [t]$, it follows that c_t is standard on $x_0 + (\langle g_1, g_2 \rangle_{N_1, \infty} \setminus \langle g_1, g_2 \rangle_{N_1-3, N_1-3})$, and that it is standard outside of $P + \langle g_1, g_2 \rangle_{N_3, N_1}$.

Hence, the square $\blacksquare(x, g_1, g_2)$ with base point $x = x_0 + (N_1 - 2)g_1 + N_1g_2$ is a standard $(1, 2)$ -square such that

- the edge $(x, x + g_1)$ lies on the finite 1-cycle of c_t ,
- the edge $(x + g_2, x + g_2 + g_1)$ lies on a standard 1-double-ray $\leftrightarrow(x + g_2, g_1)$ (lying completely outside of $P + \langle g_1, g_2 \rangle_{N_3, N_1}$) of c_t , and
- the edges $(x, x + g_2)$ and $(x + g_1, x + g_2 + g_1)$ lie on distinct standard 2-double-rays $\leftrightarrow(x, g_2)$ and $\leftrightarrow(x + g_1, g_2) \subseteq x_0 + (\langle g_1, g_2 \rangle_{N_1, \infty} \setminus \langle g_1, g_2 \rangle_{N_1-3, N_1-3})$.

Therefore, we may perform a colour switching on $\blacksquare(x, g_1, g_2)$, which results, by Lemma 8.4.8, in an almost-standard colouring of G such that a single 1-component covers $P + \langle g_1, g_2 \rangle_{N_1, N_1}$, and hence X . \square

8.5. Hamiltonian decompositions of products

The techniques from the previous section can also be applied to give us the following general result about Hamiltonian decompositions of products of graphs.

THEOREM 8.1.4. *If G and H are countable multi-graphs which both have Hamilton decompositions, then so does their product $G \square H$.*

PROOF. Suppose that $\{R_i : i \in I\}$ and $\{S_j : j \in J\}$ form decompositions of G and H into edge-disjoint Hamiltonian double-rays, where I, J may be finite or countably infinite. Note that, for each $i \in I, j \in J$, $R_i \square S_j$ is a spanning subgraph of $G \square H$, and is isomorphic to the Cayley graph of $(\mathbb{Z}^2, +)$ with the standard generating set.

Let $\pi_G : G \square H \rightarrow G$ and $\pi_H : G \square H \rightarrow H$ the projection maps from $G \square H$ onto the respective coordinates. As our *standard colouring* for $G \square H$ we take the map

$$c : E(G \square H) \rightarrow I \dot{\cup} J, e \mapsto \begin{cases} i & \text{if } e \in \pi_G^{-1}(E(R_i)), \\ j & \text{if } e \in \pi_H^{-1}(E(S_j)). \end{cases}$$

Then each $R_i \square S_j$ is 2-coloured (with colours i and j), and this colouring agrees with the standard colouring of $C_{\mathbb{Z}^2} = G((\mathbb{Z}^2, +), \{(1, 0), (0, 1)\})$ from Section 8.3. We also define an *almost-standard colouring* of $G \square H$ as in Definition 8.2.2.

We may suppose that $V(G) = \mathbb{N} = V(H)$. Fix a surjection $f : \mathbb{N} \rightarrow I \cup J$ such that every colour appears infinitely often.

By starting with $c_0 = c$ and applying Lemma 8.3.1 recursively inside the spanning subgraphs $R_{f(k)} \square S_1$, if $f(k) \in I$, or inside $R_1 \square S_{f(k)}$, for $f(k) \in J$, we find a sequence of almost-standard edge-colourings $c_k : G \square H \rightarrow I \cup J$ and natural numbers $M_k \leq N_k < M_{k+1}$ such that

- c_{k+1} agrees with c_k on the subgraph of $G \square H$ induced by $[0, M_{k+1}]^2$,
- there is a finite path D_k of colour $f(k)$ in c_k covering $[0, N_k]^2$, and
- M_{k+1} is large enough such that $D_k \subseteq [0, M_{k+1}]^2$.

To be precise, suppose we already have a finite path D_k of colour $f(k)$ in c_k covering $[0, N_k]^2$, and at stage $k+1$ we have say $f(k+1) \in I$, and so we are considering $R_{f(k+1)} \square S_1 \cong C_{\mathbb{Z}^2}$. We choose

- $M_{k+1} > N_k$ large enough such that $D_k \subseteq [0, M_{k+1}]^2 \subseteq G \square H$, and
- $N_{k+1} > M_{k+1}$ large enough such that $Q_1 = [0, N_{k+1}]^2 \subseteq G \square H$ contains all edges where c_k differs from the standard colouring c .

Next, consider an isomorphism $h: R_{f(k+1)} \square S_1 \cong C_{\mathbb{Z}^2}$. Pick a ‘square’ $Q_2 \subseteq R_{f(k+1)} \square S_1$ with $Q_1 \subseteq Q_2$, i.e. a set Q_2 such that h restricted to Q_2 is an isomorphism to the subgraph of $C_{\mathbb{Z}^2}$ induced by $[-\tilde{N}_{k+1}, \tilde{N}_{k+1}]^2 \subseteq \mathbb{Z}^2$ for some $\tilde{N}_{k+1} \in \mathbb{N}$, and then apply Lemma 8.3.1 to $R_{f(k+1)} \square S_1$ and Q_2 to obtain a finite path D_{k+1} of colour $f(k+1)$ in c_{k+1} covering Q_2 .

It follows that the double-rays $\{T_i: i \in I\} \cup \{T_j: j \in J\}$ with $T_\ell = \bigcup_{k \in f^{-1}(\ell)} D_k$ give the desired decomposition of $G \square H$. \square

8.6. Open Problems

As mentioned in Section 8.2, the finitely generated abelian groups can be classified as the groups $\mathbb{Z}^n \oplus \bigoplus_{i=1}^r \mathbb{Z}_{q_i}$, where $n, r, q_1, \dots, q_r \in \mathbb{Z}$. Theorem 8.1.2 shows that Alspach’s conjecture holds for every such group with $n \geq 2$, as long as each generator has infinite order. The question however remains as to what can be said about Cayley graphs $G(\Gamma, S)$ when S contains elements of finite order.

PROBLEM 8.6.1. Let Γ be an infinite, finitely-generated, one-ended abelian group and S be a generating set for Γ which contains elements of finite order. Show that $G(\Gamma, S)$ has a Hamilton decomposition.

Alspach’s conjecture has also been shown to hold when $n = 1$, $r = 0$, and the generating set S has size 2, by Bryant, Herke, Maenhaut and Webb [40]. In a paper in preparation [65], the first two authors consider the general case when $n = 1$ and the underlying Cayley graph is 4-regular. Since the Cayley graph is 2-ended, it can happen for parity reasons that no Hamilton decomposition exists. However, this is the only obstruction, and in all other cases the Cayley graphs have a Hamilton decomposition. Together with the result of Bermond, Favaron and Maheo [22] for finite abelian groups, and the case $\Gamma \cong (\mathbb{Z}^2, +)$ of Theorem 8.1.2, this fully characterises the 4-regular connected Cayley graphs of finite abelian groups which have Hamilton decompositions. A natural next step would be to consider the case of 6-regular Cayley graphs.

PROBLEM 8.6.2. Let Γ be a finitely generated abelian group and let S be a generating set of Γ such that $C(\Gamma, S)$ is 6-regular. Characterise the pairs (Γ, S) such that $G(\Gamma, S)$ has a decomposition into spanning double-rays.

Hamilton cycles in infinite cubic graphs

Investigating a problem of B. Mohar, we show that every one-ended Hamiltonian cubic graph with end degree 3 contains a second Hamilton cycle. We also construct two examples showing that this result does not extend to give a third Hamilton cycle, nor that it extends to the two-ended case.

9.1. Introduction

In this note we investigate whether results about the Hamiltonicity of finite cubic graphs extend to the infinite setting. The term ‘graph’ in this paper is reserved for simple graphs; when allowing parallel edges or loops, we explicitly use the term ‘multi-graph’. Our terminology follows [54].

9.1.1. Hamiltonicity in finite regular graphs. The starting point of this paper are the following results and conjectures for finite regular graphs.

THEOREM 9.1.1 (Smith ’46, see [156]). *Every Hamiltonian finite cubic graph has at least two Hamilton cycles.*

Here, a graph is *Hamiltonian* if it contains a Hamilton cycle. A graph with precisely one Hamilton cycle is also called *uniquely Hamiltonian*. Sheehan conjectured that finite cycles are the only examples of uniquely Hamiltonian regular graphs.

CONJECTURE 9.1.2 (Sheehan ’75, [142]). *Every d -regular Hamiltonian finite graph with $d \geq 3$ has at least two Hamilton cycles.*

For more details on Sheehan’s conjecture, we refer the reader to [152]. Using a nice parity argument, the so-called “lollipop technique”, Thomason extended Smith’s result in a different direction as follows:

THEOREM 9.1.3 (Thomason ’78, [150]). *Every edge in a finite graph with odd degrees only lies on an even number of Hamilton cycles. Hence, every Hamiltonian such graph has at least three Hamilton cycles.*

In particular, every finite Hamiltonian cubic graph contains at least three Hamilton cycles.

9.1.2. Infinite Hamilton circles. For a locally finite graph G , which can be considered as a topological space using the 1-complex topology, we let $|G|$ denote its Freudenthal compactification. Extending the notion of cycles, one defines *circles* in $|G|$ as homeomorphic images of the unit circle in $|G|$, see [54, §8]. A circle of $|G|$ containing all vertices (and all ends) of G is a *Hamilton circle*. A *Hamilton cycle* is a subgraph of G given by the edge set of a Hamilton circle of $|G|$.

In one-ended graphs, Hamilton cycles correspond to spanning double rays. In a two-ended graph G , a Hamilton cycle consists of two vertex-disjoint double rays R_1 and R_2 which together span G , such that the two tails of each R_i belong to different ends of G . For example, the 2-ended double ladder in Figure 9.1 has a unique Hamilton cycle comprised of all horizontal edges.




FIGURE 9.1. The infinite double ladder and its unique Hamilton cycle.

9.1.3. Questions on Hamiltonicity in infinite regular graphs. In 2007, Mohar asked to what extent the above results about Hamiltonicity in finite regular graphs generalise to the infinite setting. While the infinite double ladder in Figure 9.1 witnesses that Theorem 9.1.1 fails to extend verbatim to the infinite case, Mohar suggested two possible solutions.

First, we might restrict our attention to one-ended graphs, and second, we might say that the double ladder is not truly regular, as its ends have degree 2. Here, we take the *degree of an end* to be the maximum number of edge-disjoint rays leading to that end, see [39] or Section 9.2 below for details.

QUESTION 9.1.4 (Mohar '07, [119]). Does there exist a uniquely Hamiltonian, one-ended, d -regular graph for $d \geq 3$?

QUESTION 9.1.5 (Mohar '07, [119]). Does there exist a uniquely Hamiltonian, d -regular graph for $d \geq 3$ where also all ends have degree d ?

K. Heuer [95] has recently constructed a uniquely Hamiltonian cubic graph with continuum many ends where all ends have degree 3, thus answering Question 9.1.5. He left open the natural question whether simultaneously restricting the number of ends plus the end-degrees allows us to extend finite theorems to the infinite setting.

9.1.4. Results. In this note, we establish the following extension of Smith's Theorem 9.1.1 about second Hamilton cycles to the infinite setting, providing a partial answer to Mohar's questions.

THEOREM 9.1.6. *Every Hamiltonian one-ended cubic graph with end degree at most 3 has at least two Hamilton cycles.*

Our proof of Theorem 9.1.6 combines the stronger of the finite results, namely Thomason's Theorem 9.1.3, and a sequence of parity arguments. Interestingly, Thomason's Theorem 9.1.3 itself does *not* extend to the above setting: we construct one-ended cubic graphs with end-degree 2 or 3 that have precisely two Hamilton cycles, see Examples 9.4.1 and 9.4.3.

Improving on Heuer's example, we also construct in Example 9.4.4 a two-ended, uniquely Hamiltonian, cubic graph where both ends have degree 3. This shows that in general, it is only in the one-ended case where one could hope for an affirmative result about second Hamilton circles in cubic graphs.

Finally, we remark that we do not know whether every Hamiltonian one-ended cubic graph with end-degree 4 has a second Hamilton cycle—this seems to be the next crucial case in attacking Question 9.1.4.

9.2. Two facts about end degrees

In our proofs below we need two facts about end degrees in locally finite graphs. Given a graph $G = (V, E)$ and a set of vertices $S \subseteq V$, we denote by $E(S, V \setminus S) \subseteq E$ the set of edges of G with one endvertex in S and the other in the complement of S . We also abbreviate $E(v) = E(\{v\}, V \setminus \{v\})$.

Following [39], for an end ω of some locally finite graph G we take its *degree* (to be precise: its *edge-degree*) to be the maximum number of edge disjoint rays in G leading to ω , and its *vertex-degree* to be the maximum number of vertex-disjoint rays in G leading to ω .

LEMMA 9.2.1 ([39, Lemma 10]). *Let ω be an end of a locally finite graph G and $S \subseteq V(G)$ a finite vertex set. Then the maximal number of edge-disjoint rays to ω starting in S equals the minimum cardinality of an edge cut separating S from ω .*

LEMMA 9.2.2. *In cubic graphs, edge- and vertex-degree of ends coincide.*

PROOF. In any locally finite graph, the vertex-degree of a given end is at most its edge-degree. Conversely, any family $\{R_i : i \in I\}$ of edge disjoint rays in a cubic graph have to be internally vertex-disjoint, as otherwise there would be a vertex of degree ≥ 4 . Thus, if R'_i denotes the ray R_i minus its initial vertex, then $\{R'_i : i \in I\}$ is a family of vertex-disjoint rays of the same cardinality as our initial family. \square

9.3. Affirmative results for second Hamilton cycles

In this section, we present our positive results about the existence of additional Hamilton cycles in one-ended cubic graph with end-degree 2 or 3.

THEOREM 9.3.1. *Every Hamiltonian one-ended cubic graph with end-degree 2 has at least two Hamilton cycles.*

PROOF. Let C be a Hamilton cycle of G . Since the end of G has degree 2, by Lemma 9.2.1 there is a finite vertex set $S \subseteq V$ with $|E(S, V \setminus S)| = 2$.

Consider the minor \hat{G} of G where we contract $V \setminus S$ to a single ‘dummy’ vertex. Then $C \upharpoonright \hat{G}$ witnesses that \hat{G} is a finite Hamiltonian graph. Moreover, \hat{G} is *nearly-cubic*, that is all vertices of \hat{G} have degree 3, with the exception of the dummy vertex, which has degree 2. By [64, Theorem 1], every nearly cubic Hamiltonian graph has a second Hamilton cycle. By combining the two Hamilton cycles of \hat{G} with $C \setminus E(\hat{G})$, we have found two distinct Hamilton cycles of G . □




FIGURE 9.2. The Tutte fragment T .

For the end-degree 3 case, we employ the following auxiliary multi-graph which encodes how Hamilton cycles choose incident edges of certain vertices of a graph.

DEFINITION 9.3.2 (Hamilton incidence multi-graph). Let v and w be distinct vertices of a Hamiltonian graph G . The *Hamilton incidence multi-graph* $H = H(G, v, w)$ of G with respect to v and w is the bipartite multi-graph with bipartition

$$V(H) = [E(v)]^2 \sqcup [E(w)]^2$$

where the multiplicity of an edge $pq \in E(H)$ corresponds to the number of Hamilton cycles D of G with $p \cup q \subseteq D$.

As a concrete example of a Hamilton incidence multi-graph (which we shall meet again in Section 9.4 below), consider the Tutte fragment T (invented by Tutte in [156]) with leaves l_x, l_y and l_z as depicted in Figure 9.2.




FIGURE 9.3. The six Hamilton cycles of T' .

Let $T' = T/\{\ell_x = \ell_y = \ell_z\}$ be the graph obtained from T by identifying its three leaves. Then T' is a cubic graph with precisely six Hamilton cycles (see [48, 95, 156]), which we may visualise as follows:

The first two Hamilton cycles use the edge pair $e_x = \ell_x x$ and $e_z = \ell_z z$, and the other four Hamilton cycles use the edge pair $e_y = \ell_y y$ and e_z . In particular, there are no Hamilton cycles of T' using the edge pair $\{e_x, e_y\}$. Writing w for the contracted vertex $\{\ell_x = \ell_y = \ell_z\}$ in T' , and letting v and its incident edges f_a, f_b and f_c be as indicated in Figure 9.2, we see that the Hamilton incidence graph $H = H(T', w, v)$ as in Definition 9.3.2 is given by the multigraph in Figure 9.4.




FIGURE 9.4. The Hamilton incidence multi-graph $H(T', w, v)$.

Note that all vertices of our example $H(T', w, v)$ have even degree. In the following two lemmas, we show that this parity condition holds in general.

LEMMA 9.3.3. *Let v and w be distinct vertices of a finite cubic graph G . Then the sum of the degrees of any pair of vertices in the Hamilton incidence multi-graph $H(G, v, w)$ from the same side of its vertex bipartition is always even.*

PROOF. Indeed, if say $p \neq q \in [E(v)]^2$, we have $p \cap q = \{e\}$ for some edge $e \in E(v)$, as G is cubic. So the sum of degrees $d(p) + d(q)$ equals the number of Hamilton cycles in G using the edge e , which is even by Theorem 9.1.3. \square

LEMMA 9.3.4. *If v and w are distinct vertices of a finite cubic graph G , then all vertex degrees in $H(G, v, w)$ are of the same parity.*

PROOF. Suppose one vertex in $[E(v)]^2$ has odd (even) degree. Since $|[E(v)]^2| = 3$, we can apply Lemma 9.3.3 twice to conclude that all degrees on the $[E(v)]^2$ side of our

bipartite graph $H = H(G, v, w)$ are odd (even). Hence,

$$\sum_{p \in [E(v)]^2} d_H(p) = |E(H)| = \sum_{p \in [E(w)]^2} d_H(p)$$

is odd (even). Applying Lemma 9.3.3 twice again, we see that also all degrees on the $[E(w)]^2$ side of our bipartite graph H must be odd (even). Thus, all vertex degrees in $H(G, v, w)$ are of the same parity. \square

THEOREM 9.3.5. *Every Hamiltonian one-ended cubic graph with end-degree 3 has at least two Hamilton cycles.*

PROOF. Let C be a Hamilton cycle (i.e. a spanning double ray) of G . By assumption on the degree of our end together with Lemma 9.2.1, there is a sequence of pairwise disjoint edge cuts $F_n = E(S_n, V \setminus S_n)$ with S_n finite, $|F_n| = 3$, $S_n \subsetneq S_{n+1}$, and $\bigcup_{n \in \mathbb{N}} S_n = V(G)$.

Let $F_n = \{e_n, f_n, g_n\}$. As every double ray in a one-ended locally finite graph intersects each finite cut in a positive, even number of edges, we may suppose that $e_n, f_n \in E(C)$ and $g_n \notin E(C)$ for all $n \in \mathbb{N}$. Let G_n be the minor of G where we identify $V \setminus S_n$ to a single dummy vertex x_n , and let $G_{n,n+1}$ be the minor of G where we identify S_n and $V \setminus S_{n+1}$ to dummy vertices v_n and w_n respectively.

While a priori, G_n and $G_{n,n+1}$ are multi-graphs (with possibly parallel edges at dummy vertices), we may assume they are simple: By Lemma 9.2.2, there are three vertex-disjoint rays R_1, R_2 and R_3 leading to the single end ω . Choose $N \in \mathbb{N}$ such that $E(R_i) \cap F_n \neq \emptyset$ for all $n \geq N$ and all i . Since the R_i are vertex-disjoint, it follows that all x_n, v_n and w_n have three distinct neighbours for all $n \geq N$.

So by moving to a suitable subsequence, we may assume that all our minors G_n and $G_{n,n+1}$ are simple cubic graphs. Moreover, in all cases, the corresponding restriction of C witnesses that these minors are in fact Hamiltonian.

Now, if some G_n has two distinct Hamilton cycles both using the edge set $\{e_n, f_n\}$, then, following the same strategy as in Theorem 9.3.1, we may combine both with $C \upharpoonright (V \setminus S_n)$ to obtain two distinct Hamilton cycles of G . Hence, we may assume for the remainder of the proof that for all $n \in \mathbb{N}$, the restriction $C \upharpoonright G_n$ is the only Hamilton cycle of G_n that uses $\{e_n, f_n\}$. In particular, we are in the case where the assumptions of the following claim are satisfied for all $n \in \mathbb{N}$:

CLAIM. *If G_n and G_{n+1} have unique Hamilton cycles using the edge set $\{e_n, f_n\}$ and $\{e_{n+1}, f_{n+1}\}$ respectively, then every Hamilton cycle of G_n extends to a Hamilton cycle of G_{n+1} .*

To see why the claim implies the theorem, note that by Theorem 9.1.3, the edge e_0 is contained in an even number of Hamilton cycles of G_0 , and hence there must be a second Hamilton cycle A_0 of G_0 which uses the edge set say $\{e_0, g_0\}$. Applying the claim recursively, we obtain a sequence of Hamilton cycles A_n of G_n such that A_{n+1} extends

A_n for all $n \in \mathbb{N}$. Then $A = \bigcup_{n \in \mathbb{N}} A_n$ a Hamilton cycle of G , which is distinct from C witnessed by $g_0 \in E(A) \setminus E(C)$.

It remains to prove the claim. Assume that G_n and G_{n+1} have unique Hamilton cycles using the edge sets $\{e_n, f_n\}$ and $\{e_{n+1}, f_{n+1}\}$ respectively, and consider the Hamilton incidence graph $H_n = H(G_{n,n+1}, v_n, w_n)$ of $G_{n,n+1}$ with respect to its two dummy vertices.

STEP 4. We have $d_{H_n}(\{e_{n+1}, f_{n+1}\}) = 1$.

This is where we use the assumption that G_n and G_{n+1} have unique Hamilton cycles using the edge sets $\{e_n, f_n\}$ and $\{e_{n+1}, f_{n+1}\}$ respectively. Indeed, note first that $C \upharpoonright G_{n,n+1}$ witnesses that $d_{H_n}(\{e_{n+1}, f_{n+1}\}) \geq 1$. Next, since there is a unique Hamilton cycle A of G_n that uses $\{e_n, f_n\}$, Theorem 9.1.3 implies that G_n must have two further Hamilton cycles B and C using the edge sets $\{e_n, g_n\}$ and $\{f_n, g_n\}$ respectively. Thus, if $d_{H_n}(\{e_{n+1}, f_{n+1}\}) \geq 2$, i.e. if there are two distinct Hamilton cycles of $G_{n,n+1}$ using the edge set $\{e_{n+1}, f_{n+1}\}$, then we can combine them suitably with either A , B or C to obtain two distinct Hamilton cycles of G_{n+1} both using the edge set $\{e_{n+1}, f_{n+1}\}$, a contradiction.

STEP 5. Every vertex of H_n has odd degree.

Since Step 4 implies in particular that $d_{H_n}(\{e_{n+1}, f_{n+1}\})$ is odd, Step 5 is immediate from Lemma 9.3.4.

STEP 6. Every Hamilton cycle of G_n extends to a Hamilton cycle of G_{n+1} .

Suppose we have a Hamilton cycle A of G_n using the edge set $p \in [F_n]^2$. By Step 5, we know that in particular $d_{H_n}(p) \geq 1$, which means there is a Hamilton cycle B of $G_{n,n+1}$ using the edge set p . Then $A \cup B$ is a Hamilton cycle of G_{n+1} extending A . This completes the proof of the final step of the claim, and so the theorem follows. \square

9.4. Examples witnessing optimality

In the previous section, we have seen that Smith's Theorem 9.1.1 extends to the one-ended cubic case where the end has degree at most 3. In this section, we show that Theorem 9.1.1 does not extend to the two-ended case, and that Thomason's Theorem 9.1.3 does not extend to the infinite case at all.

9.4.1. Ends with degree two.

EXAMPLE 9.4.1. There is a one-ended cubic graph with end degree 2 that has precisely two Hamilton cycles. In particular, there are edges which do not lie on an even number of Hamilton circles.

CONSTRUCTION. Consider the cubic, one-way infinite ladder as in Figure 9.5. Clearly, it has precisely one end, which has degree 2. Moreover, it is not hard to check that this graph has precisely two Hamilton cycles. In particular, there are edges which do not lie




FIGURE 9.5. The infinite cubic ladder.

on an even number of Hamilton circles. In our example, these are the edges e_1, e_2, f_1 and f_2 . □

For completeness, we record again:

EXAMPLE 9.4.2. The double ladder is a uniquely Hamiltonian, two-ended cubic graph with both ends of degree 2.

9.4.2. Ends with degree three.

EXAMPLE 9.4.3. There is a one-ended cubic graph with end degree 3 that has precisely two Hamilton cycles. In particular, there are edges which do not lie on an even number of Hamilton circles.

CONSTRUCTION. Let $\{T_n : n \in \mathbb{N}\}$ be a family of disjoint graphs such that $T_0 \cong T'$ and $T_n \cong T$ for all $n \geq 1$. Here, T is the Tutte fragment from Figure 9.2, and T' is its cubic quotient. We use the same of vertices in T and T' as above, and by $v_n, a_n, b_n, c_n \in T_n$ etc. we refer to the respective copies of the vertices $v, a, b, c \in T$.

We now construct a sequence $\{G_n : n \in \mathbb{N}\}$ of finite cubic graphs as follows: Put $G_0 = T_0$, and define

$$G_1 = (G_0 - v_0 \sqcup T_1) / \sim \text{ where } a_0 \sim l_{x_1}, b_0 \sim l_{y_1}, c_0 \sim l_{z_1}.$$

We think of this operation as replacing the vertex v_0 and its incident edges by a new copy of T , where the leaves of the new T are suitably identified with the old neighbours of v_0 . Similarly, assuming G_n has already been defined, let

$$G_{n+1} = (G_n - v_n \sqcup T_{n+1}) / \sim \text{ where } a_n \sim l_{x_{n+1}}, b_n \sim l_{y_{n+1}}, c_n \sim l_{z_{n+1}}.$$

In other words, in every step, we replace the most recent copy of the vertex v by a new copy of T .

Note that $G_n - v_n \subseteq G_{n+1} - v_{n+1}$ for all n , so we may denote by G be the direct limit of these graphs. (Alternatively, $|G|$ can be viewed as the inverse limit of the G_n under natural minor relation $G_n \preceq G_{n+1}$, cf. [54, §8.8], and so G as a 1-complex is given by the space $|G|$ minus its unique end).

Since T' is 3-edge connected, it follows that G is a one-ended cubic graph with end-degree 3. Writing $S_n = V(G_n) \setminus \{v_n\}$, we see that the end-degree of G is witnessed by the 3-edge cuts

$$F_n = E(S_n, V(G) \setminus S_n).$$




FIGURE 9.6. The incidence multi-graph for Hamilton cycles of G .

Moreover, if we define, as in the proof of Theorem 9.3.5, the graphs $G_{n,n+1}$ to be the minors of G where we identify S_n and $V(G) \setminus S_{n+1}$ to dummy vertices α_n and β_n respectively, then our construction of G guarantees the existence of isomorphisms

$$\varphi_n: T' \rightarrow G_{n,n+1} \text{ such that } \varphi_n(w) = \alpha_n \text{ and } \varphi_n(v) = \beta_n$$

such that, due to our choice of the quotient patterns \sim ,

$$(\dagger) \quad \varphi_n(f_a) = \varphi_{n+1}(e_x), \quad \varphi_n(f_b) = \varphi_{n+1}(e_y) \text{ and } \varphi_n(f_c) = \varphi_{n+1}(e_z)$$

for all $n \in \mathbb{N}$.

Next, recall that every Hamilton cycle C of G restricts, for any $n \in \mathbb{N}$, to a Hamilton cycle of $G_{n,n+1}$, and therefore looks locally like one of the six Hamilton cycles of Figure 9.3. Pasting together the individual Hamilton incidence graphs of $G_{n,n+1}$ (cf. Figure 9.4) using the identities provided in (\dagger) gives the picture of Figure 9.6. And since for every Hamilton cycle C of G we have

$$E(C \upharpoonright G_{n,n+1}) \cap E(\beta_n) = E(C \upharpoonright G_{n+1,n+2}) \cap E(\alpha_{n+1})$$

we see that Hamilton cycles of G are in 1-1 correspondence with those rays in the multi-graph in Figure 9.6 that pick a single edge from each level.




FIGURE 9.7. The incidence multi-graph for Hamilton cycles of H .

To complete the construction of Example 9.4.3, we now consider the graph

$$H = (T \sqcup G - w_0) / \sim \text{ where } \ell_x \sim z_0, \ell_y \sim y_0, \ell_z \sim x_0.$$

Figure 9.7 shows the analogue of Figure 9.6 for our new graph H .

By the same reasoning as above, Hamilton cycles of H are in 1-1 correspondence with those rays in the multi-graph in Figure 9.7 that pick a single edge from each level. But this means that H has precisely two Hamilton cycles: Only the two left-most red edges can be extended to a ray through the Hamilton incidence multi-graph using a single edge from each level, and both these extensions are unique. \square

EXAMPLE 9.4.4. There is a uniquely Hamiltonian, two-ended cubic graph with both ends of degree 3.

CONSTRUCTION. For the construction, take a disjoint copy G' of G from the graph as constructed in the previous construction (cf. Figure 9.6). By $w'_0, x'_0, y'_0, z'_0 \in G'$ etc. we refer to the respective copies of the vertices $w_0, x_0, y_0, z_0 \in G$. Now consider the graph

$$H' = (G' - w'_0 \sqcup G - w_0) \text{ with three added edges } x'_0z_0, y'_0y_0, \text{ and } z'_0x_0.$$

Then H' is a 2-ended cubic graph with both ends of degree 3. Figure 9.8 shows the analogue of Figure 9.7 for our new graph H' .




FIGURE 9.8. The incidence graph for Hamilton cycles of H' .

By the same reasoning as before, Hamilton cycles of H' correspond in a 1-1 fashion to those double rays in the multi-graph in Figure 9.8 that pick a single edge from each level. But then it is obvious that H' has a unique Hamilton cycle, which corresponds to the double ray formed by the middle horizontal edges. \square

Circuits through prescribed edges

We prove that a connected graph contains a circuit—a closed walk that repeats no edges—through any k prescribed edges if and only if it contains no odd cut of size at most k .

10.1. Introduction

Finding a cycle¹ containing certain prescribed vertices or edges of a graph is a classical problem in graph theory. When specifying vertices, already Dirac [61, Satz 9] observed that, in a k -connected graph, any k vertices lie on a common cycle, and that this is not necessarily true for $k + 1$ distinct vertices. Dirac’s results marked the starting point for a number of results giving conditions under which a set of vertices lies on a common cycle, and we refer the reader to Gould’s survey [83] for a detailed overview of results in this direction.

When trying to find a cycle containing some specified edges, research has been driven by a number of conjectures due to Lovász [116] (1973) and Woodall [169] (1977). The strongest of these is the following:

CONJECTURE 10.1.1 (Lovász-Woodall Conjecture). *Let S be a set of k independent edges in a k -connected graph G . If k is even or $G - S$ is connected, then there is a cycle in G containing S .*

Building on earlier work by Woodall, in particular on a technique of Woodall from [169] called the *Hopping Lemma*, Häggkvist and Thomassen [84] (1982) and Kawarabayashi [99, Theorem 2] (2002) established the following variants of the Lovász-Woodall Conjecture. First, in the case of Häggkvist and Thomassen, by setting out from the stronger assumption of $(k + 1)$ -connectedness, and second, in the case of Kawarabayashi, by obtaining a weaker conclusion, namely, two cycles instead of one.

THEOREM 10.1.2 (Häggkvist and Thomassen). *For any set S of k independent edges in a $(k + 1)$ -connected graph, there is a cycle in G containing S .*

THEOREM 10.1.3 (Kawarabayashi). *Let S be a set of k independent edges in a k -connected graph G . If k is even or $G - S$ is connected, then S is contained in one or a union of two vertex disjoint cycles of G .*

¹This paper follows the notation in Bollobás’ *Graph theory*, [27].

In the present paper, we are interested in a further variant of the problem, where instead of a cycle we aim to find a circuit—a closed walk that repeats no edges (but may repeat vertices)—containing a set of prescribed edges. Clearly, for this variant, it is no longer necessary to assume our edges to be independent. If one aims for results similar in spirit to the cycle case above, it seems natural to consider edge-connectivity instead of vertex connectivity. But whereas in the above cases, vertex connectivity is a far-from necessary condition, the corresponding version for circuits admits a complete characterisation in terms of edge cuts, which is the main result of our paper.

THEOREM 10.1.4. *A connected graph contains a circuit through any k prescribed edges if and only if it contains no odd cut of size at most k .*

COROLLARY 10.1.5. *If for some $k \in \mathbb{N}$ a connected graph contains a circuit through any $2k - 1$ prescribed edges, then it also contains a circuit through any $2k$ prescribed edges.*

While all the graphs treated in this paper are simple, one can easily derive the same characterisation for multigraphs, since subdividing every edge of a multigraph once does not give rise to new odd cuts.

To see that the condition in Theorem 10.1.4 is necessary, recall that the graph given by the vertices and edges of a circuit is Eulerian, i.e. even and connected, and so a necessary requirement for finding a circuit through a set of edges is that it can be extended to an even subgraph. The latter has been characterised by Jaeger [97] in 1979.

THEOREM 10.1.6 (Jaeger). *A set of edges in a graph G is contained in an even subgraph of G if and only if it contains no odd cut of G .*

However, while Jaeger's theorem immediately shows the necessity of our characterising condition in Theorem 10.1.4, it does not yield its sufficiency, as Jaeger's even subgraph is not necessarily connected (even if G is). This issue was also overlooked by Lai [108]. See Section 10.5 for further discussion when Jaeger's condition does give rise to a circuit.

EXAMPLE 10.1.7 (Counterexample to [108, Theorem 1.1 & 4.1]). Let $k \geq 3$, let G be the ladder with $k + 1$ rungs, and S be a set of rungs of G of size $3 \leq |S| \leq k$. Then S extends to an even subgraph of G , but every such even subgraph has at least $\lceil \frac{|S|}{2} \rceil \geq 2$ components.

PROOF. Since $G - S$ is connected, the set S does not contain any cut of G (regardless of its parity), and so S extends to an even subgraph by Theorem 10.1.6.

Now let $e_1, e_2, e_3 \in S$ be three edges ordered from left to right (cf. Figure 10.1), and suppose for a contradiction there is an even, connected subgraph H of G containing e_1, e_2, e_3 . Let C and C' be the edge cuts consisting of the two incident edges to the left and to the right of e_2 respectively (cf. Figure 10.1). Since H is connected and contains e_1




FIGURE 10.1. A ladder with specified rungs $S = \{e_1, \dots, e_k\}$.

and e_3 , H meets both cuts C and C' . Since H is even, it meets every cut of G in an even number of edges, and so $C \cup C' \subseteq E(H)$. But then both end vertices of e_2 have degree three in H , a contradiction to H being even.

In particular, if H is an even subgraph containing S , then every component of H contains at most two rungs from S , and so H has at least $\lceil \frac{|S|}{2} \rceil$ components. \square

So instead of referring to Jaeger's theorem for proving the sufficiency of the characterising condition in Theorem 10.1.4, we once more build on the technique of Woodall's hopping lemma.

Finally, let us mention the survey by Catlin [45] for related research on the existence of spanning circuits in a graph. Lai [108, Theorem 3.3] established the following sufficient condition for a graph to contain a spanning circuit through any k prescribed edges:

THEOREM 10.1.8 (Lai). *For $k \in \mathbb{N}$ let $f(k)$ be the smallest even integer $\geq \max(k, 4)$. If G is $f(k)$ -edge-connected, then G contains a spanning circuit through any k prescribed edges.*

A related variant is to find spanning trails (not necessarily closed) containing a given set of edges, see e.g. [170] and the references therein.

10.2. Preliminaries

All graphs in this paper are finite and simple. We let $\mathbb{N} = \{0, 1, 2, \dots\}$ and use $[n] = \{1, 2, \dots, n\}$ and $[0, n] = \{0, 1, \dots, n\}$. For our use of the terms *cycle*, *walk*, *trail* and *circuit*, we follow [27]. Let us clarify the use of technical terms now.

DEFINITION 10.2.1. Let $G = (V, E)$ be a graph. For a set of vertices $A \subseteq V$, we write

- $\partial_G A := \{uv \in E : u \in A, v \notin A\}$ for the *edge boundary* of A in G .

For $F \subseteq E$, we call

- F a *cut* of G , if there is an $A \subseteq V$ such that $\partial_G A = F$, and
- a cut F *odd*, if $|F|$ is odd. Otherwise, we call F *even*.

Recall that all cuts of some graph are even if and only if all its vertices have even degree.

DEFINITION 10.2.2. Let $G = (V, E)$ be a graph, and let $T = v_0 \dots v_r$ a walk in G .

- T is a *trail* in G , if all of its edges are distinct. Further, v_0 is the *start vertex* and v_r is the *end vertex* of T , and all other vertices are called *inner vertices* of T .
- T is *closed*, if its start and end vertex agree. A closed trail is also called *circuit*.
- $V(T)$ and $E(T)$ denote the vertices and edges of the underlying subgraph of T .

DEFINITION 10.2.3. Let $G = (V, E)$ be a graph. For $x, y \in V$, $X, Y \subseteq V$ and trails $P = p_0 \dots p_r$ and $Q = q_0 \dots q_w$ in G , we define

- PQ or $p_0 P p_r Q q_w$ is the *concatenated trail* $p_0 \dots p_r q_1 \dots q_w$ (only when P and Q are edge-disjoint and $p_r = q_0$),
- P is an X – Y *trail*, if $p_0 \in X$, $p_r \in Y$ and no inner vertex is in X or Y . For singletons write x – y trail instead of $\{x\}$ – $\{y\}$ trail,
- P is a *subtrail* of Q with *witnessing interval* $I_P = \{t_P, \dots, t_P + r\} \subseteq [0, w]$, if $p_h = q_{t_P+h}$ for every $h \in [0, r]$ or $p_h = q_{t_P+r-h}$ for every $h \in [0, r]$, and
- $\bar{Q} = q_w \dots q_0$ is the *reversed trail* of Q .

FACT 10.2.4. If P is a subtrail of Q and P uses at least one edge, then the witnessing interval I_P of P in Q is unique.

PROOF. Let $P = p_0 \dots p_r$ and $Q = q_0 \dots q_w$ with $r \geq 1$. Note that while for a single vertex p_i there might be several q_j with $p_i = q_j$, for every edge $p_{i-1}p_i$ there is a unique $j = j(i) \in [w]$ with $p_{i-1}p_i = h_{j-1}h_j$ (since our graphs are simple). From this, it follows that $I_P = \bigcup_{i \in [r]} \{j(i) - 1, j(i)\}$, and so the witnessing interval I_P of P in Q is unique. \square

DEFINITION 10.2.5. Let $(X, <_X)$ be a finite linear order. For $a \leq_X b \in X$, we define

- $[a, b]_{<_X} := \{\ell \in X : a \leq_X \ell \leq_X b\}$ as the *closed interval* from a to b .

Further, for a subset $Y \subseteq X$, we write

- $\max_{<_X} Y$ for the *greatest element* of Y with respect to $<_X$, and
- $\min_{<_X} Y$ for the *smallest element* of Y with respect to $<_X$.

10.3. A reduction to the bridge case

The proof of our characterisation theorem of graphs containing a circuit through any k prescribed edges will proceed via induction on k . For the induction step, suppose we have $k + 1$ edges e_1, \dots, e_{k+1} of G and may assume inductively that any k edges lie on a common circuit in G . Let H be such a circuit through e_1, \dots, e_k in G . Our task is then to also incorporate the last edge e_{k+1} into a circuit.

As our first result, we will show that it suffices to consider the case where e_{k+1} is a bridge in $G - E(H)$. More precisely, we claim that it suffices to prove the following theorem:

THEOREM 10.3.1. *Let G be a graph containing no odd cut of size at most $k + 1$, let $\{e_1, \dots, e_{k+1}\}$ be a collection of $k+1$ edges in G , and H be a circuit in G through e_1, \dots, e_k such that e_{k+1} is a bridge in $G - E(H)$.*

Then there exists a circuit H' in G through e_1, \dots, e_{k+1} . Moreover, if an end vertex of e_{k+1} is not in $V(H)$, then we may assume that H' passes it exactly once.

We defer the proof of Theorem 10.3.1 until the next section, and first show how to complete the proof of the Characterisation Theorem 10.1.4 given Theorem 10.3.1.

PROOF OF THEOREM 10.1.4 GIVEN THEOREM 10.3.1. As announced, the proof of the sufficiency of the characterisation in Theorem 10.1.4 will go via induction on k . The base case is easy: A connected graph without odd cuts of size at most $k = 1$ is evidently the same as a bridgeless connected graph. But any edge in such a graph lies on a circuit.

Now assume inductively that Theorem 10.1.4 holds for some integer $k \in \mathbb{N}$. To prove Theorem 10.1.4 in the case $k + 1$, let G be a graph containing no odd cut of size at most $k + 1$, and $S = \{e_1, \dots, e_{k+1}\}$ a collection of $k + 1$ edges in G . By induction, we may find a circuit H in G through e_1, \dots, e_k . If $e_{k+1} \in E(H)$, we are done.

So assume that $e_{k+1} \notin E(H)$. If e_{k+1} is a bridge of $G - E(H)$, then we are done by Theorem 10.3.1 (the moreover-part is not needed in this case). Otherwise, e_{k+1} is not a bridge in $G - E(H)$, and we may pick D as the maximal 2-edge-connected subgraph of $G - E(H)$ containing e_{k+1} .

Note that D and H are edge-disjoint, but might share vertices. If they do, choose $v \in V(D) \cap V(H)$ arbitrarily. To see that there is a circuit H^* in D containing v and e_{k+1} , construct an auxiliary graph D' from D by subdividing e_{k+1} by a new vertex w . Since D is 2-edge-connected, so is D' . By Menger's theorem, there are two edge-disjoint $w-v$ paths in D' translating to the desired circuit H^* in D . Since H and H^* are edge-disjoint and intersect in v , it is clear that $E(H) \cup E(H^*)$ is the edge set of a circuit covering S .

Thus, we may assume that $V(D) \cap V(H) = \emptyset$. Let $F := \partial_G(V(D)) \subseteq E \setminus E(H)$ and observe that every edge in F is a bridge in $G - E(H)$. Since G is connected, F is non-empty, and we choose $e_F \in F$ arbitrarily. Write $e_F = uw$ with $u \in V(D)$. Next, we contract D in G . Let G' be the resulting graph and $v_D \in V(G')$ be the vertex corresponding to the contracted D .

Observe that H is still a circuit through e_1, \dots, e_k in G' , that v_D is not contained in $V(H)$ and that G' is simple. Furthermore, every cut of G' is also a cut in G (after uncontracting v_D), and so G' contains no odd cut of size at most $k + 1$. Hence, we may apply Theorem 10.3.1 to G' , H and e_F to find a circuit $H' \subseteq G'$ through e_1, \dots, e_k and e_F , such that H' passes v_D exactly once (by the moreover-part). Let $e = u'w'$ with $u' \in V(D)$ be the edge in F corresponding to the other edge in H' incident with v_D . The circuit H' in G' corresponds to an $u'-u$ trail H^* in $G - E(D)$. By subdividing e_{k+1} in D once and using Menger's theorem in the resulting 2-edge-connected graph D' , we find an $u-u'$ trail

Q in D through e_{k+1} . Since Q and H^* are edge-disjoint, it follows that $uQu'H^*u$ is the desired circuit in G through e_1, \dots, e_{k+1} . \square

10.4. Proving the bridge case

In this section, we prove Theorem 10.3.1, completing the proof of the characterisation stated in Theorem 10.1.4. As indicated in the introduction, our proof of Theorem 10.3.1 is based on the so-called Hopping Lemma due to Woodall [169].

Throughout this section, when describing our set-up and stating our auxiliary results, we work in a fixed 2-edge connected graph $G = (V, E)$, with $S = \{e_1, \dots, e_{k+1}\}$ a collection of $k + 1$ edges of G , and H a shortest circuit through e_1, \dots, e_k in G . Any remaining assumptions featuring in Theorem 10.3.1 will only be used in the final proof of Theorem 10.3.1 itself at the very end of this section.

If e_1, \dots, e_k lie on a cycle C , then $C - \{e_1, \dots, e_k\}$ naturally falls apart into components, each of which is a path. If as in our situation e_1, \dots, e_k lie on a common circuit H , then $H - \{e_1, \dots, e_k\}$ also falls apart into segments: subtrails H_1, \dots, H_k of H such that (after relabelling our edges) we have $H = H_1e_1H_2e_2 \dots e_{k-1}H_k e_k$. Note, however, that different segments of $H - \{e_1, \dots, e_k\}$ are no longer vertex-disjoint (and so do not correspond to components of the *subgraph* $H - \{e_1, \dots, e_k\}$, cf. Figure 10.2).

DEFINITION 10.4.1. Given the circuit $H = H_1e_1H_2e_2 \dots e_{k-1}H_k e_k$, we call H_j the j -th *segment* of H . Since H is shortest possible, every segment H_j is a path. We let \prec_j denote the path order on $V(H_j)$ induced by the circuit H .




FIGURE 10.2. A circuit $H = H_1e_1H_2e_2H_3e_3$ with segments H_1, H_2, H_3 .

DEFINITION 10.4.2. Given the circuit H with segments $\{H_j : j \in [k]\}$, for $U \subseteq V$ and $j \in [k]$, we define (cf. Definition 10.2.5)

- (1) $\text{On}_j(U) := U \cap V(H_j)$ as the *vertices of U on the j -th segment of H* ,
- (2) $\text{Cl}_j(U) := [\min_{\prec_j} \text{On}_j(U), \max_{\prec_j} \text{On}_j(U)]_{\prec_j}$ as the *closure of U on the j -th segment of H* ,
- (3) $\text{Cl}(U) := \bigcup_{\ell \in [k]} \text{Cl}_\ell(U)$ as the *closure of U in H* ,

- (4) $\text{Fr}_j(U) := \{\min_{<_j} \text{On}_j(U), \max_{<_j} \text{On}_j(U)\}$ as the *frontier of U on the j -th segment of H* and
 (5) $\text{Fr}(U) := \bigcup_{\ell \in [k]} \text{Fr}_\ell(U)$ as the *frontier of U in H* .

Note that due to the fact that different segments can intersect, the set inclusions $\text{Cl}(U) \subseteq \text{Cl}(\text{Cl}(U))$, $\text{Cl}_j(U) \subseteq \text{On}_j(\text{Cl}(U))$ and $\text{Fr}_j(U) \subseteq \text{On}_j(\text{Fr}(U))$ might be proper.

FACT 10.4.3. For $j \in [k]$ and $U \subseteq V$, we have $\text{Cl}_j(U)$ is a subtrail of H_j .

DEFINITION 10.4.4. For $x, y \in V(G)$ and $X \subseteq V$, we say

- (1) an x - y trail P is *admissible*, if it is in $G - E(H) - e_{k+1}$ and $V(P) \cap V(H) \subseteq \{x, y\}$, and
 (2) $\text{R}(X) := \{y' \in V(H) : \exists x' \in X \exists \text{ admissible } x'-y' \text{ trail}\}$ as *reach of X after H* .

We stress that the inner vertices of an admissible x - y trail are not in $V(H)$.

DEFINITION 10.4.5. We define an increasing sequence $(A_i)_{i \in \mathbb{N}}$ recursively by

- (1) $A_0 := \emptyset$,
 (2) $A_1 := \text{R}(\{a\})$, and
 (3) if A_i is already defined for some $i \geq 1$, then $A_{i+1} := \text{R}(\text{Cl}(A_i))$.

Further, we set $A := \bigcup_{i \in \mathbb{N}} A_i$. Analogously, we define an increasing sequence $(B_i)_{i \in \mathbb{N}}$ and B by interchanging a with b .

The idea behind this definition is the simple observation that if A_1 and B_1 intersect the same segment of H , then we clearly would be done. This will not always be possible, and so we iterate this procedure again and again, until we do find one vertex in A and one vertex in B that are contained in the same segment of H , as Lemma 10.4.6 below shows.

We remark that Definition 10.4.5 of $(A_i)_{i \in \mathbb{N}}$ differs from Woodall's in that Woodall's admissible paths (see $x \star y$ in [169]) from A_i to new vertices of A_{i+1} are not allowed to start from the frontier of A_i .

LEMMA 10.4.6. *If $\text{On}_j(A) = \emptyset$ or $\text{On}_j(B) = \emptyset$ for every $j \in [k]$, then G contains an odd cut of size at most $k + 1$.*

PROOF. First of all, since G is 2-edge-connected, both A and B are non-empty: Since $G - e_{k+1}$ is connected, any a - $V(H)$ path in $G - e_{k+1}$ is an admissible trail which witnesses the non-emptiness of $A_1 \subseteq A$, and similarly for B .

Since $A, B \subseteq V(H)$ and $\text{On}_j(A) = \emptyset$ or $\text{On}_j(B) = \emptyset$ for every $j \in [k]$, A and B are disjoint. Further, from the pigeonhole principle it follows without loss of generality, that $|\{j \in [k] : \text{On}_j(A) \neq \emptyset\}| \leq \lfloor \frac{k}{2} \rfloor$. Then

$$|\partial_H A| = \left| \bigcup_{j \in [k]} \partial_{(e_{j-1}H_j e_j)} \text{On}_j(A) \right| \leq \sum_{j \in [k]} |\partial_{(e_{j-1}H_j e_j)} \text{On}_j(A)| = 2 \cdot \lfloor \frac{k}{2} \rfloor,$$

and since H induces an even subgraph, $|\partial_H A|$ is even. Thus, $C := \partial_H A \cup \{e_{k+1}\}$ is odd and has size $|C| \leq 2 \cdot \lfloor \frac{k}{2} \rfloor + 1 \leq k + 1$.

To complete the proof, it remains to show that C is a cut in G . For this, we consider

$$D = \{v \in V(G) : \exists a' \in A \cup \{a\} \exists \text{ admissible } a'-v \text{ trail}\},$$

and claim that $\partial_G D = C$.

To see $C \subseteq \partial_G D$, note that $\partial_H A \subseteq \partial_G D$ by definition of A . For $e_{k+1} \in \partial_G D$, suppose to the contrary that $b \in D$. Then there exists an admissible $b-(A \cup \{a\})$ trail T . Since $B \neq \emptyset$, T combined with an admissible $b-B$ trail witnesses that $A \cap B \neq \emptyset$, a contradiction.

To prove $\partial_G D \subseteq C$, let us suppose for a contradiction that there exists some edge $e = uv \in (\partial_G D) \setminus C$ with say $u \in D$ and $v \notin D$. Since $u \in D$, we can pick an admissible trail T starting in some $a' \in A \cup \{a\}$ and ending in u . If $e \in E(H)$, then $u \in V(H)$ and thus $u \in A$ by Definition 10.4.5. Now $e \in \partial_H(A)$, which contradicts $e \notin C$. So, we assume $e \notin E(H)$. If $u \in V(H)$, then $u \in A$ and the trail uv is a witness for $v \in D$. Otherwise, Tuv is a witness. In any case, this contradicts $v \notin D$. \square

Now that we know that A and B intersect the same segment H_j of H , it is clear that there is a natural trail in H starting at a vertex of A , ending at a vertex of B , and containing all of the edges e_1, \dots, e_k . If we consider the ‘first time’ that A_n and B_m intersect a given segment H_j , then this trail has the following three crucial properties of Definition 10.4.7, as Lemma 10.4.8 shows.

DEFINITION 10.4.7. For $n, m \in \mathbb{N}$, we say a trail $Q = q_0 \dots q_w$ is A_n-B_m -coherent, if

- (C₁) $e_1, \dots, e_k \in E(Q)$, $q_0 \in A_{n+1}$ and $q_w \in B_{m+1}$,
- (C₂) for every $s \in [w]$ with $q_{s-1}q_s \in E \setminus E(H)$, there exist $r, t \in [0, w]$ with $q_r, q_t \in V(H)$ and $r < s \leq t$ such that $q_r Q q_t$ is an admissible q_r-q_t trail and each of the sets A_{n+1} and B_{m+1} contains at most one of q_r and q_t , and
- (C₃) for every $j \in [k]$, $\text{Cl}_j(A_n)$ and $\text{Cl}_j(B_m)$ are subtrails of Q with witnessing intervals $I_{A_n, j}$ and $I_{B_m, j}$ such that $I_{X, j} \cap I_{Y, j'} = \emptyset$ for every $X, Y \in \{A_n, B_m\}$ and every two distinct $j \neq j' \in [k]$.

LEMMA 10.4.8. *If $\text{Cl}_j(A_{n^*}) \neq \emptyset \neq \text{Cl}_j(B_{m^*})$ for some $j \in [k]$, then there exists an A_n-B_m -coherent trail for some $n < n^*$, $m < m^*$.*

PROOF. Let j be in $[k]$ such that $\text{Cl}_j(A_{n^*}) \neq \emptyset \neq \text{Cl}_j(B_{m^*})$. Choose $n < n^*$ and $m < m^*$ minimal such that $\text{Cl}_j(A_{n+1}) \neq \emptyset \neq \text{Cl}_j(B_{m+1})$ and pick $a_{n+1} \in \text{On}_j(A_{n+1})$ and $b_{m+1} \in \text{On}_j(B_{m+1})$.² We claim that the trail Q with start vertex a_{n+1} and end vertex

²One could make a stronger minimality assumption by choosing n, m minimal so that $\text{Cl}_j(A_n) \neq \emptyset \neq \text{Cl}_j(B_n)$ for some $j \in [k]$. Following the same proof, this gives rise to a trail Q which satisfies the following stronger variant of (C₃), namely $I_{X, j} \cap I_{Y, j'} = \emptyset$ for every $X, Y \in \{A_n, B_m\}$ and every two (not necessarily distinct) $j, j' \in [k]$. However, we do not need this stronger conclusion for the remainder of our proof.

b_{m+1} along the circuit H through e_1, \dots, e_k and $H_{j'}$ as subtrail for every $j' \in [k] \setminus \{j\}$ is $A_n - B_m$ -coherent as desired.

Indeed, (C_1) holds by construction and (C_2) is an empty condition. Lastly, since $\text{Cl}_j(A_n) = \emptyset = \text{Cl}_j(B_m)$, and all other segments $H_{j'}$ for $j' \in [k] \setminus \{j\}$ are subtrails of Q with pairwise disjoint witnessing intervals by construction, also (C_3) holds for Q . \square

While conditions (C_1) and (C_2) are straightforward adaptations from Woodall's notion of coherence [169, §III] from paths to trails, a word on (C_3) might be in order. Given the 'time-minimal' subtrail Q of H constructed in Lemma 10.4.8, we aim to modify Q while preserving as much structure of Q , and hence of H , as possible. Since segments of H may intersect, the correct notion of 'structure preserving' is to think about the trail in terms of time: Our initial trail Q constructed in Lemma 10.4.8 spends disjoint time intervals to cover the different segments of H that contain vertices from $\text{Cl}(A_n) \cup \text{Cl}(B_n)$. When modifying Q , however, we can no longer require to completely cover all these segments. So instead, we only preserve the property that if T and S are subpaths of *distinct* segments H_j and $H_{j'}$ of the form $T \in \{\text{Cl}_j(A_n), \text{Cl}_j(B_m)\}$ and $S \in \{\text{Cl}_{j'}(A_n), \text{Cl}_{j'}(B_m)\}$, then we continue to spend disjoint time intervals to cover T and S .

THEOREM 10.4.9. *If there exists an $A_n - B_m$ -coherent trail for some $n, m \in \mathbb{N}$, then there also exists an $A_0 - B_0$ -coherent trail.*

For the proof, we need two easy lemmas.

LEMMA 10.4.10. *Let $n, m \in \mathbb{N}$ and $Q = q_0 \dots q_w$ be an $A_n - B_m$ -coherent trail. If $n \geq 1$ and $q_0 \in A_n$, then Q is $A_{n-1} - B_m$ -coherent, and if $m \geq 1$ and $q_w \in B_m$, then Q is $A_n - B_{m-1}$ -coherent.*

PROOF. Due to the symmetry of the statements, we just check the conditions for Q being $A_{n-1} - B_m$ -coherent for $n \geq 1$. Property (C_1) is clear, and (C_2) is immediate from the fact that $(A_i)_{i \in \mathbb{N}}$ is an increasing sequence.

Finally, (C_3) follows from the fact that since $(A_i)_{i \in \mathbb{N}}$ is increasing, $\text{Cl}_j(A_{n-1})$ is a subtrail of $\text{Cl}_j(A_n)$, and hence we have $I_{A_{n-1}, j} \subseteq I_{A_n, j}$ for the respective witnessing intervals for all $j \in [k]$. Since $I_{X, j} \cap I_{Y, j'} = \emptyset$ for every $X, Y \in \{A_n, B_m\}$ and every two distinct $j, j' \in [k]$ holds by assumption, it follows that the same holds for every $X, Y \in \{A_{n-1}, B_m\}$. \square

LEMMA 10.4.11. *Let $n, m \in \mathbb{N}$ and $v \in (\text{Cl}(A_n) \cup \{a\}) \cup (\text{Cl}(B_m) \cup \{b\})$. If $Q = q_0 \dots q_w$ is an $A_n - B_m$ -coherent trail and P is an admissible $v - V(H)$ trail, then Q and P are edge-disjoint.*

PROOF. By symmetry we may assume that $v \in \text{Cl}(A_n) \cup \{a\}$. Suppose for a contradiction that P and Q are not edge-disjoint. Choose $s \in [w]$ such that $q_{s-1}q_s$ is the first edge of P that is also in $E(Q)$. Since $q_{s-1}q_s \in E(P) \subseteq E \setminus E(H)$ by Definition 10.4.4, it follows from property (C_2) of $A_n - B_m$ -coherent that there are $r, t \in [0, w]$ with $r < s \leq t$

and $q_r, q_t \in V(H)$ such that $q_r Q q_t$ is an admissible q_r – q_t trail and each set A_{n+1} and B_{m+1} contains at most one of q_r and q_t . But since $q_{s-1} q_s$ is the first edge of P in $E(Q)$, both $v P q_{s-1} \bar{Q} q_r$ and $v P q_{s-1} Q q_t$ are admissible trails witnessing that $q_r, q_t \in A_{n+1}$ (cf. Definition 10.4.5(3)), a contradiction. \square

PROOF OF THEOREM 10.4.9. Let n, m be minimal such that there is an A_n – B_m –coherent trail $Q = q_0 \dots q_w$ with start vertex $q_0 = a_{n+1} \in A_{n+1}$ and end vertex $q_w = b_{m+1} \in B_{m+1}$. We claim that $n = m = 0$. Otherwise, without loss of generality we may assume $n \geq 1$. By Lemma 10.4.10 and the minimality assumption, we have $a_{n+1} \in A_{n+1} \setminus A_n$. We write Q as $a_{n+1} Q q_c Q q_d Q b_{m+1}$ where $c, d \in [0, w]$ are defined as follows:

- (a) Since $a_{n+1} \in A_{n+1} \setminus A_n$ and by Definition 10.4.5(3) of A_{n+1} , there is an $x \in \text{Cl}(A_n)$ such that there exists an admissible x – a_{n+1} trail P (which might be trivial). From Definition 10.4.2(3) of the closure it follows that there is an $j \in [k]$ such that $x \in \text{Cl}_j(A_n)$. By property (C₃) of A_n – B_m –coherent, $\text{Cl}_j(A_n)$ is a subtrail of Q with witnessing interval $I_{A_n, j} \subseteq [0, w]$. Now, we choose $d \in I_{A_n, j}$ as the unique index with $q_d = x$.
- (b) Next, choose $c := \max\{r \in [0, w] : q_r \in \bigcup_{i \in [n]} \text{Fr}_j(A_i) \wedge r \leq d\}$. If $r := \min I_{A_n, j}$, then $q_r \in \text{Fr}_j(A_n)$ and obviously $r \leq d$. Hence, c exists.




FIGURE 10.3. Obtaining the rerouted trail Q' from Q .

Further, we set $n' := \min\{i \in [n] : q_c \in \text{Fr}_j(A_{i+1})\}$ and observe

- (1) $I_{A_{n'}, j} \cap [c, d] = \emptyset$,
- (2) if $I_{B_m, j} \cap [c, d] \neq \emptyset$, then $\text{Cl}_j(A_n) \cap \text{Cl}_j(B_{m+1}) \neq \emptyset$,
- (3) if $q_d \in B_{m+1}$, then $\text{Cl}_j(A_n) \cap \text{Cl}_j(B_{m+1}) \neq \emptyset$, and
- (4) P and Q are edge-disjoint.

PROOF OF (1). We assume for a contradiction that $I_{A_{n'}, j} \cap [c, d] \neq \emptyset$. Then, either choosing r as $\min(I_{A_{n'}, j} \cap [c, d])$ or $\max(I_{A_{n'}, j} \cap [c, d])$ will lead to $q_r \in \text{Fr}_j(A_{n'})$, which is a contradiction to the choice of c or n' because $c \leq r \leq d$.

PROOF OF (2). Let $I_{B_m, j} \cap [c, d] \neq \emptyset$. So, $I_{A_n, j} \cap I_{B_m, j} \neq \emptyset$ because $[c, d] \subseteq I_{A_n, j}$. Further, $\text{Cl}_j(A_n) \cap \text{Cl}_j(B_m) \subseteq \text{Cl}_j(A_n) \cap \text{Cl}_j(B_{m+1})$ implies then that $\text{Cl}_j(A_n) \cap \text{Cl}_j(B_{m+1}) \neq \emptyset$.

PROOF OF (3). If $q_d \in B_{m+1}$, then, $q_d \in \text{Cl}_j(A_n) \cap \text{Cl}_j(B_{m+1}) \neq \emptyset$.

PROOF OF (4). Since $q_d \in \text{Cl}_j(A_n) \subseteq \text{Cl}(A_n)$ and Q is $A_n - B_m$ -coherent, this follows from Lemma 10.4.11.

If $I_{B_m, j} \cap [c, d] \neq \emptyset$ or $q_d \in B_{m+1}$, then (2) or (3) imply that $\text{Cl}_j(A_n) \cap \text{Cl}_j(B_{m+1}) \neq \emptyset$, which by Lemma 10.4.8 gives rise to a coherent trail that contradicts the minimality of n and m . Hence, we assume $I_{B_m, j} \cap [c, d] = \emptyset$ and $q_d \notin B_{m+1}$.

Now we reroute Q and obtain $Q' := q_c \bar{Q} a_{n+1} \bar{P} q_d Q b_{m+1}$, see Figure 10.3. From (4) it follows that Q' is a trail. We show that Q' is $A_{n'} - B_m$ -coherent, contradicting the minimality of n and m :

- (C₁) Since $E(q_c \dots q_d) \subseteq E(H_j)$ and since all our edges satisfy $e_i \notin E(H_j)$, the fact that Q satisfied (C₁) implies that Q' uses e_1, \dots, e_k . Also, the start vertex q_c is in $\text{Fr}_j(A_{n'+1}) \subseteq A_{n'+1}$ and the end vertex b_{m+1} is still in B_{m+1} .
- (C₂) Because $a_{n+1} \notin A_n \supseteq A_{n'+1}$ and $q_d \notin B_{m+1}$, each of the sets $A_{n'+1}$ and B_{m+1} contains at most the start or the end vertex of P . Also, the $q_d - a_{n+1}$ trail P is admissible. This implies that (C₂) is true for edges that are in P . For edges that are not in P , it follows directly from Q 's (C₂) and $q_c, q_d \in V(H)$.
- (C₃) Due to (1) and $I_{B_m, j} \cap [c, d] = \emptyset$, the trails $\text{Cl}_{j'}(A_{n'})$ and $\text{Cl}_{j'}(B_m)$ are subtrails of $q_1 \dots q_c$ or $q_d \dots q_w$ for every $j' \in [k]$. Hence, Q' inherits property (C₃) from Q . □

We are now ready to complete the proof of Theorem 10.3.1.

PROOF OF THEOREM 10.3.1. Since G contains no odd cut of size at most $k + 1$, Lemma 10.4.6 implies that $\text{On}_j(A) \neq \emptyset \neq \text{On}_j(B)$ for some $j \in [k]$. By Lemma 10.4.8 there is an $A_n - B_m$ -coherent trail in $G - e_{k+1}$ for some $n, m \in \mathbb{N}$, and so by Theorem 10.4.9 there also exists an $A_0 - B_0$ -coherent trail Q from a vertex $a_1 \in A_1$ to a vertex $b_1 \in B_1$ in $G - e_{k+1}$.

By Definition 10.4.5(2) of A_1 and B_1 , there is an admissible $a - a_1$ trail P_a and an admissible $b - b_1$ trail P_b . Since e_{k+1} is a bridge in $G - E(H)$,³ the trails P_a and P_b are vertex-disjoint. Thus, P_a, P_b, Q and e_{k+1} are edge-disjoint by Lemma 10.4.11 and Definition 10.4.4(1). Together with property (C₁) of Q , it follows that $H' := baP_a a_1 Q b_1 \bar{P}_b b$ is the desired circuit in G through e_1, \dots, e_{k+1} .

To see the moreover-part of Theorem 10.3.1, observe that if $a \notin V(H)$, then $a \notin V(Q)$ due to (C₂) and Definition 10.4.5(2) of A_1 . Thus, the circuit H' passes a once, since P_a and P_b are vertex disjoint. The same holds for b . □

10.5. Concluding remarks and an open question

To find a circuit through any k prescribed edges we employed a global property by forbidding all odd cuts of bounded size. However, if we are only interested in one specific

³We remark that this is the only place in our argument where we use that e_{k+1} is a bridge in $G - E(H)$.

edge set, forbidding all bounded sized odd cuts seems unnecessarily strong: For example, if our k edges are contained in a $(k + 1)$ -edge-connected subgraph, then it is irrelevant whether the whole graph contains some further small odd cuts. Hence, the following natural question arises:

QUESTION 4. When can a given edge set of a graph G be covered by a circuit in G ?

One line of investigation could be whether a condition similar to the one in Jaeger's theorem 10.1.6 could be of additional help:

DEFINITION 10.5.1. For any $k \in \mathbb{N}$, let $g(k)$ be the smallest integer such that a set of at most k edges in a $g(k)$ -edge-connected graph G is covered by a circuit in G if and only if it contains no odd cut of G .

LEMMA 10.5.2. For any $k \in \mathbb{N}$,

- (1) $g(k) \leq m \leq k + 1$, where m is the smallest even integer $\geq k$, and
- (2) for $k \geq 4$, $g(k) > \ell$, where ℓ is the greatest odd integer $\leq \frac{1}{2}(\sqrt{8k - 7} + 1)$.

PROOF. The first part follows directly from Theorem 10.1.4.

For the lower bound of $g(k)$, let ℓ is the greatest odd integer $\leq \frac{1}{2}(\sqrt{8k - 7} + 1)$, and consider H_i to be a K_ℓ with $V(H) = \{v_{i,1}, \dots, v_{i,\ell}\}$ for $i \in [2]$. Further, we define $G := H_1 + H_2 + \{v_{1,j}v_{2,j} : j \in [\ell]\}$. We remark that G is ℓ -connected. Now, we pick $S := E(H_1) \cup \{e\}$ where e is some edge of $E(H_2)$. We calculate

$$|S| = \binom{\ell}{2} + 1 = \frac{\ell(\ell - 1)}{2} + 1 \leq k$$

where the inequality holds for $\ell \leq \frac{1}{2}(\sqrt{8k - 7} + 1)$. By Theorem 10.1.6, S contains no odd cut of G , because S is contained in the even subgraph $H_1 + H_2$. But clearly there exists no circuit H' in G that covers S . \square

FACT 10.5.3. We have

$$g(1) = 0, g(2) = 2, g(3) = 3 \text{ and } g(4) = 4.$$

PROOF. To see $g(1) = 0$, observe that any edge not being a bridge of its component must lie on a cycle.

For $g(2) = 2$, note that $g(2) \leq 2$ by Lemma 10.5.2, and $g(2) > 1$ by considering two disjoint cycles connected by an edge, and letting S consist of one edge from each cycle.

Next, Example 10.1.7 shows $g(3) > 2$. For $g(3) \leq 3$, let G be a 3-edge-connected graph and S be a 3-set of edges which contains no odd cut of size at most three. By Theorem 10.1.6, there exists an even subgraph H of G . We choose H subgraph-minimal, and so H has at most three components.

First, we assume that H has three components C_1, C_2, C_3 , and reduce it to the case where H has two components by considering the three edge-disjoint $V(C_1)-V(C_2 + C_3)$ paths in G which exist by Menger's theorem.

Now, we assume that H has two components C_1, C_2 where without loss of generality $|E(C_1) \cap S| = 1$. Again there are three edge-disjoint $V(C_1)-V(C_2)$ paths in G . At least two of them meet the same segment of C_2 such that we can construct a cycle in G which goes through all three edges.

Finally, $g(4) = 4$ follows from Lemma 10.5.2. □

Thus, by adding Jaeger's condition, for odd $|S|$ it appears we need less edge connectivity than before. It might be an interesting problem to find the precise values for the function f , or at least to improve any of the bounds given in Lemma 10.5.2. In particular, we were not able to find an example witnessing $g(5) > 4$.

n -Arc connected graphs

Given a graph G , of arbitrary size and unbounded vertex degree, denote by $|G|$ the one-complex associated with G . The topological space $|G|$ is *n -arc connected* (*n -ac*) if every set of no more than n points of $|G|$ are contained in an arc (a homeomorphic copy of the closed unit interval).

For any graph G , we show the following are equivalent: (i) $|G|$ is 7-ac, (ii) $|G|$ is n -ac for all n , and (iii) G is a subdivision of one of nine graphs. A graph G has $|G|$ 6-ac if and only if either G is one of the nine 7-ac graphs, or, after suppressing all degree-2-vertices, the graph G is 3-regular, 3-connected, and removing any 6 edges does not disconnect G into 4 or more components.

Similar combinatorial characterizations of graphs G such that $|G|$ is n -ac for $n = 3, 4$ and 5 are given. Together these results yield a complete classification of n -ac graphs, for all n .

11.1. Introduction

Graphs are typically considered as combinatorial objects: a set of vertices, along with a set of un-ordered pairs of vertices, forming edges abstractly connecting the vertices; but it is equally natural to consider graphs as geometric objects with a set of vertices and some pairs of vertices literally connected by an arc (a homeomorphic copy of the closed unit interval). Indeed right at the birth of graph theory, with Euler's solution of the Königsberg Bridges problem – asking for a particular kind of physical path – and a little later with Hamilton's solution of his Icosian problem – requesting an arc, or circle, in the skeleton of a dodecahedron, containing all vertices – the geometric view is the most immediate. When we think of a graph geometrically (a 1-complex), then the points on edges become *first class citizens*, and this change in perspective opens up new classes of problems. In this paper we always consider a given combinatorial graph as a geometric graph with the natural underlying set, topologized in any way so that each arc forming an edge has its usual topology as a subspace (the exact topology on the graph will not turn out to be important here).

A natural extension of Hamilton's problem is to ask, for some n , which graphs G are n -arc Hamilton (respectively, n -Hamilton) that for any choice of at most n vertices there is an arc (respectively, a circle) in G containing the specified points. For example, a classical theorem in graph theory of Dirac [61, Satz 9] says that a n -connected graphs are n -Hamilton. (Recall that a combinatorial graph is *k -connected* if deleting at most

$k - 1$ vertices does not result in a disconnection.) However, high connectivity isn't always necessary, indeed every cycle is n -Hamilton, for all n , despite it being only 2-connected. Dirac also noted in [61] that n -connected graphs are not necessarily $(n + 1)$ -Hamilton. A characterization was found in [163]: let G be an n -connected graph, $n \geq 3$, then G is $(n + 1)$ -Hamilton if and only if no set T of vertices of G of size n separates G into more than n components. As is well-known, despite the existence of simple sufficient conditions, there is still no characterization of Hamiltonicity. Let us also note that Egawa, Glas & Locke [163] gave a sufficient condition for an n -connected graph to be $(n + 1)$ -arc Hamilton, but the authors are not aware of characterizations of the $(n + 1)$ -arc Hamilton, n -connected graphs.

Taking the geometric viewpoint we are led to consider connecting *arbitrary* points in a graph by arcs or circles. Let G be a graph, considered as a topological space, and S a subset of G , then (S, G) is *n -arc connected* (or, S is *n -ac* in G) if for any choice of at most n elements of S there is an arc in G containing the specified points, while (S, G) is *n -circle connected* (or, S is *n -cc* in G) if for any choice of at most n elements of S there is a simple closed curve in G containing the specified points. Further, we say S is ω -ac (respectively, ω -cc) in X if it is n -ac (n -cc) in X for all integers $n \in \mathbb{N}$. Observe that a graph G with vertices V has a Hamiltonian path (respectively, cycle) if and only if V is $|V|$ -ac (respectively, $|V|$ -cc) in G , and is n -arc Hamilton (respectively, n -Hamilton) if and only if V is n -ac (respectively, n -cc) in G .

Define a graph G to be *n -ac* (resp., *n -cc*) if G is n -ac (resp., n -cc) in G – in other words, for any choice of at most n points of G there is an arc (resp., circle) in G containing the specified points. In this paper we give a complete solution to the problem of characterizing which graphs are n -ac or n -cc. By ‘complete’ we mean for any n , and for any graph, without restriction on the number of vertices, or edges, or the degree of any vertex. Our characterizations give tests for a graph to be n -ac or n -cc which are combinatorial in nature, only referring to vertices and edges, and which are polynomial in the number of vertices for finite graphs. The proofs largely rely on Menger-type results, and arguments based on (and in some cases, extending) the theory of alternating walks.

In describing our results, we should start by stating that it is straightforward to see that a graph is 2-cc if and only if it is 2-connected, while the only 3-cc graphs are cycles (see Theorem 11.3.1 and preceding discussion). Hence our focus is on n -ac graphs, for some n . It is also clear that a graph G is 2-ac if and only if it is connected (combinatorially). In [69] it was shown that a non-degenerate finite connected graph G is 7-ac if and only if it is ω -ac if and only if G is homeomorphic to one of 6 graphs (arc, circle, lollipop, ϑ -curve, figure-of-eight, dumbbell). Extending that argument shows that this behaviour occurs also for arbitrary graphs. Indeed (see Theorem 11.3.15), a non-degenerate graph G is 7-ac if and only if it is n -ac for all n , and if and only G is homeomorphic to a finite list

of graphs, namely homeomorphic to one of the six finite 7-ac graphs, or one of the finite 7-ac graphs with some endpoints removed (giving three more possibilities).

It remains, then, to characterize the n -ac graphs for $n = 3, 4, 5$ and 6 . In [69] infinite families of finite graphs which are n -ac but not $(n + 1)$ -ac were given for all $2 \leq n \leq 6$, and the problem of characterizing finite n -ac graphs for $n = 3, 4, 5$ and 6 was raised. Theorems 11.3.1 and 11.3.2, 11.3.3 and 11.3.5, 11.3.6 and 11.3.7, and 11.3.10 solve that problem for arbitrary graphs. As a sample: a graph G is 6-ac if and only if either G is one of the nine 7-ac graphs mentioned above, or, after suppressing all degree-2-vertices, the combinatorial graph G is 3-regular, 3-connected, and removing any 6 edges does not disconnect G into 4 or more components.

11.2. Preliminaries

In this paper, the term *graph* refers to a combinatorial graph $G = (V, E)$ where V is a (possibly infinite) set and $E \subseteq [V]^2$. However, to every combinatorial graph $G = (V, E)$ we associate the topological space, $|G|$, which is the 1-complex of G , namely the quotient space $(V \oplus \bigoplus_{e \in E} [0, 1]_e) / \sim$ where V carries the discrete topology and for an edge $e = \{v, w\} \in E$ we identify v in V with the $0 \in [0, 1]_e$ and w with $1 \in [0, 1]_e$. In fact our results hold whenever the set $|G|$ is given a topology in which the image under the quotient of each $[0, 1]_e$ is homeomorphic to $[0, 1]$. (For example, the metric topology induced by vertex distance, extended to interior points of edges in the natural manner, would work equally well. The quotient topology is simply the finest one satisfying this property.) Where no confusion can arise – when we discuss purely topological notions, for example – we abuse notation and simply write G for $|G|$.

11.2.1. Notation and Conventions. Let $G = (V, E)$ be a graph. An edge, $e = \{v, w\}$ is often abbreviated, $e = vw$. By convention we label *fixed* vertices by a, b, \dots , and *general* vertices as v, w *et cetera*. Let $V = A \dot{\cup} B$ a partition of its vertex set. The set $E(A, B)$ of edges of G with one endpoint in A and the other in B is called an *edge-cut* of G .

For $A \subseteq V$ or $F \subseteq E$ we write $G[A]$ and $G[F]$ for the induced subgraph of G .

A subset e of $|G|$ is called a *closed edge* if it is the image under the quotient map of some $[0, 1]_e$, and is called an *edge* if it is a closed edge minus its endpoints. In particular note that edges are open sets. If $e = vw \in E(G)$ then $\bar{e} = \{v, w\} \cup e \subseteq |G|$ is the closed edge in $|G|$ naturally associated with the combinatorial edge e in G . By convention we label points in the space $|G|$ by x, y, \dots

11.2.2. Background from graph theory. Our main tools from graph theory will be the block-cutvertex decomposition of (possibly infinite) connected graphs, and certain variants of Menger's theorem, especially the ones involving the concept of alternating

walks. A good overview of these techniques is given, for example, in the chapter on connectivity of Diestel's book, [54, § 3.1 & 3.3].

A graph $G = (V, E)$ is *cyclically connected* if every two vertices lie on a cycle, and, recall, is *2-connected* if removing any single vertex does not disconnect the graph. Note that a graph is cyclically connected if, in our terminology, the vertices are 2-cc in the graph. Observe that, according to this definition, the complete graph on two vertices is 2-connected but not cyclically connected. However, by Menger's theorem stated below, any 2-connected graph with at least 3 vertices is cyclically connected.

11.2.2.1. *Block-cutvertex decomposition.* Let G be a connected graph. A *block* of G is an inclusion-maximal 2-connected subgraph. Every edge is contained in a block, and by their maximality, different blocks overlap in at most one vertex, which then must be a cut-vertex of G . Therefore, the blocks form an edge-disjoint decomposition of G . Let $C \subseteq V$ denote the set of cut-vertices of $G = (V, E)$ and \mathcal{B} the collection of blocks. The *block graph* $B(G)$ of G is the bipartite graph formed on the vertex set $C \dot{\cup} \mathcal{B}$ with an edge $cB \in E(B(G))$ if and only if $c \in B$. For the proof of the next lemma see [54, Lemmas 3.1.3 & 3.1.4].

LEMMA 11.2.1. *The block graph of a (possibly infinite) connected graph is a (possibly infinite) graph-theoretic tree.*

11.2.2.2. *Chain graphs and cycle graphs.* A connected graph is called a *chain graph* if it is a linearly ordered union (possibly just one) of subgraphs (called *links*) such that only consecutive graphs meet, and their intersection consists of a single vertex only (called the *linking vertex*). Thus, a connected graph G is a chain graph of 2-connected links if and only if its block graph $B(G)$ is a finite path, a ray, or a double ray.

Similarly, a connected graph G is a *cycle graph* if G is a union of graphs L_0, L_1, \dots, L_{n-1} (called *links*) for some $n \in \mathbb{N}$ where (a) $L_i \cap L_j = \emptyset$ if $|i - j| > 1 \pmod n$ and (b) L_i meets L_j at a single vertex (called the *linking vertex*) if $|i - j| = 1 \pmod n$.

11.2.2.3. *Menger's theorem and alternating walks.* We say a graph $G = (V, E)$ is k -connected for some $k \in \mathbb{N}$ if the induced subgraph $G - W$ is connected for all $W \subseteq V$ with $|W| < k$. We note that even if a graph G is k -connected for some large k , the underlying topological 1-complex $|G|$ is never 3-connected in the topological sense, as removing two endpoints of an edge disconnects the interior of that edge from the rest of the graph.

Given sets A, B of vertices, we call a path [walk] $P = x_0, \dots, x_n$ an $A - B$ *path* [walk] if $V(P) \cap A = \{x_0\}$ and $V(P) \cap B = \{x_n\}$, i.e. if the path [walk] starts in A , ends in B and is otherwise disjoint from $A \cup B$. Two or more paths are *independent* if none of them contains an inner vertex of another. If $A, B \subseteq V$ and $X \subseteq V$ is such that every $A - B$ path in G contains a vertex from X , then we say that X *separates* the vertex set A from B . This implies $A \cap B \subseteq X$, i.e. X does not have to be disjoint from A or B .

THEOREM 11.2.2 (Menger's Theorem). *Let $G = (V, E)$ be a (potentially infinite) graph and $A, B \subseteq V$. Then the minimum number of vertices separating A from B in G is equal to the maximum number of disjoint $A - B$ paths in G .*

PROOF. See [54, Theorem 3.3.1 & Prop 8.4.1] respectively for the proof for finite graphs and its infinite extension. \square

However, at several points in this paper we shall employ the following more algorithmic version of Menger's theorem using the notion of alternating walks. Recall that for given sets A, B of vertices and \mathcal{P} a collection of disjoint $A - B$ paths, a walk $W = x_0e_0x_1e_1 \dots e_{n-1}x_n$ in G with no repeated edges (but possibly repeated vertices) is *alternating* with respect to \mathcal{P} if

- W starts in $A \setminus \bigcup \mathcal{P}$,
- the only repeated vertices of W lie on paths in \mathcal{P} ,
- whenever W hits a vertex of some path $Q \in \mathcal{P}$ (meaning say $e_i \notin Q$ but $x_{i+1} \in Q$) then W follows Q back towards the direction of A for at least one edge of Q , and
- whenever W uses an edge e from a path $Q \in \mathcal{P}$, then W traverses this edge backwards.

We refer the reader to [54, Fig. 3.3.2] and the surrounding discussion for further information. The following two lemmas list the two crucial properties of alternating paths in the context of Menger's theorem. For the proof of the first see [54, Lemma 3.3.3].

LEMMA 11.2.3. *If no alternating walk ends in $B \setminus \bigcup \mathcal{P}$, then G contains an $A - B$ separator X on \mathcal{P} with $|X| = |\mathcal{P}|$.*

LEMMA 11.2.4. *If an alternating walk W ends in $B \setminus \bigcup \mathcal{P}$, then G contains a set of disjoint paths \mathcal{P}' with $|\mathcal{P}'| = |\mathcal{P}| + 1$. Moreover, the alternating walk W can be chosen such that*

- (1) $E(\mathcal{P}')$ is precisely the symmetric difference $E(\mathcal{P}) \Delta E(W)$, and
- (2) every path $P' \in \mathcal{P}'$ traverses the edges of $E(\mathcal{P}') \cap E(W)$ in the same order as W .

PROOF. The first assertion of the lemma is proved in [54, Lemma 3.3.2]. However, in order to see the moreover-part, we need to recall the main idea of [54, Lemma 3.3.2]. First, the definition of an alternating walk ensures that after taking the symmetric difference $\Delta = E(\mathcal{P}) \Delta E(W) := \bigcup_{P \in \mathcal{P}} E(P) \Delta E(W)$, every vertex outside $A \cup B$ will have degree 0 or 2 in $G[\Delta]$, and vertices in A or B lying on \mathcal{P} of W will continue to have degree 1. Thus, every component of $G[\Delta]$ containing a vertex $a \in A$ will have to be a finite path starting at a and ending at $A \cup B$. To see that any such path in fact has to end at a vertex of B , one proves the additional fact that the path traverses an edge in the symmetric difference always in the forward direction with respect to \mathcal{P} or W . This, clearly, yields the first assertion of the lemma.

However, there might be further components of $G[\Delta]$ which are finite cycles not incident with $A \cup B$. To eliminate the occurrence of such finite cycles (and hence to establish property (1)), we choose an alternating walk W ending in $B \setminus \bigcup \mathcal{P}$ such that $|E(W) \setminus \bigcup_{P \in \mathcal{P}} E(P)|$ is minimal. Now suppose for a contradiction that there is a cycle C in $G[\Delta]$. By the argument above, we know that traversing the edges of C in the forward direction with respect to \mathcal{P} or W induces a cyclic order $<$ on $E(C)$. Observe that since $(C, <)$ is a cyclic order, there are edges e_k, e_{k+i} for $i \geq 0$ of $E(W)$ such that $e_{k+i} < f_1 < f_1 < \dots < f_\ell < e_k$ is a segment of $(E(C), <)$, where $\ell \geq 0$ and f_0, \dots, f_ℓ is a subpath of some path $P \in \mathcal{P}$ (here, $\ell = 0$ allows for the possibility that $e_{k+i} < e_k$ are successors in $(E(C), <)$).

If $\ell > 0$, then the walk $W' = x_0 e_0 \dots e_{k-1} x_k f_\ell \dots f_1 x_{k+i+1} e_{k+i+1} \dots e_{n-1} x_n$ is an alternating walk contradicting the minimality of W . Otherwise, if $\ell = 0$ (so $i > 0$), write x for the vertex incident with both e_{k+i} and e_k . Then by definition of alternating, there is a path $Q \in \mathcal{P}$ with $Q = y_0 f_0 y_1 f_1 \dots f_{m-1} y_m$ such that $e_{k+i+1} = f_r$, $e_{k-1} = f_{r+1}$ and $x = y_{r+1}$. But then the walk $W' = x_0 e_0 \dots e_{k-1} x e_{k+i+1} \dots e_{n-1} x_n$ is an alternating walk contradicting the minimality of W . This contradiction shows that Δ is cycle-free, establishing (1).

It remains to argue that by choosing $|E(W) \setminus \bigcup_{P \in \mathcal{P}} E(P)|$ minimal, we also have property (2). But if (2) fails for some path $P' \in \mathcal{P}'$, there a segment on the path P' of the form $e_{k+i} < f_1 < f_1 < \dots < f_\ell < e_k$ where $\ell \geq 0$ and f_0, \dots, f_ℓ is a subpath of some path $P \in \mathcal{P}$, which yields a contradiction to the minimality of W as before. \square

From Lemmas 11.2.3 and 11.2.4 we immediately deduce:

THEOREM 11.2.5. *Let $G = (V, E)$ be a (potentially infinite) k -connected graph, $A, B \subseteq V$ be disjoint sets of vertices each of size at least k , and \mathcal{P} a collection of disjoint $A - B$ paths with $|\mathcal{P}| = i < k$. Then there exists an alternating $A - B$ walk W such that the symmetric difference $E(\mathcal{P}) \Delta E(W)$ is precisely the edge set of a collection of $i + 1$ many disjoint $A - B$ paths.*

11.2.3. Background on n -arc connectedness.

LEMMA 11.2.6 ([69, Lemma 2.6]). *Let v be a vertex of a graph G of degree at least 3. If x_0, x_1, x_2 are interior points of distinct edges of G incident with v , then any arc α containing $\{x_0, x_1, x_2\}$ satisfies $v \in \text{int}(\alpha)$, and one of its endpoints lies in $[v, x_0] \cup [v, x_1] \cup [v, x_2]$.* \square

The next lemma can be seen as a partial extension of the previous lemma:

LEMMA 11.2.7. *Let G be a graph, and $E(A, B) = \{e_1 = a_1 b_1, \dots, e_n = a_n b_n\}$ be an edge cut of G with $a_i \in A$ and $b_i \in B$. Suppose that $x_i \in e_i$ are interior points.*

Then any arc containing x_1, \dots, x_n with endpoints x_1 and x_n contains either both $[a_1, x_1]$ and $[a_n, x_n]$ or both $[x_1, b_1]$ and $[x_n, b_n]$ if n is even, and it contains $[a_1, x_1]$ and $[x_n, b_n]$ or vice versa if n is odd.

PROOF. From the given arc, fix an embedding $\alpha : [0, 1] \rightarrow G$ such that $\alpha(0) = x_1$, $\alpha(1) = x_n$ and x_2, \dots, x_{n-1} are in $\alpha([0, 1])$. For concreteness let us suppose that n is even and the arc α travels from x_1 along the edge e_1 to a_1 (rather than b_1), and so the given arc contains $[a_1, x_1]$. We show the arc also contains $[a_n, x_n]$. The other cases are similar.

After a_1 , which is in A , the arc α passes through the even number of points x_2, \dots, x_{n-1} in some order, before ending at x_n . As it does so the arc must cross backwards and forwards between A and B an even number of times. Hence α must enter e_n from A , in other words by passing through a_n , and thus the arc contains $[a_n, x_n]$, as claimed. \square

Note that when picking an arc witnessing that points $x_0, x_1, \dots, x_{n-1} \in X$ lie on a common arc, we may assume that endpoints of the arc are among x_0, x_1, \dots, x_{n-1} .

LEMMA 11.2.8. *Let X be a topological space. If there exists $A \subseteq X$ such that $|A| \leq n \in \mathbb{N}$ and $X \setminus A$ has at least $n + 2$ components, then X is not $(n + 2)$ -ac.*

PROOF. Pick $n + 2$ points x_0, x_1, \dots, x_{n+1} each belonging to a distinct component of $X \setminus A$. Suppose there is an arc containing x_0, x_1, \dots, x_{n+1} . Relabeling x_0, x_1, \dots, x_{n+1} if necessary, we can fix an embedding $\alpha : I \rightarrow X$ and $0 = t_0 < t_1 < \dots < t_{n+1} = 1$ such that $x_i = \alpha(t_i)$ for each i . Then $\alpha((t_i, t_{i+1})) \cap A \neq \emptyset$ for each $i = 0, 1, \dots, n$, which is a contradiction since $|A| \leq n$ and α is injective. \square

LEMMA 11.2.9. *Let G be a graph. If for some $n \in \mathbb{N}$, the condition*

$$(\star_n) \quad \text{no } n \text{ points of } |G| \text{ cut } |G| \text{ into at least } n + 2 \text{ components}$$

holds, then (\star_n) implies (\star_m) for all $1 \leq m \leq n$.

PROOF. Let $A \subseteq |G|$ be finite of size $m \geq 1$ witnessing the failure of (\star_m) . Then $|G| \setminus A$ contains at least one half-open edge, i.e. an open set U such that $U \cong (0, 1)$ and $\overline{U} \cong [0, 1)$. By picking $n - m$ many points from U and adding them to the set A , we obtain a set A' witnessing the failure of (\star_n) . \square

Our last lemma in this section says that when verifying whether a graph G is n -ac, it suffices to consider points on the interior of edges of $|G|$.

LEMMA 11.2.10. *For $n \in \mathbb{N}$, a graph G is n -ac if and only if $(|G| \setminus V, |G|)$ is n -ac.*

PROOF. Only the backwards implication requires proof. Assume that $(|G| \setminus V, |G|)$ is n -ac and let $x_0, x_1, \dots, x_n \in G$ be arbitrary. Pick $y_0, y_1, \dots, y_n \in |G| \setminus V$ as follows: if x_i lies on the interior of an edge, then $y_i = x_i$; otherwise, if $x_i \in V$, let y_i be a point on the interior of some edge incident with x_i . By assumption, there is an arc α containing y_0, y_1, \dots, y_n which we may assume to have endpoints y_0 and y_n . Therefore, $x_1, x_2, \dots, x_{n-1} \in \alpha$. We will show that α can be extended to include x_0 and the same argument will work for x_n as well. If $y_0 = x_0$ we are done. If $y_0 \neq x_0$ then x_0 is an endpoint of the edge containing y_0 , say e . The arc α contains one of these endpoint and if $x_0 \in \alpha$ we are done. Otherwise, $\alpha \cup e \cup \{x_0\}$ is also an arc and contains x_0 . \square

11.3. Characterizing n -ac Graphs

11.3.1. Characterizing 2-ac, 2-cc, 3-ac and 3-cc graphs. A graph G is n -strongly arc connected, abbreviated n -sac (see [68]) if for any list of no more than n elements of $|G|$ there is an arc in G containing the points in the specified order. We note that no graph is 4-sac (pick four points x_1, x_3, x_2, x_4 in that order along any edge). It is evident that the following are equivalent for a graph G : (i) G is 2-ac, (ii) G is 2-sac, (iii) $|G|$ is connected, and (iv) G is connected (combinatorially).

It is also clear that a graph is 3-cc if and only if it is a cycle. Indeed a graph G is not 3-cc if (i) it contains a vertex of degree one (that vertex is not in any circle), or (ii) a vertex of degree at least 3 (consider three points from the interior of three edges exiting the vertex), or (iii) is a chain.

We characterize 3-sac and 2-cc graphs. The equivalence of (1) through (4) below for finite graphs was established in [68, Prop. 6].

THEOREM 11.3.1. *For a (possibly infinite) graph G , the following are equivalent:*

- (1) G is 3-sac,
- (2) G is cyclically connected,
- (3) any three points of $|G|$ lie on a circle or a ϑ -curve,
- (4) $G \neq K_2$ is 2-connected, and
- (5) G is 2-cc.

PROOF. The equivalence of (2) \Leftrightarrow (4) follows from Menger's Theorem 11.2.2.

For (2) \Rightarrow (3), pick three points $x_0, x_1, x_2 \in G$. Now for every each 2-element subset A_i of $\{x_0, x_1, x_2\}$, use the fact that G is cyclically connected to find a (finite) cycle $C_i \subseteq G$ containing the two points of A_i . Then consider the finite connected subgraph $H = \bigcup_i C_i$ of G . By construction, any two points of H lie on a cycle, so H is cyclically connected. By the finite case, the three points x_0, x_1, x_2 of H lie on a circle or a ϑ -curve in H , and hence in also G .

The implication (3) \Rightarrow (1) follows from the finite case (see [68, Prop. 6]), and to see (1) \Rightarrow (4), note that if a topological space has a cut-point, then it fails to be 3-sac, [68, Lemma 1].

Finally, evidently 2-cc graphs are cyclically connected, while (3) \Rightarrow (5) since the circle and ϑ -curve are clearly 2-cc. \square

Thus cyclically connected graphs are (strongly) 3-ac, and this extends naturally to a characterization of 3-ac graphs.

THEOREM 11.3.2. *A (potentially infinite) graph G is 3-ac if and only if it is a chain graph of 2-connected links, or, equivalently, if and only if its block graph is connected and contains no vertex of degree at least 3.*

PROOF. To see that the conditions are necessary, consider the block-cutvertex decomposition of G and its associated block graph $B(G)$, which is a (potentially infinite) tree by Lemma 11.2.1. To prove that G is a chain graph of 2-connected links, it suffices to show every vertex of $B(G)$ has degree at most 2. It follows from Lemma 11.2.8 that no cut-vertex of G can have degree strictly bigger than 2 in T . And if there is a block B of G with contains at least three cut vertices c_0, c_1, c_2 , then picking three points x_i each on the interior of edges $e_i \in G \setminus E(B)$ incident with c_i easily shows that G cannot be 3-ac.


For the converse direction, suppose G is a chain graph of 2-connected graphs. Pick x_0, x_1, x_2 in G . Then there is a minimal finite ‘convex’ part of that chain, say L_0, \dots, L_n with $L_i \cap L_{i+1} = v_i$, covering x_0, x_1, x_2 .

If all x_0, x_1, x_2 lie in the same link L_0 , we are done by Theorem 11.3.1, as L_0 is 2-connected. If x_0, x_1 lie in same link, say L_0 , and x_2 lies in L_n , then we may find

- an arc α_0 in L_0 that picks up $\{x_0, x_1, v_0\}$ ending at v_0 (clear if $L_0 = K_2$, and by Theorem 11.3.1 otherwise),
- arcs α_i in L_i with endpoints v_{i-1} and v_i for $0 < i < n$, and
- an arc α_n in L_n with endpoints v_{n-1} and x_3 .

It is then clear that the concatenation of the α_i is an arc though our three points x_0, x_1, x_2 . Finally, in the case where $x_0 \in L_0, x_1 \in L_i$ for $0 < i < n$ and $x_2 \in L_n$, the same approach extends straightforwardly. □

11.3.2. Characterizing 4-ac graphs. Let us say that a graph G is a *basic 4-ac graph* if G is (a subdivision of) a circle, a ϑ -curve, a cycle graph of two circles and an arc (‘happy-face curve’), or if it is (a subdivision of) a cycle graph of alternating two circles and two arcs (‘baguette curve’). See the following sketch for the latter two basic 4-ac graphs.



As any four edges of these graphs either lie on a common ϑ -curve, a figure-8-curve or a dumbbell, these graphs are indeed 4-ac. The purpose of our next theorem is to prove a ‘converse’ of this observation for cyclically connected 4-ac graphs.

THEOREM 11.3.3. *For a cyclically connected graph G , the following are equivalent:*

- (1) G is 4-ac,
- (2) no two vertices cut G into 4 or more components, and
- (3) any four edges of G are contained in a basic 4-ac subgraph of G .

PROOF. The implication (1) \Rightarrow (2) is Lemma 11.2.8, and (3) \Rightarrow (1) is clear.

For (2) \Rightarrow (3), let G be a 4-ac cyclically connected graph such that no two point set cuts it into at least 4 components. Let $x_0, \dots, x_3 \in G$. To show that G is 4-ac, we may assume, by Lemma 11.2.10, that all x_i are interior points of edges.

By Theorem 11.3.1, the three points x_0, x_1, x_2 lie on a common circle or a common ϑ -curve X . In the first case, Menger's Theorem 11.2.2—applied with the two endvertices of the edge containing x_3 against $V(X)$ —shows that there are two disjoint $x_3 - X$ arcs, and so there is a ϑ -curve contain $\{x_0, \dots, x_3\}$ and we are done.

Otherwise, let us write a and b for the two degree-3-vertices of the ϑ -curve X , and e_0, e_1, e_2 for its three edges. Further, as x_0, x_1, x_2 do not lie on a common cycle, we may label the edges of X such that $x_i \in e_i$. Since G is cyclically connected, it follows from Menger's Theorem 11.2.2 as before that there is an arc α such that $x_3 \in \alpha$ and $X \cap \alpha = \{\alpha(0), \alpha(1)\}$. Up to symmetry, the following cases can occur:

- (1) $\alpha(0), \alpha(1) \in e_0$,
- (2) $\alpha(0) \in e_0, \alpha(1) \in e_1$,
- (3) $\alpha(0) = a, \alpha(1) \in e_0$, or
- (4) $\alpha(0) = a, \alpha(1) = b$.

In the first case, $Y = X \cup \alpha$ is homeomorphic to a baguette curve. In the second case, Y is homeomorphic to a K_4 , where removing any edge not containing a point x_i reduces it to a ϑ -curve. In the third case, Y is a happy-face-curve. Finally, in the last case, Y consists of the vertices a and b with four parallel edges e_0, \dots, e_3 between them. Since by assumption, $|G| \setminus \{a, b\}$ consists of at most three components, there is an arc δ in $G \setminus \{a, b\}$ internally disjoint from Y with say $\delta(0) \in e_0$ and $\delta(1) \in e_1$. One checks that any four points in $Z = Y \cup \delta$ lie on either a ϑ -curve or on a happy-face-curve, which completes the proof. \square

Next, we extend our characterization of 4-ac graphs to graphs which are no longer necessarily cyclically connected. For this, the following lemma gives us additional control over arcs in our four basic 4-ac graphs.

LEMMA 11.3.4. *Let G be one of our basic 4-ac graphs. If w is a point in the interior of an edge of G , then for any three further points in G , there exists an arc in G that contains those three points and has w as an endpoint.*

PROOF. If G is either a circle or a ϑ curve, then it is easy to see that the assertion of the lemma holds.

So let G be the happy-face curve with cycles C_1, C_2 , degree-4 vertex $a \in C_1 \cap C_2$, degree-3-vertices $b \in C_1 \setminus C_2$ and $c \in C_2 \setminus C_1$ and edges $\{e_0, e_1\} = E(C_1)$, $\{e_2, e_3\} = E(C_2)$ and $e_4 = bc$. Pick points $w, x_0, x_1, x_2 \in G$. Since removing one of e_0, \dots, e_3 reduces G to a ϑ -curve, we only have to consider the case when one of w, x_0, x_1, x_2 belongs to the interior of each e_i for $0 \leq i \leq 3$. But now, since $G \setminus e_4$ is a figure-8-curve, the assertion of the lemma is clear.

Finally, assume G is the baguette curve with cycles C_1, C_2 , degree-3 vertices $a, b \in C_1$ and $c, d \in C_2$ and edges $\{e_0, e_1\} = E(C_1)$, $\{e_2, e_3\} = E(C_2)$ and $e_4 = ac, e_5 = bd$ between C_1 and C_2 . Pick points $w, x_0, x_1, x_2 \in G$. Since removing one of e_0, \dots, e_3 reduces G to a ϑ -curve, we only have to consider the case when one of w, x_0, x_1, x_2 belongs to the interior of each e_i for $0 \leq i \leq 3$. But now, since $G \setminus e_5$ is a dumbbell with w lying on one of its cycles, the assertion of the lemma is again clear. \square

THEOREM 11.3.5. *A graph G is 4-ac if and only if it is a chain graph such that*

- (1) *all links are 2-connected and 4-ac,*
- (2) *all interior links are edges,*
- (3) *if v is a cut vertex and L a link of G with $v \in L$, then $\deg_L(v) \leq 2$.*

PROOF. Suppose G is a 4-ac graph. Then G is 3-ac, and so a chain graph of 2-connected links by Theorem 11.3.2. Item (1) is now clear.

For (2), suppose for a contradiction, there is a chain graph G of 2-connected links with decomposition $\{L_n : n \in J\}$ for $J \subseteq \mathbb{Z}$ with $|J| \geq 3$ that is 4-ac but one of the interior links, say L_0 , is not an arc. Consider the subgraph $G' = L_{-1} \cup L_0 \cup L_1$ where $L_{-1} \cap L_0 = \{u\}$ and $L_0 \cap L_1 = \{v\}$. Pick x_0 in $L_{-1} \setminus \{u\}$ and x_3 in $L_1 \setminus \{v\}$. Observe that any arc containing x_0 and x_3 and any two further points in L_0 must (without loss of generality) start at x_0 and end at x_3 . So it suffices to show that we can choose x_1, x_2 in $L_0 \setminus \{u, v\}$ so that there is no arc in L_0 starting at u , ending at v and containing both x_1 and x_2 . Consider u in L_0 . By 2-connectedness of L_0 and the fact that L_0 is not an arc, we must have $\deg_{L_0}(u) \geq 2$. Pick x_1 and x_2 from the interior of distinct edges of L_0 incident with u . Now it is clear that no arc in L_0 starting at u and containing x_1 and x_2 , can end at v , a contradiction.

For (3), suppose there are links L_0 and L_1 with $L_0 \cap L_1 = \{v\}$ and $\deg_{L_0}(v) \geq 3$. Then picking three vertices on the interior of different edges incident with v in L_0 , and picking a fourth vertex on the interior of an edge incident with v in L_1 shows that G is not 4-ac, a contradiction.

For the converse, assume that G is a chain graph satisfying properties (1)–(3). We may suppose that G contains a non-trivial link L not isomorphic to K_2 . If the block graph of G is infinite, it follows from (2) that G is isomorphic to L with a one-way infinite ray R attached to a vertex v of L . But then it is clear that it suffices to show that L with a single extra edge attached at v is 4-ac. Thus we may assume, by (1) and the foregoing discussion, that G consists of finite number ≥ 2 of links. So let L_1, \dots, L_n with $n \geq 2$ and $L_i \cap L_{i+1} = \{v_i\}$ be the decomposition of G into links according to properties (1)–(3). Pick four points $x_0, \dots, x_3 \in G$. If all four points are contained in the same link, then we are done by (1). Otherwise, find basic 4-ac subgraphs H and H' in L_1 and L_n containing $v_1 \cup (\{x_0, \dots, x_3\} \cap L_1)$ and $v_{n-1} \cup (\{x_0, \dots, x_3\} \cap L_n)$ respectively. Since by (3), v_1 and v_{n-1} have degree 2 in H and H' respectively, it follows from Lemma 11.3.4 that there are arcs α and α' in H and H' picking up all vertices $\{x_0, \dots, x_3\} \cap L_1$ and $\{x_0, \dots, x_3\} \cap L_n$


and starting at v_1 and v_{n-1} respectively. Since all middle links are arcs by (2), the arc $\alpha \cup L_2 \cup \dots \cup L_{n-1} \cup \alpha'$ witnesses that G is 4-ac. \square

11.3.3. Characterizing 5-ac graphs. Our first theorem reduces the problem of characterizing all 5-ac graphs to the cyclically connected case.

THEOREM 11.3.6. *Let G be a graph which is not cyclically connected. Then G is 5-ac if and only if G is homeomorphic to one of the following graphs:*

(a) a finite path (equivalently, an arc), a ray or a double ray; (b) a lollipop with or without the endpoint; (c) the dumbbell graph, or (d) the figure-of-eight-graph.

See the following diagram for sketches of the lollipop graph, the dumbbell graph, and the figure-of-eight graph.



PROOF. It is straightforward to check that each of the listed graphs is indeed 5-ac. So suppose G is a non-cyclically connected but 5-ac graph. By Theorem 11.3.5, G is a chain graph with multiple links such that all interior links are arcs. If G is not a double ray, then we may suppose that G has a block decomposition $\{L_n : n \in J\}$ where J is an interval in $\{0\} \cup \mathbb{N}$ containing 0. We show that L_0 is either an arc or circle, for then any end-link of the block decomposition of G is a circle or an arc, and all interior links are arcs – and the theorem follows immediately.

Claim: L_0 is either a circle or an arc. Indeed, let v_0 be the linking vertex $L_0 \cap L_1$. By Theorem 11.3.5 (3), if L_0 is not an arc, v_0 has degree 2 in L_0 . Let e_0 and e_1 be the two edges of L_0 incident with v_0 . If L_0 is not a circle, we may suppose without loss of generality that $e_0 = v_0w$ where w has degree 3 in L_0 . Write e_2, e_3 for the other two edges incident with w . Pick four points x_0, \dots, x_3 with x_i on the interior of e_i for $0 \leq i \leq 3$ and pick x_4 on the interior of an edge e_4 in L_1 incident with v_0 . Then these five points witness that G is not 5-ac: By Lemma 11.2.7, any arc would need to start and end on the same side of the edge cut $\{e_0, e_1\} = E(L_0 \setminus \{v_0\}, G \setminus L_0)$, but also needs to start and end in a neighbourhood of w and a neighbourhood of v_0 respectively by Lemma 11.2.6, a contradiction. \square


Thus, we may concentrate on cyclically connected graphs. Here, we have the following characterization.

THEOREM 11.3.7. *A cyclically connected graph G is 5-ac if and only if*

- (1) G has maximum degree 4,
- (2) no 3, or fewer, vertices of G cut $|G|$ into 5 or more components,

- (3) G is not a cycle graph of three (non-trivial) links L_0, L_1, L_2 such that the linking vertex $v \in L_0 \cap L_1$ has both $\deg_{L_0}(v) = 2 = \deg_{L_1}(v)$, and
- (4) G is not the union of three (edge-disjoint) connected subgraphs L_0, L_1, L_2 with two linking vertices v, a such that $L_0 \cap L_1 = L_0 \cap L_2 = L_1 \cap L_2 = \{v, a\}$ and $\deg_{L_2}(v) = 2$.

Note that the combinatorial condition (2) is equivalent to the topological statement (2') 'no 3 points of $|G|$ cut $|G|$ into 5 or more components', and this is what we use below. (To see this equivalence, observe that (2') is automatically stronger than (2), and for the converse, replace any point of $|G|$ in the interior of an edge with one of the vertices at the ends of the edge.) We also remark that the three graphs sketched below witness that even given (1), conditions (2)–(4) are mutually independent. A K_5^- , i.e. a K_5 with one of the edges removed, violates (2) but satisfies (3) and (4). Similarly, the second graph is a non-5-ac graph which fails (3) (as the diagram shows, it is a cycle graph of the type excluded by (3)) but satisfies (2) and (4). Finally, the third graph below satisfies (2) and (3) but not condition (4). (To see that (4) is violated, consider the decomposition as shown in the diagram, where the restriction imposed by (4) on degrees fails.)



We split the proof of Theorem 11.3.7 into two parts. First, in Proposition 11.3.8, we will show that the four conditions listed in the characterization are necessary (replacing (2) with (2') where convenient). In Proposition 11.3.9 further below, we will then show the converse direction.

PROPOSITION 11.3.8. *Any cyclically connected 5-ac graph satisfies properties (1)–(4) above.*

PROOF. For (1), suppose that v is a vertex of degree at least 5, and x_0, \dots, x_4 are chosen from the interior of distinct edges of G incident with v . Suppose for a contradiction that there is an arc α in G through all five points. Applying Lemma 11.2.6 to $[v, x_0] \cup [v, x_1] \cup [v, x_2]$, we know that v is an interior point of α and may assume that one endpoint of α lies say on $(v, x_0]$. Next, applying Lemma 11.2.6 with $[v, x_1] \cup [v, x_2] \cup [v, x_3]$ we may assume that the second endpoint of α lies say on $(v, x_1]$. But applying Lemma 11.2.6 once again with $[v, x_2] \cup [v, x_3] \cup [v, x_4]$, we see that α is forced to have a third endpoint, a contradiction.

Condition (2) follows from Lemma 11.2.8.

For (3), suppose G is a cycle graph of three (non-trivial) links L_0, L_1, L_2 with $L_0 \cap L_1 = \{v\}$, $L_0 \cap L_2 = \{v_1\}$ and $L_1 \cap L_2 = \{v_2\}$ such that the linking vertex $v \in L_0 \cap L_1$ say has both $\deg_{L_0}(v) = 2 = \deg_{L_1}(v)$. Pick points x_0, x_1 on the interior of the distinct edges in L_0 incident with v , points x_2, x_3 on the interior of the distinct edges in L_1 incident with v , and x_4 on the interior of some edge in L_2 . We claim that these five points witness that G cannot be 5-ac. To see this, observe first that Lemma 11.2.6 implies that any potential arc $\alpha: [0, 1] \rightarrow G$ containing x_0, \dots, x_4 has to start and end inside $[v, x_0] \cup [v, x_1] \cup [v, x_2] \cup [v, x_3]$. In particular, x_4 lies on the interior of α , and so also v_1 and v_2 lie on the interior of α . Without loss of generality, let $0 < t_1 < t_2 < 1$ be the points where $\alpha(t_i) = v_1$ for $i = 1, 2$. Now following the arc $\alpha \upharpoonright [0, t_1]$ backwards in time, we will first encounter say x_0 at time $0 \leq s_0 < t_1$. Similarly, following the arc $\alpha \upharpoonright [t_2, 1]$ forwards in time, we will first encounter say x_2 at time $t_2 < s_2 \leq 1$. Now, however, the points x_0, \dots, x_3 are contained in the space $Y = G \setminus \alpha \upharpoonright (s_0, s_2)$. But v is a 4-cut point of Y with all x_i contained in different components of $Y - v$. As in Lemma 11.2.8, it follows that the set $\{x_1, \dots, x_4\}$ cannot be covered by the two disjoint arcs $\alpha \upharpoonright [0, s_0]$ and $\alpha \upharpoonright [s_2, 1]$, a contradiction.

For (4), the argument is somewhat similar to the previous case. Pick points x_0, x_1 on the interior of the edges incident with v in L_0 and L_1 respectively, points x_2, x_3 on the interior of the distinct edges in L_2 incident with v , and x_4 on the interior of some edge $e = ab$ in L_2 incident with a (where, without loss of generality, we assume that $b \neq v$). We claim that these five points witness that G cannot be 5-ac. To see this, observe first that Lemma 11.2.6 implies that any potential arc $\alpha: [0, 1] \rightarrow G$ containing x_0, \dots, x_4 has to start and end inside $[v, x_0] \cup [v, x_1] \cup [v, x_2] \cup [v, x_3]$. In particular, x_4 lies on the interior of α , and so also a and b lie on the interior of α , too. Without loss of generality, let $0 < t_1 < t_2 < 1$ be the points where $\alpha(t_i) = v_1$ for $i = 1, 2$. Now following the arc $\alpha \upharpoonright [0, t_1]$ backwards in time, if α continues in L_0 or in L_1 , we can argue similar to the previous case. If, however, α stays in L_2 , then without loss of generality there are s_2, s_3 with $0 < s_2 < t_1 < t_2 < s_3 < 1$ with $\alpha(s_2) = x_2$ and $\alpha(s_3) = x_3$, and again we can argue that v is a 4-cut point of $Y = G \setminus \alpha \upharpoonright (s_0, s_2)$ with all x_i contained in different components of $Y - v$, and we get a contradiction as before. \square

PROPOSITION 11.3.9. *Let G be a cyclically connected graph such that*

- (1) G has maximum degree 4,
- (2) no 3 points of $|G|$ cut $|G|$ into 5 or more components,
- (3) G is not a cycle graph of three (non-trivial) links L_0, L_1, L_2 such that the linking vertex $v \in L_0 \cap L_1$ has both $\deg_{L_0}(v) = 2 = \deg_{L_1}(v)$, and
- (4) G is not the union of three (edge-disjoint) connected subgraphs L_0, L_1, L_2 with two linking vertices v, a such that $L_0 \cap L_1 = L_0 \cap L_2 = L_1 \cap L_2 = \{v, a\}$ and $\deg_{L_2}(v) = 2$.

Then $|G|$ is 5-ac.

PROOF. Consider $x_0, \dots, x_4 \in |G|$. By Lemma 11.2.10, to show that G is 5-ac, it suffices to consider points x_i which lie on the interior of edges of G . Applying Lemma 11.2.9 and Theorem 11.3.3, we see that condition (2) implies in particular that x_0, \dots, x_3 lie on a basic 4-ac space X , i.e. either on a cycle, a ϑ -curve, a baguette- or a happy-face-curve. Using 2-connectedness, we may connect x_4 to X via two internally disjoint paths α_1, α_2 , i.e. paths with $\alpha_i(0) = x_4$, $\alpha_i(1) \in X$, and $\alpha \upharpoonright [0, 1) \cap X = \emptyset$. If X was a cycle, then all five points lie on a ϑ -curve, so in particular they lie on a common arc, and we are done. Thus, only the three cases remain where X is a ϑ -curve, a baguette-curve, or a happy-face-curve. We now analyze each case separately.

Case 1. X a ϑ -curve.

Write a and b for the two degree-3-vertices of the ϑ -curve, and e_0, e_1, e_2 for the three edges of the ϑ -curve. As we may suppose that no 4 vertices of $\{x_0, \dots, x_4\}$ lie on a cycle, we may label the edges of our ϑ -curve such that $x_0 \in e_0$, $x_1 \in e_1$ and $x_2, x_3 \in e_2$. Since vertices in G have degree at most 4, the following cases for how that arcs α_1 and α_2 connect up to X can occur:

- (i) $\alpha_1(1) = a$ and $\alpha_2(1) = b$,
- (ii) $\alpha_1(1) = a$ and $\alpha_2(1) \in e_i$, or
- (iii) both α_1 and α_2 hit X on interior points of edges.


In (iii), either $\alpha_1(1)$ and $\alpha_2(1)$ lie on the same edge $e_i \subseteq X$, in which case $Y = X \cup \alpha_1 \cup \alpha_2$ is a baguette-curve containing $\{x_0, \dots, x_4\}$, or, by symmetry, we may assume that $\alpha_1(1) \in (a, x_0) \subseteq e_0$ and $\alpha_2(1) \in e_1 \cup e_2$, in which case $Y = (X \setminus (a, \alpha_1(1))) \cup \alpha_1 \cup \alpha_2$ is a ϑ -curve containing $\{x_0, \dots, x_4\}$. In both cases, our five points x_0, \dots, x_5 lie on a 5-ac subspace, and we are done.

Next, we claim that—similar to the proof of Theorem 11.3.3—case (i) reduces to case (ii). Indeed, suppose that $\alpha_1(1) = a$ and $\alpha_2(1) = b$. Then $Y = X \cup \alpha_1 \cup \alpha_2$ is a graph with vertices a and b and four parallel edges e_0, \dots, e_3 with $x_0 \in e_0$, $x_1 \in e_1$, $x_2, x_3 \in e_2$ and $x_4 \in e_3$. By assumption (2) and Lemma 11.2.9, the points a and b do not cut G into 4 or more components, and hence there is an arc δ in G between two different edges of Y . By symmetry, we may assume that $\delta(0) \in (a, x_0) \subseteq e_0$. But then $Y \setminus (a, \delta(0))$ is homeomorphic to a ϑ -curve $X' = \{a, b\} \cup e_1 \cup e_2 \cup e_3$ with the point x_0 joined to X' via two arcs attaching to $\delta(1)$ and b , i.e. the configuration of subcase (ii).

Thus, it remains to work through case (ii). By symmetry, the following possibilities can occur:

$$\begin{aligned} \alpha_2(1) \in (a, x_0) \subseteq e_0 \subseteq X, \quad \alpha_2(1) \in (x_0, b) \subseteq e_0 \subseteq X, \quad \alpha_2(1) \in (a, x_2) \subseteq e_2 \subseteq X, \\ \alpha_2(1) \in (x_2, x_3) \subseteq e_2 \subseteq X, \quad \text{or} \quad \alpha_2(1) \in (x_3, b) \subseteq e_2 \subseteq X. \end{aligned}$$

Write $Y = X \cup \alpha_1 \cup \alpha_2$. In the first and third case, $Y \setminus (a, \alpha_2(1))$ is a ϑ -curve containing $\{x_0, \dots, x_4\}$, and in the second and fourth case, $Y \setminus (\alpha_2(1), b)$ is a figure-8-curve containing $\{x_0, \dots, x_4\}$. Thus, in the first four cases, our five points x_0, \dots, x_5 lie on a common ω -ac subspace, and we are done. In the fifth case, we see that Y is a happy-face curve, i.e. a cycle graph consisting of two cycles and an arc with a point x_i on every single edge. For convenience, let us relabel all edges and points of Y as in the picture.



By condition 2, the three points a, b, t do not disconnect G into 5 or more components, and therefore there is an arc δ internally disjoint from Y connecting some pair of edges e_i and e_j for $i \neq j$. Again we differentiate several subcases (up to symmetry) depending on the attaching points of δ .

- (a) $\delta(0) \in (a, x_1)$ and $\delta(1) \in (a, x_2)$,
- (b) $\delta(0) \in (a, x_1)$ and $\delta(1) \in (x_2, t)$,
- (c) $\delta(0) \in (a, x_2)$ and $\delta(1) \in (x_3, t)$,
- (d) $\delta(0) \in (a, x_2)$ and $\delta(1) \in (b, x_3)$,
- (e) $\delta(0) \in (x_2, t)$ and $\delta(1) \in (x_3, t)$,
- (f) $\delta(0) \in (a, x_2)$ and $\delta(1) \in (a, x_5)$,
- (g) $\delta(0) \in (a, x_2)$ and $\delta(1) \in (x_5, b)$,
- (h) $\delta(0) \in (x_2, t)$ and $\delta(1) \in (a, x_5)$,
- (i) $\delta(0) \in (x_2, t)$ and $\delta(1) \in (x_5, b)$.

Note in case (a), for example, we additionally know that $(a, x_1) \subseteq e_1$ and $(a, x_2) \subseteq e_2$, and similarly for all the cases. Write $Z = Y \cup \delta$. Now in (b) and (c), $\{x_0, \dots, x_4\}$ lie on the common ϑ -curve $Z \setminus ((a, \delta(0)) \cup (\delta(1), t))$ respectively. In (d), $Z \setminus ((a, \delta(0)) \cup (b, \delta(1)))$ is a figure-8-curve containing $\{x_0, \dots, x_4\}$. In (e), $Z \setminus ((\delta(0), t) \cup (\delta(1), t))$ is a figure-8-curve containing $\{x_0, \dots, x_4\}$. In (g), $\{x_0, \dots, x_4\}$ lie on the common figure-8-curve $Z \setminus ((a, \delta(0)) \cup (\delta(1), b))$. In (h), $\{x_0, \dots, x_4\}$ lie on the common ϑ -curve $Z \setminus ((\delta(0), t) \cup (a, \delta(1)))$. And in (i), $Z \setminus ((\delta(0), t) \cup (\delta(1), b))$ is a dumbbell containing $\{x_0, \dots, x_4\}$.

Thus, it remains to check cases (a) and (f). Note first these cases are isomorphic (after relabeling $a := \delta(1)$ in (f), and so forth). So without loss of generality, we may assume we are in case (a). Note that by assumption 3, there must exist some additional arcs connecting different parts of the subgraph. Let us work in the (connected) space $G' = |G| \setminus \{t\}$. First, assume there is no cut-vertex separating $A = e_1 \cup \{a\} \cup e_2$ from $B = e_3 \cup \{b\} \cup e_4$ in G' (in particular, we assume b is not such a cut-vertex). Then by Menger, there exists an $A - B$ walk β in $G - t$ which is alternating with respect to \bar{e}_5 such that $a, b \notin \beta$.

Claim 1: *If $\beta \cap (x_5, b) \neq \emptyset$, then we are done.*

To see the claim, let $t \in (0, 1)$ be minimal such that $\beta(t) \in e_5$. Since we have excluded (h) and (i) above, we may assume that $\beta(0) \in (a, x_2) \subseteq e_2$. Next, let $t' \in (0, 1)$ be minimal such that $x = \beta(t') \in (x_5, b)$, and consider the arc $\beta' = \beta \upharpoonright [0, t']$. Then β' is an $A - x$ walk disjoint from B and alternating with respect to $[a, x] \subseteq \bar{e}_5$ such that

$x_5 \notin \beta'$ (this follows from the definition of ‘alternating’), and so we find two independent $A - x$ -paths γ_1 and γ_2 in the symmetric difference of β' and $[a, x]$ with starting vertices a and $\beta(0)$ respectively such that $x_5 \in \gamma_i$ for precisely one i , see Lemma 11.2.4. But then $[Y \setminus ((a, \beta(0)) \cup e_5)] \cup \gamma_1 \cup \gamma_2$ is a figure-8-curve containing $\{x_0, \dots, x_4\}$.

Claim 2: *If $\beta \cap (x_5, b) = \emptyset$, then we are also done.*

To see this, note that $\beta \cap (x_5, b) = \emptyset$ implies, by the definition of alternating, that $x_5 \notin \beta$. Therefore, by taking the symmetric sum of \bar{e}_5 and β , we obtain two disjoint $A - x$ -paths γ_1 and γ_2 with starting vertices a and $\beta(0)$ respectively such that $x \in \gamma_i$ for precisely one i (if we choose the walk β according to the moreover-part of Lemma 11.2.4). Next, since we have excluded cases (c), (d) and (e) above, we may assume that $\beta \cap (a, x_5) \neq \emptyset$, and since we have excluded cases (h) and (i) above, we may further assume that $\beta(0) \in (a, x_2) \subseteq e_2$. Thus, up to symmetry, the following four arrangements can occur:

- | | |
|--|--|
| (1) $\gamma_1(1) = b, \gamma_2(1) \in (b, x_3) \subseteq e_3,$ | (3) $\gamma_2(1) = b, \gamma_1(1) \in (b, x_3) \subseteq e_3,$ |
| (2) $\gamma_1(1) = b, \gamma_2(1) \in (x_3, t) \subseteq e_3,$ | (4) $\gamma_2(1) = b, \gamma_1(1) \in (x_3, t) \subseteq e_3.$ |

In the first case, $(Y \setminus [e_5 \cup (a, \gamma_2(0)) \cup (b, \gamma_2(1))]) \cup \gamma_1 \cup \gamma_2$ is a figure-8-curve containing $\{x_0, \dots, x_4\}$. In the second case, $(Y \setminus [e_5 \cup (a, \gamma_2(0)) \cup (\gamma_2(1), t)]) \cup \gamma_1 \cup \gamma_2$ is a ϑ -curve containing $\{x_0, \dots, x_4\}$. In the third case, $(Y \setminus [e_5 \cup (a, \gamma_2(0)) \cup (b, \gamma_1(1))]) \cup \gamma_1 \cup \gamma_2$ is a figure-8-curve containing $\{x_0, \dots, x_4\}$. And in the last case, the subgraph $(Y \setminus [e_5 \cup (a, \gamma_2(0)) \cup (\gamma_1(1), t)]) \cup \gamma_1 \cup \gamma_2$ is a ϑ -curve containing $\{x_0, \dots, x_4\}$.

This completes the case checks for when there was no cut-vertex between A and B in G' . So now, we may assume that some vertex $v \in e_5 \cup \{b\}$ is a cut-vertex of G' . Without loss of generality, v is chosen as close to a on \bar{e}_5 as possible.

Claim 3: *If $v \in (a, x_5)$, then we are done.*

Indeed, the existence of a further cut point v' separating x_5 from B in $G - \{t, v\}$ would contradict condition (3). Therefore, by Menger (cf. Corollary 11.2.5), there exists an $x_5 - B$ walk β in $G - \{t, v\}$ which is alternating with respect to $[x_5, b] \subseteq e_5$. Write z for the endpoint of β on $e_3 \cup e_4$. By the excluded cases (h) and (i) above, we may assume that $\beta(1) \in (b, x_3) \subseteq e_3$. By taking the symmetric difference of $[x_5, b]$ and β , we see as in Claim 1 above that our set $\{x_0, \dots, x_4\}$ lies on a figure-8-curve.

Claim 4: *If $v \in (x_5, b)$, then then we are also done.* This case follows as in Claim 1 (using the fact that v was chosen left-most).

Claim 5: *Can deal with the case $v = b$.*

Again, since v was chosen left-most, our paths β and \bar{e}_5 witness that b has degree 2 in $L_2 := G \setminus (e_3 \cup e_4)$. Now at this point, $\{t, b, w\}$ would give rise to a decomposition of G into

$$L_0 = \bar{e}_3, \quad L_1 = \bar{e}_4, \quad \text{and} \quad L_2 = G \setminus (e_3 \cup e_4),$$

contradicting condition (4). Therefore, since $v = b$ was assumed to be a cut-vertex, we are forced to conclude there there must be an additional arc δ' between e_3 and e_4 .

As we have excluded (b) and (c) above, we may assume that $\delta'(0) \in (b, x_3) \subseteq e_3$ and $\delta'(1) \in (b, x_4) \subseteq e_4$.

Finally, since we have dealt with Claim 1 already, we may assume that $x_5 \notin \beta$. Taking the symmetric difference of β and \bar{e}_5 gives us two disjoint $A - b$ paths γ_1 and γ_2 such that $x \in \gamma_i$ for precisely one i (again assuming that we choose the walk β according to the moreover part of Lemma 11.2.4), with say $\gamma_1(0) = a$ and $\gamma_2(0) \in (a, x_2)$. But then $[Y \setminus (e_5 \cup (a, \gamma_2(0)) \cup (b, \delta'(0)) \cup (b, \delta'(1)))] \cup \gamma_1 \cup \gamma_2$ is a figure-8-curve containing $\{x_0, \dots, x_4\}$, and we are done.

Case 2. X a baguette-curve.

This case is fairly easy in comparison. If X is a baguette curve with cycles C_1, C_2 , degree-3-vertices $a, b \in C_1$ and $c, d \in C_2$ and edges $\{e_0, e_1\} = E(C_1)$, $\{e_2, e_3\} = E(C_2)$ and $e_4 = ac, e_5 = bd$ between C_1 and C_2 , we may assume that $x_i \in e_i$ for $0 \leq i \leq 3$, as otherwise we are back in the ϑ -curve case. Now consider where the arcs α_1 and α_2 attaching x_4 hit X . Note first that if say $\alpha_i(1) \in C_i$, then $\{x_0, \dots, x_4\}$ lie on a common dumbbell, and we are done. Thus, up to symmetry, the following cases remain:

- | | |
|--|--|
| (a) $\alpha_1(1) \in e_3 \cup e_4$, | (c) $\alpha_1(1) \in e_4, \alpha_2(1) \in e_5 \cup \{d\}$, or |
| (b) $\alpha_1(1) \in e_4, \alpha_2(1) = c$, | (d) $\alpha_1(1) = c, \alpha_2(1) = d$. |

In case (a), we may assume by symmetry that $\alpha_1(1) \in (c, x_3)$. Then $(X \cup \alpha_1) \setminus ((c, \alpha_1(1)) \cup e_5)$ is a lollipop containing $\{x_0, \dots, x_4\}$. Next, let $Y = X \cup \alpha_1 \cup \alpha_2$. In case (b), $Y \setminus (\alpha_1(1), \alpha_2(1))$ is a baguette-curve containing $\{x_0, \dots, x_4\}$. In case (c), $Y \setminus (\alpha_1(1), c) \cup (b, \alpha_2(1))$ is a lollipop containing $\{x_0, \dots, x_4\}$. Thus, it remains to analyze case (d) more closely. In this case, the subgraph $Y = X \cup \alpha_1 \cup \alpha_2$ is the baguette curve X with an extra edge e_6 with endpoints c and d . In particular, removing the vertices c and d from $|Y|$ would leave 4 connected components.

Thus, using the condition that no two vertices split $|G|$ into 4 different components (by condition (2) and Lemma 11.2.9), we know that there must be a further arc δ internally disjoint from Y and connecting different components of $|Y| \setminus \{a, b\}$. Up to symmetry (as there is no structural difference between e_2, e_3 and e_6), we may assume that $\delta(0) \in (x_2, d) \subseteq e_2$. Then for the other endpoint of δ , the following cases can occur:

- | | | |
|----------------------------|---|---|
| (i) $\delta(1) \in e_4$, | (iv) $\delta(1) = b$, | (vii) $\delta(1) \in (a, x_0) \subseteq e_0$, |
| (ii) $\delta(1) \in e_5$, | (v) $\delta(1) \in (x_3, d) \subseteq e_3$, | (viii) $\delta(1) \in (x_0, b) \subseteq e_0$. |
| (iii) $\delta(1) = a$, | (vi) $\delta(1) \in (c, x_3) \subseteq e_3$, | |

In all cases, it is straightforward to see to verify that our five points x_0, \dots, x_4 lie on a common dumbbell. This completes the proof of Case 2.

Case 3. X a happy-face-curve.

Again, this case is fairly easy in comparison. If X is a happy face curve with cycles C_1, C_2 , degree-4 vertex $a \in C_1 \cap C_2$, degree-3-vertices $b \in C_1 \setminus C_2$ and $c \in C_2 \setminus C_1$ and

edges $\{e_0, e_1\} = E(C_1)$, $\{e_2, e_3\} = E(C_2)$ and $e_4 = bc$, we may assume that $x_i \in e_i$ for $0 \leq i \leq 3$, as otherwise we are back in the ϑ -curve case. Now consider where the arcs α_1 and α_2 attaching x_4 hit X . Note that the α_i cannot hit on any x_j (as they were chosen to lie on the interior of edges of G), nor on the center vertex a , by condition (1).

If α_1 and α_2 hit the same segment of $C_i \setminus \{a, x_j, x_k\}$, then ignoring the edge e_4 , we see that all our 5 points lie on a figure-8-curve.


Next, if α_1 hits C_1 say, and α_2 doesn't, then it's easy to see that we are back in the discussion as in Case 1, where all our five points lie on the different edges of a happy face curve, so we are done, as we have solved this arrangement already.

Lastly, we assume that α_1 and α_2 hit different segments of say $C_1 \setminus \{a, x_0, x_1\}$. Let us view C_1 as a cycle ax_0fx_1ga with vertices a, x_0, x_1 and three edges. After removing the edge e_5 , we see that up to symmetry, the following three cases can occur: (i) $\alpha_1(1) \in e$ and $\alpha_2(1) \in f$, (ii) $\alpha_1(1) \in e$ and $\alpha_2(1) \in g$, or (iii) $\alpha_1(1) \in f$ and $\alpha_2(1) \in g$. In all three cases, we see that $X \setminus (e_5 \cup (a, \alpha_1(1)))$ is a dumbbell containing our five points x_0, \dots, x_4 . This completes the proof. \square

11.3.4. Characterizing 6-ac graphs. Our characterization of 6-ac graphs, the main result of this section, is as follows.

THEOREM 11.3.10. *A graph G is 6-ac if and only if either G is one of the nine 7-ac graphs of Theorem 11.3.15 or, after suppressing all degree-2-vertices, the graph G is 3-regular, 3-connected, and removing any 6 edges does not disconnect G into 4 or more components.¹*

Note that the last condition in particular implies that G must be triangle-free. However, the stronger condition we chose is necessary for the characterization, as demonstrated by the following 3-regular 3-connected, triangle-free graphs, which both fail to be 6-ac (in both cases consider the six points labeled \blacktriangleright).



¹Equivalently, if G is 3-regular, 3-connected and not an *inflated* K_4 : there is no partition of $V(G)$ into four non-empty subsets V_1, \dots, V_4 such that each $G_i = G[V_i]$ is connected and there is precisely one $G_i - G_j$ edge in G for every pair $i \neq j$.

We split the proof of Theorem 11.3.10 into two parts. First, in Proposition 11.3.11, we will show that the three conditions mentioned in the characterization are necessary. In Proposition 11.3.14 further below, we will then show the converse direction.

PROPOSITION 11.3.11. *Let G be a 6-ac graph which different from the nine 7-ac graphs. Then G is 3-regular, 3-connected, and removing any 6 edges does not disconnect G into 4 or more components.*

PROOF. Let G be a 6-ac graph which different from the nine 7-ac graphs.

To see that G is 3-regular, note that G contains no vertices of degree 1, since Theorem 11.3.6 implies that G is cyclically connected. We suppress vertices of degree 2. Suppose for a contradiction that v is a vertex of G of degree ≥ 4 . Since G is cyclically connected, it follows that G must have another branch point. Then one of the edges incident with v have a branch point as its other endpoint, say u . Let this edge be e . Pick arcs $\alpha_1, \alpha_2, \alpha_3$ interior-disjoint from each other and e with $\{v\} = \alpha_i \cap \alpha_j$, such that for each i , $\alpha_i \setminus \{v\}$ contains no branch points of G . Also pick arcs β_1, β_2 interior-disjoint from the α_i , each other and e , $\beta_1 \cap \beta_2 = \{u\}$ and $\beta_i \setminus \{u\}$ contain no branch points of G . Pick one point from the interior of each of α_i, β_j and e . Then, by Lemma 11.2.6, there is no arc going through these points.

To see that G is 3-connected, it suffices to show, since G is 3-regular, that it is 3-edge connected, i.e. that there is no partition $V(G) = A \dot{\cup} B$ with $|E(A, B)| \leq 2$. Note that cyclical connectedness implies that $|E(A, B)| \geq 2$. So suppose for a contradiction that there is a 2-edge cut $E(A, B) = \{e_1, e_2\}$. Let $e_i = a_i b_i$ with $a_i \in A$ and $b_i \in B$. Note that since G is cyclically connected and 3-regular, all four endpoints of e_1 and e_2 are distinct. In particular, a_1 is incident with two further edges e_3, e_4 which both have all their endpoints in A , and b_2 is with two further edges e_5, e_6 which both have all their endpoints in B . Pick six points $x_i \in e_i$. Since any arc α picking up x_1 and x_2 has to have, without loss of generality, both its endpoints on the A -side of $G \setminus \{x_1, x_2\}$ by Lemma 11.2.7, it follows that it cannot pick up x_5 and x_6 without violating Lemma 11.2.6.

Finally, suppose deleting edges e_1, \dots, e_6 from G leaves components C_1, \dots, C_k . We claim that $k \leq 3$. First, observe that every edge e_i is incident with at most 2 different components, and by 3-connectedness, every component C_i is incident with at least 3 distinct edges. By double counting, it follows $k \leq 4$.

So assume that $k = 4$. Then every component must be incident with precisely 3 of the 6 edges. We claim that the four components and the 6 edges are arranged like a K_4 . For this, it suffices to show that for any two components there is only one edge incident with both components. If there were two components that share three incident edges, then G would be disconnected, a contradiction. And if there are two components that share two incident edges, then the other two components must also share two further incident

edges, from which we conclude that the remaining two edges form a disconnection of the G , contradicting once again 3-connectedness.

Thus, the 4 components together with the 6 edges are arranged like a K_4 . But then it follows from Lemma 11.2.7 that if we choose an interior point x_i on each of the six edges e_i for $1 \leq i \leq 6$, there is no arc α in the graph picking up these 6 points. Indeed, suppose that the arc α starts at x_1 , traverse x_2 up to x_5 in the given order and ends and x_6 . Write v for the first vertex on α and assume $v \in V(C_1)$.

If e_6 is not incident with C_1 , consider the cut $E(C_1, G \setminus C_1) = \{e_1, e_i, e_j\}$ of G with $1 < i < j < 6$. Let $\beta := \alpha \upharpoonright [0, \alpha^{-1}(x_j)]$ and $\gamma = \alpha \upharpoonright [\alpha^{-1}(x_j), 1]$ denote the subarcs of α from x_1 to x_j and from x_j to x_6 respectively. By Lemma 11.2.7, it follows that $[x_j, w]$ with $w \in C_j \subseteq \beta$ is the final segment of β . Pick $y \in (x_j, w)$. Then $\{x_1, x_i, y\}$ is a separation of G separating x_j from x_6 , contradicting the fact that γ is an arc in $G \setminus \{x_1, x_i, y\}$ between these very two points. Finally, if e_6 is also incident with C_1 , then say C_2 is incident with edges e_i, e_j, e_ℓ with $1 < i < j < \ell < 6$. Considering the arcs $\beta := \alpha \upharpoonright [\alpha^{-1}(x_i), \alpha^{-1}(x_\ell)]$ and $\gamma = \alpha \upharpoonright [\alpha^{-1}(x_\ell), 1]$, we may arrive at a similar contradiction as before. \square

Before we start proving the converse, we need the following two lemmas. Note also that the properties 3-connected and 3-regular imply that our graph is simple, i.e. (even after suppressing all degree-2-vertices) it contains no loops or parallel edges.

LEMMA 11.3.12. *Any four points of a 3-regular, 3-connected graph lie on a circle or a ϑ -curve.*

PROOF. Let G be a 3-regular, 3-connected graph. It is easy to check that 3-regularity and 2-connectedness imply that any 4 points x_1, \dots, x_4 of $|G|$ lie on a circle, a theta curve, or a baguette curve.

In the first two cases, we are done, so it remains to show that if our four vertices lie on a baguette curve, they also lie on a ϑ -curve. Let C_1 and C_2 be the two cycles of the baguette curve. Note that we may assume that x_1, x_2 lie on C_1 and x_3, x_4 on C_2 . Now by Menger's theorem (using 3-connectedness of G and the fact that $|V(C_i)| \geq 3$), there are 3 vertex disjoint paths $\alpha_1, \alpha_2, \alpha_3$ from C_1 to C_2 , each meeting $C_1 \cup C_2$ only in their endpoints. Note that $C_1 \setminus \{x_1, x_2\}$ consists of two segments, so one of these segments meets both say α_1 and α_2 . But then the cycle C_2 together with α_1 , then walking around C_1 picking up x_1 and x_2 , and then following back along α_2 gives us a ϑ -curve containing the four points x_1, \dots, x_4 . \square

LEMMA 11.3.13. *Any five points of a 3-regular, 3-connected graph lie on a circle or a ϑ -curve.*

PROOF. Let G be a 3-regular, 3-connected graph and consider five points x_1, \dots, x_5 of $|G|$. If any four of them lie on a circle, then we are done.

Thus, by the previous lemma, we may assume that x_1, \dots, x_4 lie on a ϑ -curve with edges e_1, e_2, e_3 and vertices a and b . By symmetry, we may assume that $x_1, x_2 \in e_1$, $x_3 \in e_2$ and $x_4 \in e_3$. Connect the last point x_5 to the ϑ -curve via two new independent arcs α_1 and α_2 . Since G is 3-regular, the two arcs α_1 and α_2 cannot hit the ϑ -curve in a or b . If the two arcs connect to different edges of the ϑ -curve, then in particular either e_2 or e_3 is hit, and by deleting a suitable part of e_2 or e_3 not containing x_3 or x_4 we have found a ϑ -curve containing x_1, \dots, x_5 . Thus, we may assume that the two arcs hit the same edge e_i , and then we have found a baguette curve of G containing all five points x_1, \dots, x_5 . We will show that in this case, they also lie on a ϑ -curve.

Let C_1 and C_2 be the two cycles of the baguette curve. Up to symmetry, the following cases can occur:

- (1) $x_4, x_5 \notin C_1 \cup C_2$, (2) $x_1, x_2, x_3 \in C_1$, $x_4 \in C_2$ and $x_5 \notin C_1 \cup C_2$,
- (3) $x_1, x_2 \in C_1$ and $x_3, x_4, x_5 \in C_2$, or
- (4) $x_1, x_2 \in C_1$, $x_3, x_4 \in C_2$ and $x_5 \notin C_1 \cup C_2$,

In case (1), if two vertices lie outside of $C_1 \cup C_2$, then it's easy to find a circle inside the baguette curve containing four of the vertices. In case (2), we may again find a circle inside the baguette curve picking up x_4, x_5 and two of the remaining three vertices on C_1 .

In case (3) we follow a strategy similar to the previous lemma. By Menger and 3-connectedness, there are 3 vertex disjoint paths $\alpha_1, \alpha_2, \alpha_3$ from C_1 to C_2 , each meeting $C_1 \cup C_2$ only in their endpoints. Note that $C_1 \setminus \{x_1, x_2\}$ has two components, so one of these segments meets say α_1 and α_2 . But then we can follow α_1 , then walking around C_1 picking up x_1 and x_2 , and then following α_2 and then walk around C_2 back to the endpoint of α_1 in the correct direction so as to pick up two out of the three vertices on C_2 . So we have found four points on a circle.

In case (4), let us denote by β the $C_1 - C_2$ -edge of our baguette curve containing x_5 . As before, by Menger and 3-connectedness, there are three vertex disjoint $C_1 - C_2$ paths α_1, α_2 and α_3 .

Subcase (4a). If it is possible to choose arcs $\alpha_1, \alpha_2, \alpha_3$ such that one of them contains x_5 , then we do so. Assume that α_1 contains x_5 . If a second path say α_2 hits $C_1 \setminus \{x_1, x_2\}$ in the same segment as α_1 , then first using α_1 , then picking up x_1, x_2 on C_1 , then using α_2 , and then returning to α_1 on C_2 picking up at least one more point say x_4 gives a circle containing four of our points, and we are done. Otherwise, by symmetry and pigeon hole principle, we may assume that α_2 and α_3 both hit $C_1 \setminus \{x_1, x_2\}$ as well as $C_2 \setminus \{x_3, x_4\}$ in the same segments, and so it is easy finding a circle containing x_1, \dots, x_4 and we are again done.

Subcase (4b). No path system between C_1 and C_2 contains x_5 . By construction (and the fact that we have excluded subcase 3a) there is a subarc $\beta' \subseteq \beta$ such that $x_5 \in \beta'$ and say $\beta'(0) \in \alpha_1$, $\beta'(1) \in \alpha_2$ and which is otherwise disjoint from $C_1 \cup C_2 \cup \alpha_1 \cup \alpha_2 \cup \alpha_3$.

Now if say α_2 hits $C_1 \setminus \{x_1, x_2\}$ in the same segment as α_3 , then by following α_3 , picking up x_1, x_2 on C_1 , then following along α_2 until we can turn into β' to pick up x_5 , and then following α_1 into C_2 , and back to the beginning of α_3 picking up one more point say x_4 on C_2 , we have found a circle containing four of our points, and are done. Otherwise, by symmetry and pigeon hole principle, we may assume that α_1 and α_2 both hit $C_1 \setminus \{x_1, x_2\}$ as well as $C_2 \setminus \{x_3, x_4\}$ in the same segments, and so it is easy finding a circle containing x_1, \dots, x_4 and we are again done. \square

We are now ready to prove the converse direction of our main characterization theorem.

PROPOSITION 11.3.14. *Let G be a simple 3-regular, 3-connected graph such that removing any 6 edges does not disconnect G into 4 or more components. Then G is 6-ac.*

PROOF. Pick six points x_1, \dots, x_6 from G which we may assume, by Lemma 11.2.10, to be interior points of edges. By Lemma 11.3.13, there is a ϑ -curve Θ containing the first five points x_1, \dots, x_5 . Write e, f, g for the edges of Θ and a, b , for the vertices of Θ . We may assume that every edge of e, f, g is incident with a point x_i , and so up to symmetry there are two cases to consider, namely

- (A) $x_1 < x_2 < x_3 \in e$ (ordered from a to b), $x_4 \in f$ and $x_5 \in g$, or
- (B) $x_1, x_2 \in e$, $x_3, x_4 \in f$ and $x_5 \in g$.

We may assume that $x_6 \notin \Theta$. Pick two independent $x_6 - \Theta$ arcs α_1 and α_2 . By 3-regularity, the arcs cannot hit Θ in a or b .

In case (A), if one of the arcs hits Θ on a segment of $\Theta \setminus \{x_1, \dots, x_5\}$ incident with a or b , then it's easy to see that all 6 points lie on a theta curve or on a dumbbell. Similarly, if the two arcs hit the same segment of $\Theta \setminus \{a, b, x_1, \dots, x_5\}$ then all 6 points lie on a theta curve. Hence, it remains to investigate the case where α_1 hits on the segment $(x_1, x_2) \subseteq e$ and α_2 hits on the segment $(x_2, x_3) \subseteq e$. In this situation, we have a baguette curve consisting of two cycles C_1 and C_2 and disjoint $C_1 - C_2$ arcs β_1 and β_2 with $x_1, x_2 \in C_1$, $x_3, x_4 \in C_2$, $x_5 \in \beta_1$ and $x_6 \in \beta_2$ (i.e. one point x_i on every edge of the baguette curve).

By 3-connectedness, and the fact that $|V(C_i)| \geq 3$, there exists a $C_1 - C_2$ path β_3 which is *alternating* with respect to β_1 and β_2 . Indeed, by Lemma 11.2.4 and the fact that in a 3-regular graph, every alternating walk is automatically a path, we may choose an alternating path β_3 such that the symmetric difference $\beta_1 \triangle \beta_2 \triangle \beta_3$ yields 3 disjoint $C_1 - C_2$ paths γ_1, γ_2 and γ_3 traversing the shared edges with β_3 in the same order as β_3 . Three subcases arise.


- (1) If there exists a $C_1 - C_2$ path containing x_5 and x_6 , then our 6 points lie on a dumbbell. So may assume that β_3 does not contain both x_5 and x_6 .
- (2) If β_3 contains none of x_5 and x_6 , then $\bigcup \gamma_i$ covers both x_5 and x_6 , so either, a single γ_i contains both x_5 and x_6 and we are back in (1), or we have say $x_5 \in \gamma_1$ and $x_6 \in \gamma_2$, and the following subcases arise.

- If γ_1, γ_2 hit the same segment of $C_1 \setminus \{x_1, x_2\}$, then find a cycle picking up 5 of our 6 points, and we are done, and similarly, if γ_1, γ_2 hit the same segment of $C_2 \setminus \{x_3, x_4\}$.

- Otherwise, note that the unique segment δ_i on $C_i \setminus \gamma_3$ between the endpoints of α_1 and α_2 contains precisely one point x_j . Thus, $\alpha_1 \cup \alpha_2 \cup \delta_1 \cup \delta_2$ is a circle containing 4 of our points, and it is then easy to see that using α_3 and suitable segments of $C_i \setminus \delta_i$, we can find a ϑ -curve containing all 6 of our points.

(3) In the final subcase, we may assume that β_3 covers x_5 but not x_6 . Then $x_6 \in \gamma_1$ say, and note that by construction of the symmetric difference, there is an arc $\delta \ni x_5$ which is internally disjoint from $C_1 \cup C_2 \cup \bigcup \gamma_i$ and has its endpoints at interior vertices of some γ_i and γ_j . Note that it follows from Lemma 11.2.4(2) that $i \neq j$. If $i = 1$ and $j \neq 1$ then we are back in case (1). And if say $i = 2$ and $j = 3$, then γ_1 and say γ_2 hit the same segment of $C_1 \setminus \{x_1, x_2\}$, and so by starting with γ_1 , picking up x_1, x_2 on C_1 , following γ_2 , switching to δ , then following γ_3 to C_2 , and move on C_2 back to γ_1 picking up at least one more vertex of x_3 and x_4 , we have found a cycle containing 5 of our points, so we are again done. This completes the argument for case (A).

In case (B), we may use the same arguments as at the beginning of (A) to see that the only critical case is where α_1 hits on the segment $(x_1, x_2) \subseteq e$ and α_2 hits on the segment $(x_3, x_4) \subseteq f$. Then x_1, \dots, x_6 lie on the 6 edges of a K_4 , where we label points and edges as in the figure.



Now consider $|G| \setminus \{x_1, \dots, x_6\}$. By the third-listed assumption on G , this space has at most 3 components, and hence there must exist an arc δ internally disjoint from K_4 between two vertices v, w of G with say $v \in (a, x_6)$ and $w \notin [x_5, a] \cup [x_4, a] \cup [x_6, a]$.

By symmetry, there are five cases to consider for the position of w , namely

- (a) $w \in (d, x_6) \subseteq e_6$, (b) $w \in (d, x_3) \subseteq e_3$, (c) $w \in (b, x_3) \subseteq e_3$,
- (d) $w \in (b, x_4) \subseteq e_4$, and (e) $w \in (b, x_1) \subseteq e_1$.

By inspection, one checks that in cases (b)–(e) there is an arc contained in $K_4 \cup \delta$ which contains all our 6 points. Thus, it remains to deal with case (a).

Let β_1, β_2 be the two disjoint $\{v\} - \{a, d\}$ paths in K_4 . Since G is 3-connected, it follows from Corollary 11.2.5 that there is a further $v - \{a, d\}$ path β_3 which is alternating with respect to $\{\beta_1, \beta_2\}$ (where again we use that any alternating walk must be a path by 3-regularity). Note that by 3-regularity, if h denotes the third edge incident with v , then β_3 and δ agree on h . Starting from v , let z be the first vertex on β_3 which lies on $K_4 \setminus e_6$, and let y be the vertex before z on β_3 . Note that we may assume that

- (1) either $y \in (x_6, d) \subseteq e_6$ and $z \in (x_2, d) \cup (d, x_3)$, or

$$(2) \ y \in (a, x_6) \subseteq e_6 \text{ and } z \in (x_5, a) \cup (a, x_4),$$

as otherwise the arc β_3 between y and z witnesses (up to symmetry) that we are in one of the cases (b) – (f). Now by the fact that β_3 has been chosen according to Corollary 11.2.5, taking the symmetric difference of $\{\beta_1, \beta_2, \beta_3\}$ yields three internally disjoint $\{v\} - \{a, d, z\}$ paths $\gamma_1, \gamma_2, \gamma_3$.

Claim: In case (1), there are two independent $\{x_6\} - \{d, z\}$ paths internally disjoint from $K_4 - e_6$. This follows from Menger’s Theorem 11.2.2 once we show that inside the subgraph $\beta_1 \cup \beta_2 \cup \beta_3$ no single point separates x_6 from the set $\{z, d\}$. So suppose for a contradiction that there is such a separating point s . Since β_2 is a $x_6 - d$ path, we must have $s \in (x_6, d)$. But also walking from x_6 along the edge e_3 to v and then along β_3 to z is a $x_6 - z$ path, it follows that $s \in (x_6, d) \cap \beta_3$, and so $s \notin \gamma_1 \cup \gamma_2 \cup \gamma_3$. But then going from x_6 to v on e_3 , and then taking a suitable γ_i to $\{d, z\}$ shows that s cannot have been a separator.

Claim: In case (2), there are two independent $\{x_6\} - \{a, z\}$ paths internally disjoint from $K_4 - e_6$. Once again this will follow from Menger’s Theorem 11.2.2 once we show that inside the subgraph $\beta_1 \cup \beta_2 \cup \beta_3$ no single point separates x_6 from the set $\{z, a\}$. So suppose for a contradiction that there is such a separating point s . Again, we must have $s \in (a, x_6)$. But also walking from x_6 along the edge e_3 to w and then along $\delta - h$ and $\beta_3 - h$ to z is a $x_6 - z$ path, it follows that $s \in (a, x_6) \cap (\beta_3 - h)$. In particular, we have $s \neq v$ and $s \notin \gamma_1 \cup \gamma_2 \cup \gamma_3$. So either $s \in (a, v)$ or $s \in (v, x_6)$. In the first case, we can walk from x_6 to v on e_6 , and then take a suitable γ_i to reach $\{a, z\}$. In the second case, we can walk from x_6 to w on e_6 , then to v on δ , and then take a suitable γ_i to $\{a, z\}$.

Thus, it follows that by substituting that segment $[a, z]$ or $[d, z]$ of $K_4 \setminus e_6$ with those two disjoint paths, we see that all of our 6 points lie on a ϑ -curve. \square

Some examples of small 6-ac graphs. By checking against a list of simple, 3-regular graphs of small order², we see that the only graph on 6 vertices satisfying our characterization is $K_{3,3}$, and that the only two graphs on 8 vertices satisfying our assumption are $K_{3,3}$ with an extra edge connecting the midpoints of two, non-adjacent edges of $K_{3,3}$ (the so-called Wagner graph), and the 3-dimensional hypercube.

11.3.5. Characterizing 7-ac and ω -ac Graphs.

THEOREM 11.3.15. *Let G be a non-degenerate graph. Then the following are equivalent:*

- (a) G is 7-ac, (b) G is ω -ac, and (c)
- G is homeomorphic to one of the 6 finite graphs which are 7-ac, or
 - G is homeomorphic to one of the finite 7-ac graphs minus possibly some endpoints.

PROOF. Since the graphs mentioned in part (c) are all ω -ac, it suffices to show (a) implies (c). So suppose G is 7-ac. The proof of Theorem 2.12 of [69] shows that G can

²See e.g. https://en.wikipedia.org/wiki/Table_of_simple_cubic_graphs

have at most two vertices of degree 3 or higher. If all vertices have degree two, then G is either homeomorphic to $(0, 1)$ or S^1 . If all vertices have degree no more than 2, but not all are degree 2, then G is either homeomorphic to $[0, 1]$ or $[0, 1)$. Otherwise, extending from the (at most two) vertices of degree at least 3, there will be a finite family of: (finite) cycles, closed intervals (i.e. finite paths) or half-open intervals (i.e. one-way infinite paths). The half-open intervals give rise to the objects mentioned in the second bullet point. \square

CHAPTER 12

n -Arc and n -circle connected graph-like spaces

A space X is *n -arc connected* (respectively, *n -circle connected*) if for any choice of at most n points there is an arc (respectively, a circle) in X containing the specified points. We study n -arc connectedness and n -circle connectedness in compactifications of locally finite graphs and the slightly more general class of graph-like continua, uncovering a striking difference in their behaviour regarding n -arc and n -circle connectedness.

12.1. Introduction

A topological space X is *n -arc connected*, abbreviated *n -ac*, if for any choice of at most n points there is an arc (a homeomorph of the closed unit interval) in X containing the specified points. Similarly, X is *n -circle connected* (abbreviated, *n -cc*) if for any choice of at most n points there is a simple closed curve (homeomorph of the unit circle) in X containing the specified points. Note that a space is arc connected if and only if it is 2-ac. A space which is n -ac (respectively, n -cc) for all n is called ω -ac (respectively, ω -cc).

Every graph is a topological space when considered as a 1-complex, and recently the authors together with A. Mamatelashvili, developing results from [69], have given a complete combinatorial characterization of which graphs (without any restriction on the number of vertices, or edges, or the degree of any vertex) are n -ac or n -cc for any $n \in \mathbb{N}$, see [76]. In particular, a non-degenerate graph G is 7-ac if and only if it is ω -ac if and only if G is homeomorphic to one of nine distinct graphs [76, Theorem 3.5.1]. For $n \leq 6$ there are infinitely many n -ac graphs (even finite), but effective characterizations are now known. For example [76, Theorem 3.4.1]: a graph G is 6-ac if and only if either G is one of the nine 7-ac graphs mentioned above, or, after suppressing all degree-2-vertices, the combinatorial graph G is 3-regular, 3-connected, and removing any 6 edges does not disconnect G into 4 or more components. When considering n -cc graphs, the situation is even simpler: the only 3-cc graphs are the finite cycles, while 2-cc graphs are those that contain no cut vertices.

Finite graphs are extremely simple continua (a *continuum* is a compact, metric and connected space), and for arbitrary continua the problem of characterizing which are n -ac or n -cc is difficult. Indeed, using ideas from descriptive set theory, it is shown in [68] that there is no characterization of n -ac rational continua simpler than the definition of n -ac

(here n is in $\mathbb{N} \cup \{\omega\}$, and a continuum is *rational* if it has a base of open sets whose boundaries are countable).

It is natural to investigate where the transition between the results for graphs – ‘7-ac implies ω -ac’ and effective characterizations for $n \leq 6$ – and the provable complexity for rational continua occurs. In [72], for each n , a regular continuum is constructed which is n -ac but not $(n + 1)$ -ac (a continuum is *regular* if it has a base of open sets whose boundaries are finite). So, in this context, regular continua are too complex.

In the present paper it is shown that the transition takes place precisely between the Freudenthal compactification of locally finite graphs and graph-like continua. Graph-like continua were introduced by Thomassen & Vella [153] as a natural abstraction of the Freudenthal compactification of locally finite graphs. They defined a *graph-like continuum* to be a continuum X which contains a closed zero-dimensional subset V , such that for some discrete index set E we have that $X \setminus V$ is homeomorphic to $E \times (0, 1)$. (Note that, unfortunately, this definition of ‘graph-like’ is *not* the one usually used in continua theory, but this terminology is now standardized among graph theorists.) Up until now all results about the Freudenthal compactification of locally finite graphs have extended naturally to graph-like continua. Thus, it was entirely unanticipated that the n -ac property behaves so differently between the Freudenthal compactification of locally finite graphs and graph-like continua.

12.1.1. Freudenthal compactification of locally finite graphs. Let G be a locally finite, countable, connected graph. Its *Freudenthal compactification*, denoted FG , is the maximal compactification of G with *zero-dimensional* remainder, $FG \setminus G$. (See the discussion immediately preceding Theorem 12.2.3 below for an alternative, constructive description of the Freudenthal compactification of a locally finite graph.) A space is *zero-dimensional* if it has a basis of open sets whose boundaries are empty, i.e. a basis of set which are simultaneously closed and open (clopen).

In the last two decades, Diestel and his students have shown that many combinatorial theorems about paths and cycles in finite graphs extend verbatim to the Freudenthal compactification of infinite, locally finite graphs if one exchanges finite paths and cycles for topological arcs and simple closed curves respectively, see [54, Chapter 8] and [53, 52, 59].

Given this evidence, it might not come as a surprise that the property of n -arc connectedness also lifts nicely to the Freudenthal compactification. Indeed, as our first main result of this paper, we show in Theorem 12.2.3 that for a locally finite, connected graph G and some $n \in \mathbb{N}$, its Freudenthal compactification FG is n -ac [n -cc] if and only if G itself is n -ac [n -cc], allowing us to lift all our characterizations from [76]. However, we also give examples that this is not generally true for all compactifications with zero-dimensional remainder, and it remains an open problem, for example, to characterize for which locally finite graphs the one-point compactification is n -ac. What remains true, though, is the

fact that there are only six different 7-ac graph compactifications, all of which all are again even ω -ac. So there is no jump in complexity happening at this point yet. These results are in Section 12.2.

12.1.2. Graph-like continua. As mentioned above, a *graph-like continuum* is a continuum X which contains a closed zero-dimensional subset V , such that for some discrete index set E we have that $X \setminus V$ is homeomorphic to $E \times (0, 1)$. The sets V and E are the *vertices* and *edges* of X respectively. Clearly a compactification of a connected, locally finite graph is graph-like if and only if the remainder is zero-dimensional. Thus the Freudenthal compactification is graph-like. The points in the remainder of a Freudenthal compactification are called *ends*.

In fact, graph-like spaces were introduced by Thomassen and Vella as a natural abstraction of the Freudenthal compactification of a graph, in order to eliminate the necessity for distinct treatments of vertices and ends in arguments about FG . Papers in which graph-like spaces have played a key role include [153] where several Menger-like results are given, and [49] where algebraic criteria for the planarity of graph-like spaces are presented. In [30], aspects of the matroid theory for graphs have been generalized to infinite matroids on graph-like spaces.

We now know from [70, Theorem A] that graph-like continua had earlier been studied by topologists under the name *completely regular continua* (continua in which every non-degenerate subcontinuum has non-empty interior), and are much closer both to finite graphs and the Freudenthal compactification of graphs than their definition ‘by analogy’ might suggest. Indeed a continuum is graph-like if and only if it is the inverse image of finite graphs under edge-contraction bonding maps (see Section 12.3.1 for details), if and only if it is a (standard) subcontinuum of a Freudenthal compactification of a graph.

Even though the graph-like continua are in complexity just a small step above compactifications of locally finite graphs, it turns out that this is already enough to give rise to completely new and surprising examples of n -ac and n -cc graph-like continua for *all* $n \geq 2$ and ω . For n -circle connectedness, our main result is as follows: while there is topologically a unique 3-cc graph compactification, namely the circle (which is even ω -cc), we show in Theorem 12.4.3 that there are in fact continuum, 2^{\aleph_0} , many pairwise non-homeomorphic ω -cc graph-like continua. For n -arc connectedness, our main result is: while there are only six different 7-ac graph compactifications (which all are even ω -ac), we show in Theorem 12.4.10 that for every $n \geq 2$ there are continuum many n -ac $[n$ -cc] graph-like continua which are not $(n + 1)$ -ac $[(n + 1)$ -cc].

These examples are presented in Section 12.4. In Section 12.3 we develop the necessary machinery to construct graph-like continua, and to check whether they are n -ac or n -cc. In addition – and as an exception to the rule – the 2-cc graph-like continua are

characterized, just like graphs, as being those without cut points, and as having inverse limit representations by finite 2-cc graphs.

12.2. Locally finite graphs, and their Freudenthal compactification

The fundamental result of this section is Theorem 12.2.3 stating that the Freudenthal compactification FG of a locally finite graph G is n -ac precisely when G is n -ac. Since the problem of determining when a graph is n -ac, or n -cc, is completely solved, so is the problem for Freudenthal compactifications of locally finite graphs.

Parts of these results can be extended to arbitrary graph-like compactifications of locally finite graphs. But examples demonstrate that Theorem 12.2.3 does not extend in full generality to graph-like compactifications of locally finite graphs.

12.2.1. Some notation. Trails and walks in graphs, and paths and arcs in graph-like continua, all start and end at points (typically vertices). It is convenient to turn this around, and given vertices v and w , by a $v - w$ -trail, path or arc we mean a trail, path or arc (respectively) starting at v and ending at w . More generally, given sets R and S , of vertices by an $R - S$ -trail (path or arc) we mean a trail (path or arc, respectively) starting at some element of R and ending at some member of S .

12.2.2. Restricting to points on edges. We begin with the following extension of [76, Lemma 2.3.5] to the class of regular continua. Since graph-like continua are regular [70, Lemma 7], its critical corollary is that in order to check whether a graph-like continuum is n -ac, it is sufficient to assume the points lie on edges. It is convenient also to extend our definitions. Let X be a space and S a subset. Then (S, X) is n -ac (respectively, n -cc) if for any choice of at most n points from S there is an arc (resp., simple closed curve) in X containing the specified points.

LEMMA 12.2.1. *Let X be a regular continuum, $D \subseteq X$ an arbitrary dense subset of X , and $n \in \mathbb{N}$. Then X is n -ac [n -cc] if and only if (D, X) is n -ac [n -cc].*

PROOF. Only the backwards implication requires proof. Assume that (D, X) is n -ac and let $x_0, x_1, \dots, x_n \in X$ be arbitrary (with $n \geq 1$). Since X is regular, there are open neighbourhoods $U_i \ni x_i$ such that

- $\overline{U_i} \cap \overline{U_j} = \emptyset$ for all $0 \leq i < j \leq n$, and such that
- $|\partial U_i| = k_i \in \mathbb{N}$ is minimal with respect to all open neighbourhoods V of x_i with $V \subseteq U_i$ for all i .

Pick points $y_i \in U_i \cap D$. By assumption, there is an arc [closed curve] α going through y_0, y_1, \dots, y_n , having two of these points as its endpoints. We are now going to argue that we can modify α inside each $\overline{U_i}$ as so to pick up x_i but still remain an arc [closed curve] in X . It suffices to give this argument for $i = 0$, so write $x = x_0$, $U = U_0$ and $k = k_0$.

Let us assume that $\partial U = \{u_1, \dots, u_k\}$. Without loss of generality, α passes through u_1, \dots, u_i in the given linear [cyclic] order (for $1 \leq i \leq k$), and doesn't use u_{i+1}, \dots, u_k . If $k = 1$, it is clear how to use local arc-connectedness of X to add x_1 to our arc α [in the n -cc case, $k = 1$ cannot occur]. Otherwise, since at least one of the endpoints of α lies outside of U [and trivially in the n -cc case], we see that $\bar{U} \cap \alpha$ consists of at most $i - 1 \leq k - 1$ connected arcs (and at least one, as $y_0 \in \bar{U} \cap \alpha$).

Next, by the fact that $|\partial U| = k \in \mathbb{N}$ was minimal with respect to all neighbourhoods of x contained in U , it follows from Menger's n -od Theorem that there is a k -fan F with center x and leaves in α contained in \bar{U} , see [118, Theorem §I] or [131, Der verschärfte n -Beinsatz]. By the pigeon hole principle, two leaves of the fan F must lie on the same connected component of $\bar{U} \cap \alpha$, and so it is clear how to include x into our arc [closed curve] α . As this procedure can be repeated for all $i = 1, \dots, n$, the proof is complete. \square

12.2.3. Freudenthal compactification of locally finite connected graphs. In the proof of the next theorem, we need the following standard lemma bounding the number of edges in a graph leaving a certain vertex set. For a subset $A \subseteq V(G)$ write $\partial A = E(A, V \setminus A)$ for the induced edge cut (we write $\partial_G A$ for ∂A when we want to emphasize we are working inside the graph G), and A^c for $V(G) \setminus A$.

LEMMA 12.2.2. *Let G be a graph, $A, A' \subseteq V(G)$. Then*

$$|\partial A| + |\partial A'| \geq \max \{ |\partial(A \cap A')| + |\partial(A \cup A')|, |\partial(A \setminus A')| + |\partial(A' \setminus A)| \}.$$

PROOF. We indicate the short argument of this folklore lemma: We have to verify that every edge e that is counted on the right will also be counted on the left, and if it is counted say in both $\partial(A \cap A')$ and $\partial(A \cup A')$ on the right, it is also counted in both sums on the left.

If $e \in \partial(A \cap A')$, then e joins a vertex $v \in A \cap A'$ to a vertex w that fails to lie in A or which fails to lie in A' . In the first case, $e \in \partial A$, and in the second case we have $e \in \partial A'$. Since $\partial(A \cup A') = \partial(A^c \cap A'^c)$, the same holds for edges in $\partial(A \cup A')$: every such edge lies in $\partial(A^c) = \partial A$ or in $\partial(A'^c) = \partial A'$.

Finally, if e is counted twice on the left, i.e., if $e \in \partial(A \cap A')$ and $e \in \partial(A \cup A') = \partial(A^c \cap A'^c)$, then e joins a vertex $v \in A \cap A'$ to some other vertex, and it also joins some $w \in A^c \cap A'^c$ to some other vertex. As $A \cap A'$ and $A^c \cap A'^c$ are disjoint, we have $e = vw$. But this means that $e \in \partial A$ as well as $e \in \partial A'$, so e is counted twice also on the left.

The other inequality, $|\partial A| + |\partial A'| \geq |\partial(A \setminus A')| + |\partial(A' \setminus A)|$, now follows from the first one by applying the fact that $|\partial B| = |\partial(B^c)|$. \square

The final ingredient for our key Theorem 12.2.3 is an alternative, and more explicit, description of the Freudenthal compactification of a locally finite graph in terms of ends.

Let G be a locally finite connected graph. A 1-way infinite path is called a *ray*, a 2-way infinite path is a *double ray*. Two rays R and S in G are *equivalent* if no finite

set of vertices separates them. Alternatively, we may say that G contains infinitely many disjoint $R - S$ -paths. The corresponding equivalence classes of rays are the *ends* of G . The set of ends of a graph G is denoted by $\Omega = \Omega(G)$.

Recall that topologically, we view G as a cell complex with the usual 1-complex topology. Adding its ends compactifies it, with the topology on $G \cup \Omega$ generated by the open sets of G and neighbourhood bases for ends $\omega \in \Omega$ defined as follows: Given any finite subset S of $V(G)$, let $C(S, \omega)$ denote the unique component of $G - S$ that contains a cofinal tail of some (and hence every) ray in ω , and let $\hat{C}(S, \omega)$ denote the union of $C(S, \omega)$ together with all ends of G with a ray in $C(S, \omega)$. As our neighbourhood basis for ω we take all sets of the form $\hat{C}(S, \omega) \cup \overset{\circ}{E}(S, C(S, \omega))$, where S ranges over the finite subsets of $V(G)$ and $\overset{\circ}{E}(S, C(S, \omega))$ denotes the interior of the edges with one endpoint in S and the other in $C(S, \omega)$. Note that in this topology, we have $\overline{\hat{C}(S, \omega)} \cap \Omega = \hat{C}(S, \omega) \cap \Omega$.

It is well known that this process of adding the ends does indeed yield the Freudenthal compactification, i.e. $FG = G \cup \Omega$. In particular it is locally connected at ends, and has neighbourhoods which restrict to zero-dimensional sets on the end space. For further details and proofs see Chapter 8 of [54].

THEOREM 12.2.3. *For the Freudenthal compactification FG of a locally finite connected graph G the following are equivalent for each $n \in \mathbb{N}$:*

- (1) FG is n -ac, (2) (G, FG) is n -ac, and (3) G is n -ac.

PROOF. The equivalence (1) \Leftrightarrow (2) is a special instance of Lemma 12.2.1. The implication (3) \Rightarrow (2) is trivial. For (2) \Rightarrow (3) consider n points $x_1, \dots, x_n \in G$ and find, by assumption, an arc α in FG going through the specified points. Our task is to modify this arc α so that it still contains x_1, \dots, x_n but does not use ends of G anymore.

Without loss of generality we may assume that start- and end-point of α are amongst the x_i . Then it follows from [39, Prop. 3] that every end $\omega \in \alpha \cap (FG \setminus G)$ has degree 2 in α , meaning that for every finite set of vertices $S \subseteq V(G)$ there is a bipartition (A_ω, B_ω) of $V(G)$ such that: (i) the induced subgraph $G[A_\omega]$ is connected, (ii) $\omega \in \overline{A}$, (iii) $S \subseteq B_\omega$, and (iv) $|E(\alpha) \cap \partial A_\omega| = 2$ (i.e. the arc α uses precisely two edges from the edge cut $E(A_\omega, B_\omega)$).

Let us call such a set A_ω with $|E(\alpha) \cap \partial A_\omega| = 2$ a *2-neighbourhood* of ω . Moreover, note that $|E(\alpha) \cap \partial A| \geq 2$ whenever $\omega \in \overline{A}$ and $A \subseteq A_\omega$ (\star). Next, let $S = \{x_1, \dots, x_n\}$ and choose for every end $\omega \in \alpha \cap (FG \setminus G)$ a bipartition (A_ω, B_ω) with the above four properties. Since $\alpha \cap (FG \setminus G)$ is compact, there are finitely many ends $\omega_1, \dots, \omega_\ell$ such that $\alpha \cap (FG \setminus G) \subseteq \overline{A_{\omega_1}} \cup \dots \cup \overline{A_{\omega_\ell}}$. We may assume that this cover is minimal, i.e. for every $i \leq \ell$ there is an end $\epsilon_i \in \alpha \cap (FG \setminus G)$ such that $\epsilon_i \in \overline{A_i} \setminus \bigcup \{\overline{A_j} : j \neq i\}$ ($\star\star$).

Claim: *Every minimal cover of $\alpha \cap (FG \setminus G)$ consisting of 2-neighbourhoods has a disjoint refinement consisting of 2-neighbourhoods.*

The proof of the claim is via induction on the size of the cover. Let us make the convention that $\partial_\alpha A := E(\alpha) \cap \partial A$ consists of those boundary edges of A that are used by α . If the cover consists of a single element only, there is nothing to show. So we may assume $\ell \geq 2$ and consider our cover $\{A_1, \dots, A_\ell\}$. Let $\tilde{A}_1 := A_1$ and $\tilde{A}_i := A_i \setminus A_1$ for all $1 < i \leq \ell$. From (\star) and $(\star\star)$ it follows that $|\partial_\alpha \tilde{A}_i| \geq 2$ for all $i \leq \ell$.

We shall use Lemma 12.2.2 to see that $|\partial_\alpha \tilde{A}_i| \leq 2$ for all $i \leq \ell$ as well. This is clear for \tilde{A}_1 . For $i \geq 2$, Lemma 12.2.2 applied to the graph $(V, E(\alpha))$ implies

$$4 = |\partial_\alpha A_1| + |\partial_\alpha A_i| \geq |\partial_\alpha(A_1 \setminus A_i)| + |\partial_\alpha(A_i \setminus A_1)| \geq 2 + |\partial_\alpha \tilde{A}_i|,$$

where $\partial_\alpha(A_1 \setminus A_i) \geq 2$ follows again from (\star) and $(\star\star)$. Thus, we have $|\partial_\alpha \tilde{A}_i| = 2$ for all $i \leq \ell$. Applying the induction assumption to the collection $\{\tilde{A}_2, \dots, \tilde{A}_\ell\}$ we obtain a disjoint refinement of 2-neighbourhoods, which together with A_1 forms the desired refinement of our original collection. This establishes the claim.

Next, we argue that for each \tilde{A}_i , there is a finite edge path P_i in $G[\tilde{A}_i]$ from one edge in $\partial_\alpha \tilde{A}_i$ to the other. Let $\alpha_i \subseteq \alpha$ be the subarc of α that lies in the closure of \tilde{A}_i in FG . By definition of the topology of the Freudenthal compactification, for every end ω in α_i , there is a finite subset $T \subseteq V(G)$ such that $C(T, \omega) \subseteq \tilde{A}_i$. By compactness, finitely many such $C(T_j, \omega_j)$ for $j \leq N$ say cover the ends used by α_i . Now since every $C(T_j, \omega_j)$ is by definition a connected graph, we may recursively in j find a finite edge-path in $C(T_j, \omega_j)$ connecting the first and last point of $\alpha_i \cap C(T_j, \omega_j)$. By doing so, we obtain a finite edge-walk in $G[\tilde{A}_i]$ from one edge in $\partial_\alpha \tilde{A}_i$ to the other, which includes the desired finite edge path P_i .

But now we are done: for each $i \leq \ell$, replace α_i by P_i . Since each replacement took place in the disjoint subsets \tilde{A}_i , this gives rise to an arc completely inside the graph G containing all n points x_1, \dots, x_n as desired. \square

12.2.4. Graph-like compactification of locally finite connected graphs. Since every 7-ac graph is one, up to homeomorphism, of a finite family, we easily deduce from Theorem 12.2.3 that the Freudenthal compactification of a locally finite graph is 7-ac only in very limited cases. However, this holds for arbitrary graph-like compactifications (i.e. for compactifications with zero-dimensional remainders).

PROPOSITION 12.2.4. *Let G be a countable, locally finite graph. Let γG be a graph-like compactification of G .*

If γG is 7-ac then γG is (homeomorphic to) a finite graph (and is one of the 6 finite graphs which are 7-ac, or equivalently ω -ac).

PROOF. The proof of Theorem 2.12 of [69] shows that the graph G can have at most two vertices of degree 3 or higher. If all vertices have degree two, then as above γG is

either an arc or a circle. If all vertices have degree no more than 2, but not all are degree 2, then G is either a finite chain, or an infinite one-way chain. In either case γG is an arc or a circle. Otherwise, extending from the (at most two) vertices of degree at least 3, there will be a finite family of: (finite) cycles, finite chains or infinite one-way chains. The infinite chains have either one or two endpoints in γG . In all scenarios, γG is homeomorphic to a finite graph. \square

Although Theorem 12.2.3, as stated, only applies to n -arc connectedness, and not n -circle connectedness, the n -cc property is completely dealt with via the next two lemmas. Indeed, as in the previous result, these apply to arbitrary graph-like compactifications of locally finite graphs.

LEMMA 12.2.5. *Let γG be a graph-like compactification of a countable, locally finite graph G . Then the following are equivalent: (a) γG is 2-cc, (b) γG has no cut points, (c) G has no cut points, (d) G is 2-cc, and (e) G is cyclically connected.*

PROOF. Since γG is graph-like, the equivalence of (a) and (b) follows from Proposition 12.3.5 below. Since no point of the remainder, $\gamma G \setminus G$, can be a cut point of γG ; while every cut point of G is a cut point of γG , we see that (b) and (c) are equivalent. Finally, the characterization of 2-cc graphs (Theorem 3.1.1 of [76]) yields the remaining equivalences. \square


LEMMA 12.2.6. *Let γG be a graph-like compactification of a countable, locally finite graph G . Then the following are equivalent: (a) γG is 3-cc, (b) γG is a circle, and (c) G is either a cycle, or a double ray and γG is its one-point compactification.*

PROOF. Suppose γG is 3-cc. The corresponding argument for finite graphs shows that every vertex of G has degree 2. So G is either a finite cycle, or a double ray. In the latter case, there are only two different graph-like compactifications: γG is either a circle, or an arc – but in the latter case, γG is not 3-cc. \square

However, Theorem 12.2.3, stating that a locally finite, countable graph G is n -ac if and only if FG is n -ac, does not extend to general graph-like compactifications for $n \leq 6$.

EXAMPLE 12.2.7.

- (a) The infinite ladder, D , is 5-ac but not 6-ac, while αD is 6-ac.
- (b) The graph C below is 4-ac but not 5-ac, while its one-point compactification, αC , is 6-ac.




Proof. For (a): Let D be the usual double ladder, i.e. $V(D) = \{0, 1\} \times \mathbb{Z}$ in which two vertices (m, n) and (m', n') are adjacent if and only if $|m - m'| + |n - n'| = 1$. Using the characterizations from [76], it follows that D is 5-ac but not 6-ac.

We focus on showing αD is 6-ac. Since we may assume our six points x_1, \dots, x_6 lie on edges, we may find $n \geq 5$ large enough such that $x_1, \dots, x_6 \in D[\{0, 1\} \times [-n, n]]$.

Set $G_1 = D[\{0, 1\} \times [-n, n]]$. Take a disjoint copy of G_1 , and modify it to form a graph G_2 as follows: first, remove the edge corresponding to $\{(0, 0), (0, 1)\}$, and second, subdivide the edges $\{(0, -2), (0, -3)\}$ and $\{(0, 3), (0, 4)\}$ by vertices a and b , and, finally, add new edges from a to $(0, 0)$ and $(0, 1)$ to b . Let us write $e = \{(1, 0), (1, 1)\}$ for the unique bridge of G_2 , and $G_2^+ := G_2[\{0, 1\} \times \{1, \dots, n\}]$ and $G_2^- := G_2[\{0, 1\} \times \{0, -1, \dots, -n\}]$ for the two components of $G_2 - e$.

Now consider the auxiliary graph $G = G_1 \sqcup G_2$ where we additionally add four new edges: (1) f^+ between the copies of $(0, n)$, (2) f^- between the copies of $(0, -n)$, (3) g^+ between the copies of $(1, n)$, and (4) g^- between the copies of $(1, -n)$.

It follows from [76, Theorem 3.4.1] that G is 6-ac, and so there is an arc α in G containing x_1, \dots, x_6 and, without loss of generality, starting and ending in points $x_i \neq x_j$. In particular, α starts and ends outside of G_2 . Moreover, note that $\partial_G G_2^+ = \{e, f^+, g^+\}$ is a 3-edge cut, and so if α contains points from G_2^+ then α will cross this cut in precisely two edges, and so $\beta^+ = \alpha \cap G_2^+$ will be a subarc of α . Similarly, $\beta^- = \alpha \cap G_2^-$ will be a subarc of α . But then it is clear that by replacing β^+ and β^- with suitable arcs in the corresponding connected components of $\alpha D \setminus G_1$ (where say an $e - f^+$ arc will be replaced by an $\infty - f^+$ -arc in αD), we may lift α to an arc in αD witnessing 6-ac.



For (b): That αC is 6-ac can be directly checked by a case-by-case analysis.

To see that C is 4-ac but not 5-ac we can apply the characterizations of [76] as follows. First note that removing the middle edge disconnects C into two components C_+, C_- which are isomorphic. Since C_\pm is cyclically connected, and no two vertices cut it into 4 or more components, it is 4-ac by [76, Theorem 3.2.1]. As C is 3-regular it follows from

[76, Theorem 3.2.3] that C is 4-ac. On the other hand, since removing the middle edge disconnects C , it is not cyclically connected. Now [76, Theorem 3.3.1] states that for C to be 5-ac it must be homeomorphic to one of: an arc, ray, double ray, lollipop with or without end point, dumbbell or figure-eight, and it is clearly not homeomorphic to any of these spaces. \square

The argument given that αD is 6-ac is straightforward, but follows from an *ad hoc* reduction to the combinatorial graph characterization of 6-ac. The direct check that αC is 6-ac is lengthy and tedious, in sharp contrast to the simple arguments, from the combinatorial characterizations, that C is 4-ac but not 5-ac. These two examples demonstrate some of the difficulties in determining when a graph-like compactification of a locally finite, connected graph G is n -ac, and also the value in having a combinatorial characterization.

PROBLEM 12.2.8. Find a combinatorial characterisation for when a graph-like compactification of a locally finite, connected graph G is n -ac.

A place to start would be to discover when the one-point compactification of a graph is 6-ac.

12.3. General graph-like continua

In this section we first develop some machinery for graph-like spaces with the aim of connecting them, via inverse limits with ‘nice’ bonding maps, to finite graphs. This machinery then yields tests for a graph-like continuum to be, or not to be, n -ac or n -cc. In Proposition 12.3.5 these tests are refined to characterize 2-cc graph-like continua. In the next section our machinery and tests for graph-likes are applied to construct various examples.

12.3.1. Graph-like spaces as inverse limits. Here we develop techniques of Espinoza and the present authors in [70], to detect when a continuum is graph-like, and characterize when a graph-like continuum is Eulerian.

For convenience let us say that a map π from one graph-like continuum, X , to another, Y , is *nice* if it is surjective, monotone (fibres, $\pi^{-1}\{v\}$, are connected) and maps vertices to vertices, and edges either homeomorphically to another edge, or to a vertex.

Let X be a graph-like continuum with vertex set V . By subdividing edges once, if necessary, we may assume that every edge of X has two distinct endpoints in V , i.e. that the graph-like continuum is *simple*.

For a clopen subsets $U, U' \subseteq V$, not necessarily different, $E(U, U')$ denotes the set of edges with one endpoint in U and the other endpoint in U' . It is not hard to see, [70, Lemma 1], that $E(U, V \setminus U)$ is always finite. A *multi-cut* is a partition $\mathcal{U} = \{U_1, U_2, \dots, U_n\}$ of V into pairwise disjoint clopen sets such that for each i , the *induced subspace* $X[U_i]$ of X , i.e. the closed graph-like subspace with vertex set U_i and edge set $E(U_i, U_i)$, is connected.

The *multigraph associated with* \mathcal{U} is the quotient $G_X(\mathcal{U}) = G(\mathcal{U}) = X/\{X[U] : U \in \mathcal{U}\}$. Let $p_{\mathcal{U}}: X \rightarrow G(\mathcal{U})$ denote the quotient mapping from X to the multigraph associated with \mathcal{U} . We note that $G(\mathcal{U})$ is indeed a finite, connected multi-graph, and that $p_{\mathcal{U}}$ is nice. Conversely, if p is a nice map of X to a finite, connected graph G , then there is a multi-cut \mathcal{U} such that $G = G(\mathcal{U})$ and $p_{\mathcal{U}}$ realizes p in the sense that they are identical on the vertices of X , and they carry the same edges of X to the same edges of G .

A sequence, $(\mathcal{U}_n)_n$, of multi-cuts of X is *cofinal* if for every multi-cut \mathcal{U} there is an \mathcal{U}_n which refines it. According to Theorem 13 of [70], for any cofinal sequence, $(\mathcal{U}_n)_n$, of multi-cuts, the graph-like continuum X is naturally homeomorphic to an inverse limit $\varprojlim G_X(\mathcal{U}_n)$, where the bonding maps are all nice. Conversely, if a space X is homeomorphic to an inverse limit, $\varprojlim G_n$, where the G_n are finite, connected graphs, and all bonding maps are nice, then (Theorem 14 of [70]) X is a graph-like continuum. Note that in this case, for every m , the projection map, typically denoted, p_m , from $\varprojlim G_n$ to G_m is nice, and so is realized as a $p_{\mathcal{U}_m}$ for some multi-cut \mathcal{U}_m .

12.3.2. Sufficient conditions. The following lemma – a special case of Lemma 12.2.1 – records that as in the case with graphs, also for graph-like continua we may choose our points x_1, \dots, x_n without loss of generality to be interior points of edges.

LEMMA 12.3.1. *Let X be a graph-like continuum with vertex set V . Let $n \in \mathbb{N}$. Then X is n -ac [n -cc] if and only if $(X \setminus V, X)$ is n -ac [n -cc].*

LEMMA 12.3.2. *Let \mathcal{U} be a multi-cut of a graph-like continuum X . Then every arc [simple closed curve] in $G = G_X(\mathcal{U})$ lifts to an arc [simple closed curve] in X .*

PROOF. Since the quotient mapping $p_{\mathcal{U}}: X \rightarrow G(\mathcal{U})$ is nice, it follows that for every vertex v of G , its fibre $p_{\mathcal{U}}^{-1}(v) = X[U]$ for some $U \in \mathcal{U}$ is an connected, and hence arc-connected subcontinuum of X , see [70, Lemma 2]. Thus, we may lift any arc [simple closed curve] α in $G = G_X(\mathcal{U})$ by filling in suitable subarcs inside each fibre $p_{\mathcal{U}}^{-1}(v) = X[U]$ for every vertex $v \in \alpha$. \square

COROLLARY 12.3.3. *Let X be a graph-like continuum. If $G_X(\mathcal{U})$ is n -ac [n -cc] for every multicut \mathcal{U} of X , then X is n -ac [n -cc].*

PROOF. By Lemma 12.3.1 it suffices to consider points x_1, \dots, x_n lying on edges of X , say $x_i \in e_i$. Since $\varprojlim G_X(\mathcal{U}_n) \cong X$, there is a multicut \mathcal{U} of X such that e_1, \dots, e_n are all displayed in the finite graph $G = G_X(\mathcal{U})$. By assumption, $G_X(\mathcal{U})$ is n -ac [n -cc], and so there is an arc [simple closed curve] in G containing the distinct points $p_{\mathcal{U}}(x_1), \dots, p_{\mathcal{U}}(x_n)$. The assertion is then immediate by Lemma 12.3.2. \square

12.3.3. Necessary conditions. Call a graph G n -E (n -Eulerian) if for every n or fewer points in G there is an edge disjoint closed trail in G containing the points. Equivalently, we may say that every n edges of G lie on a common Eulerian subgraph of G .

Observe that a finite graph is Eulerian if and only if it is n -E for all n . Call a graph G n -oE (n -open Eulerian) if for every n or fewer points in G there is an edge disjoint (possibly not closed) trail in G containing the points.

PROPOSITION 12.3.4. *Let X be a graph-like continuum.*

- (a) *If X is n -cc, then for every multi-cut \mathcal{U} of X the graph $G(\mathcal{U})$ is n -E.*
 (b) *If X is n -ac, then for every multi-cut \mathcal{U} of X the graph $G(\mathcal{U})$ is n -oE.*

PROOF. We prove (a). So suppose X is n -cc. Write X as an inverse limit $X = \varprojlim G_k$ of graphs, with nice bonding maps. We verify that each G_k is n -E.

Fix k . Let p_k be the nice projection from X to G_k . Take no more than n points from G_k , say x_1, \dots, x_n . Pick y_1, \dots, y_n in X , such that $p_k(y_i) = x_i$, for $i = 1, \dots, n$. As X is n -cc, there is a simple closed curve S in X containing these points. The projection of S under p_k into G_k is an edge-disjoint closed trail in G_k which contains all the x_i . This shows that G_k is n -E.

The proof of (b) is very similar. In place of a circle we get an arc α containing y_1, \dots, y_n . Its projection in G_k is an edge-disjoint trail which may or may not be closed, but definitely contains the points x_1, \dots, x_n . Thus G_k is n -oE. \square

12.3.4. 2-cc Graph-like Continua. A space X is 3-sac if given any three points, x_1, x_2, x_3 of X , there is an arc in X starting at x_1 , passing through x_2 , and ending at x_3 . The main result here is the following one showing that in graph-like continua being 2-cc is equivalent to being 3-sac, and characterizing these properties in terms of the standard properties of the graph-like continuum and, also, its inverse limit representation.

PROPOSITION 12.3.5. *For a graph-like continuum X , the following are equivalent:*

- (1) X is 3-sac, (2) X is 2-cc, (3) X has no cut points,
 (4) for every representation $X = \varprojlim G_m$, where each G_m is a finite, connected graph and each bonding map is nice, there is a m such that G_m has no cut-point,
 (5) X can be represented as $X = \varprojlim G_n$, where each G_n is a finite, 2-cc graph and each bonding map is nice.

The next lemma shows the equivalence of (1), (2) and (3) even among all Peano continua. Lemma 12.3.6 also shows that (5) is equivalent to (5') where '2-cc' is replaced by 'no cut points'. Then the equivalence of (3), (4) and (5') is the $k = 2$ case of Proposition 12.3.7.

LEMMA 12.3.6. *For a Peano continuum X , the following are equivalent:*

- (1) X is 3-sac, (2) X is 2-cc, and (3) X has no cut points.

PROOF. The equivalence of (1) and (2) for any continuum X was established in [68, Prop. 7], and was shown in [68, Thm. 5], evoking a result by Bellamy and Lum, to be equivalent to (3') X is arc connected, has no arc-cut point, and has no arc end points (x

is an arc end point if there are not two arcs intersecting only at x). Clearly (3') implies (3) (cut points are arc-cut points).

Suppose X is Peano. Then it is arc connected. We show if X contains an arc-cut point or an arc end point then it contains a cut point, and so (3) implies (3').

First, assume that X has an arc end point x . Recall the result by Nöbling [131] that if a point x in a Peano continuum X has order at least n (i.e. any small enough neighbourhood of x has boundary at least of size n) then X contains an n -od with center x , i.e. a union of n many arcs with only the point x in common. Thus, an arc end point must necessarily have order 1, and so we have found many cut-points.

Second, it is not hard to show that every arc-cut point x of a Peano continuum X must necessarily be a cut point. Indeed, suppose $X \setminus \{x\}$ has at least 2 arc-components. Let Y be an arc-component. Using local connectedness, it is easy to show that Y must be closed in $X \setminus \{x\}$, and further, that the collection $\{Y \subseteq X \setminus \{x\} : Y \text{ arc-component}\}$ is a locally finite collection of sets. Thus, one arc component against the union of the rest is a partition of $X \setminus \{x\}$ into non-empty closed sets. \square

In analogy to graphs, call a graph-like continuum k -connected if the deletion of any $k - 1$ vertices never disconnects it. Note that a graph-like continuum is 2-connected if and only if it has no cut points (if removing a point on an edge disconnects, then so does removing either of the end points of the edge).

A k -pre-cutting is a triple (Y, A, B) where Y is a set of vertices with $|Y| < k$, and A, B are subcontinua with $A \cup B = X$ and $A \cap B = Y$. A k -cutting of X is a non-trivial k -pre-cutting, (Y, A, B) where by *non-trivial* we mean that $A \setminus Y$ and $B \setminus Y$ are non-empty. Observe that if (Y, A, B) is a k -cutting then $X \setminus Y$ is disconnected. Conversely, if Y is a set of vertices of size $< k$, and removing Y from X disconnects X , say $X \setminus Y = U \cup V$ where U are disjoint, open and non-empty, then (Y, A, B) is a k -cutting, where $A = U \cup Y$ and $B = V \cup Y$.

If $f: Z \rightarrow W$ is a nice map from Z to another graph-like continuum, W , and (Y, A, B) is a k -pre-cutting in Z , then $(f(Y), f(A), f(B))$ is a k -pre-cutting in W .

PROPOSITION 12.3.7. *For a graph-like continuum X , the following are equivalent:*

- (a) X is k -connected,
- (b) for every representation $X = \varprojlim G_m$, where each G_m is a finite, connected graph and each bonding map is nice, there is an m such that G_m is k -connected, and
- (c) X can be represented, $X = \varprojlim G_n$, where each G_n is a finite, k -connected graph and each bonding map is nice.

PROOF. Suppose (b) holds. Fix a representation $X = \varprojlim G_m$. For any n , we have $X = \varprojlim_{m \geq n} G_m$, so, by (b), for some $m_n \geq n$ we know G_{m_n} is k -connected. Letting $H_n = G_{m_n}$, we have a representation $X = \varprojlim H_n$ where all the graphs involved are k -connected. Thus (c) follows from (b).

Next suppose (a) fails, we show (c) also fails, and so (c) implies (a). Fix a k -cutting (Y, A, B) of X . Take any representation $X = \varprojlim G_m$, where each G_m is a connected, finite graph, and each bonding map is nice. Denote, as usual, the projection map of $\varprojlim G_m$ to G_m by p_m , and recall it is nice. Pick x in $A \setminus Y$, and $b \in B \setminus Y$. Find m sufficiently large that in G_m the points $p_m(a)$ and $p_m(b)$ are distinct and not contained in $p_m(Y)$. Then, as p_m is nice, $(p_m(Y), p_m(A), p_m(B))$ is a k -cutting of G_m , which, therefore, is not k -connected.

Finally we show if (b) is false then so is (a). Fix a representation $X = \varprojlim G_m$, where each G_m is a finite, connected graph which is not k -connected, and each bonding map, $\pi_m: G_{m+1} \rightarrow G_m$, is nice. Let \mathcal{T} be the set of all finite sequences $\langle (Y_1, A_1, B_1), \dots, (Y_n, A_n, B_n) \rangle$ where each (Y_m, A_m, B_m) is a k -pre-cutting of G_m , $Y_m = \pi_m(Y_{m+1})$, $A_m = \pi_m(A_{m+1})$, $B_m = \pi_m(B_{m+1})$, and some term in the sequence is non-trivial (i.e. a k -cutting, and note all subsequent terms of the sequence are also non-trivial).

Order \mathcal{T} by extension to get a tree. Observe that every sequence in \mathcal{T} has only finitely many immediate successors (indeed there are only finitely many k -pre-cuttings, (Y_m, A_m, B_m) , of G_m , since Y_m is a set of vertices of the finite graph G_m). Further \mathcal{T} is infinite. To see this fix n . We show there is a sequence in \mathcal{T} of length n . Well, by hypothesis, G_n is not k -connected, and so contains a k -cutting (Y_n, A_n, B_n) . Then $\langle (Y_1, A_1, B_1), \dots, (Y_m, A_m, B_m), \dots, (Y_n, A_n, B_n) \rangle$ is in \mathcal{T} where

$$(Y_m, A_m, B_m) = (\pi_m(Y_{m+1}), \pi_m(A_{m+1}), \pi_m(B_{m+1})) \quad \text{for } m = n - 1, \dots, 1.$$

By König’s Lemma, see e.g. [54, Lemma 8.1.2], the tree \mathcal{T} has an infinite branch, $\sigma_1, \sigma_2, \dots, \sigma_m, \dots$. So we get an infinite sequence of k -pre-cuttings $\langle (Y_m, A_m, B_m) \rangle_m$ which are mutually compatible: $(Y_m, A_m, B_m) = (\pi_m(Y_{m+1}), \pi_m(A_{m+1}), \pi_m(B_{m+1}))$, for all $m \geq 1$. Let $A = \varprojlim A_m$, $B = \varprojlim B_m$ and $Y = A \cap B$. Then, by compatibility, A and B are subcontinua of X , $X = A \cup B$ and Y is a set of vertices. But some term of the branch is non-trivial, and so from that point on, all the k -pre-cuttings are non-trivial. Further the sets Y_m must stabilize. Thus (Y, A, B) is a non-trivial k -pre-cutting, and X is not k -connected. □

12.3.5. Distinguishing graph-like continua. Let X be a graph-like continuum. For distinct vertices v and w from X define $k_X(v, w)$, the *edge connectivity* between v and w , to be the minimal number of edges whose removal separates v and w (i.e. which form an edge-cut between v, w). Note that $k_X(v, w)$ is well-defined, and by Menger’s theorem for graph-like continua, [70, Theorem 22], $k = k_X(v, w)$ equals the maximum size of a family of edge-disjoint $v - w$ -paths.

LEMMA 12.3.8. *Let X be a graph-like continuum containing distinct vertices v and w . If Y is another graph-like continuum and f is a nice map of X to Y then $k_X(v, w) \leq k_Y(f(v), f(w))$ provided $f(v) \neq f(w)$.*

PROOF. Pick edges e_1, \dots, e_k that separate $f(v)$ from $f(w)$ in Y . Then, as f is nice, those same edges exist in X and separate $v \in f^{-1}\{f(v)\}$ from $w \in f^{-1}\{f(w)\}$. \square

LEMMA 12.3.9. *Let $X, X' \neq S^1$ be graph-like continua with standard representations $X = (V, E)$ and $X' = (V', E')$. Then every homeomorphism $f: X \rightarrow X'$ is a nice isomorphism of graph-like spaces.*

PROOF. Since the degree of a point is a topological property, and hence preserved under homeomorphisms, it follows that any homeomorphism $f: X \rightarrow X'$ must map V homeomorphically to V' and therefore, by considering complements, edges to edges. Since it is bijective, it is trivially monotone. \square

In particular, the previous two lemmas allow us to use *combinatorial information* to show that two graph-like continua X and Z are non-homeomorphic. Indeed, it suffices to find distinct v and w in X such that $k_X(v, w) \neq k_Z(v', w')$ for all distinct v', w' in Z . This is specified explicitly by the next lemma.

LEMMA 12.3.10. *Let $X = (V, E)$ be a graph-like continuum, with representation $X = \varprojlim G_k$, of connected graphs, with nice bonding maps. Let v and w be distinct vertices of X and define $s = s(v, w)$ to be minimal such that $p_s(v) \neq p_s(w)$.*

Then $k_X(v, w) = \min\{k_{G_t}(p_t(v), p_t(w)) : t \geq s\} =: \underline{k}$.

Further, the sequence $(k_{G_t}(p_t(v), p_t(w)) : t \geq s)$ is decreasing and eventually constant. It stabilizes, so $k_X(v, w) = k_{G_t}(p_t(v), p_t(w))$, at the minimal t for which there is a set \mathcal{E} of edges in X of size $k_X(v, w)$ separating v and w such that all members of \mathcal{E} exist in G_t .

PROOF. Note that $\underline{k} \geq k_X(v, w)$ if and only if for all $t \geq s$ we have $k_{G_t}(p_t(v), p_t(w)) \geq k_X(v, w)$. Now, for each $t \geq s$, apply Lemma 12.3.8 to the nice map $p_t: X \rightarrow G_t$.

Conversely, note $k_X(v, w) \geq \underline{k}$ if and only if for some $t \geq s$ we have $k_X(v, w) \geq k_{G_t}(p_t(v), p_t(w))$. Fix open edges e_1, \dots, e_k in X separating v from w . Specifically, say v is in C , w is in D , where C, D form of a clopen partition of $X \setminus \bigcup_i e_i$. Pick t sufficiently large that $t \geq s$ and p_t is a homeomorphism on each of the fixed edges (so, we can suppose e_1, \dots, e_k are edges in G_t). We claim that in G_t removing e_1, \dots, e_k separates $p_t(v)$ from $p_t(w)$. Otherwise, there is a $p_t(v) - p_t(w)$ path P in $G_t - \{e_1, \dots, e_k\}$. But then, due to the monotonicity of p_t , the subspace $p_t^{-1}(P)$ is a connected subset of $X - \{e_1, \dots, e_k\}$ containing both v and w , a contradiction.

Since every bonding map, π_n from G_n to G_{n-1} is nice, it follows from Lemma 12.3.8 that $(k_{G_n}(p_n(v), p_n(w)))_{n \geq s}$ is indeed decreasing. So it must stabilize at some t , with value $k_X(v, w)$. It follows that in X there are open edges E_1, \dots, E_k , where $k = k_X(v, w)$, separating v from w , such that these same edges exist in G_t . From the argument above we see that – as claimed – t is minimal for which there is a set \mathcal{E} of edges in X of size $k_X(v, w)$ separating v and w such that all members of \mathcal{E} exist in G_t . \square

12.4. The graph-like examples

In this section we construct families of examples which demonstrate that – with the sole exception of the characterization of 2-cc graph-like continua given in Proposition 12.3.5 – none of our positive results of Section 12.2 for n -ac and n -cc Freudenthal compactifications of locally finite graphs extend to arbitrary graph-like continua. Indeed we give continuum-sized families of examples which help demonstrate the difficulties involved in classifying n -ac and n -cc graph-like continua. Below we write K_m for the complete graph on m vertices.


12.4.1. A procedure for constructing graph-like continua.

Every graph-like continuum, X say, can be represented as an inverse limit, $\varprojlim G_k$, of connected graphs, with nice bonding maps. The k th bonding map, π_k , determines how to transition from G_{k+1} to G_k .

For the purposes of *constructing* a graph-like continuum, however, it is more convenient to have a rule for building G_{k+1} from G_k , and then specifying the bonding map. For our present purposes the following method is simple but effective.

The input data for the construction process are: (1) the first graph, G_1 , and (2) rules, one for each n , specifying how to replace a vertex, v , of degree n in a graph by a connected subgraph, G_v . Then to construct the inverse sequence, recursively apply the rules to the vertices of G_k to get G_{k+1} , and define the bonding map π_k to be the map which collapses each connected subgraph, G_v , in G_{k+1} to v in G_k . Clearly this map is nice.

By convention, if no rule is specified for vertices of degree n , then the rule is to leave the vertex alone. A typical rule for vertices of degree four is depicted below. Here each vertex of degree four is to be replaced with the complete graph on four vertices, and the four original edges are connected to one new vertex of the complete graph each. The bonding map collapses the new complete graph to the single old vertex.




12.4.2. Non-trivial ω -ac and ω -cc graph-like continua.

EXAMPLE 12.4.1. There is a graph-like continuum which is ω -cc but is not a graph (in particular, not the circle).

CONSTRUCTION. For each k we define recursively, 4-regular (multi) graphs G_k , following the procedure outlined above. The graph-like continuum $X = \varprojlim G_k$ will be ω -cc, but not a graph.

Let G_1 be any 4-regular connected multi-graph, for example the figure-eight graph (one vertex, two loops). The rules for constructing G_{k+1} from G_k are always the same: uncontract every vertex of G_k to a complete graph on four vertices, K_4 , in the natural manner (as above). This will have the effect that G_{k+1} will still be 4-regular, and so the recursion can be continued. The first three steps of the algorithm are depicted right.



It is obvious that X is not a graph. To see that $X = \varprojlim G_k$ is ω -cc, let $n \in \mathbb{N}$ be arbitrary, and note that by Lemma 12.3.1 it suffices to consider points x_1, \dots, x_n lying on (different) edges of X . Find $k \in \mathbb{N}$ sufficiently large such that x_1, \dots, x_n lie on different edges of G_k . Since G_k is 4-regular, it has an Eulerian cycle α . Since in G_{k+1} , every vertex of G_k is expanded into a K_4 , it is easy to see that the cycle α lifts to a simple closed curve α' of G_{k+1} , containing all vertices x_1, \dots, x_n . By Lemma 12.3.2, α' lifts to a simple closed curve α'' of X containing all vertices x_1, \dots, x_n , and so the proof is complete. \square

Note: for the above construction to produce an ω -cc graph-like continuum it suffices that (1) every G_k is Eulerian and (2) each vertex v in some G_k is uncontracted to G_v in G_{k+1} so that every edge in G_k incident to v is incident to distinct vertices in G_v , and those vertices are contained in a complete subgraph of G_v . (That each G_k is d_k -regular, and $(d_k)_k$ is constant, simplifies defining the expansion rules, but neither constraint is necessary.)

EXAMPLE 12.4.2. There is a graph-like continuum X , not a graph, which is ω -ac but not 2-cc.

CONSTRUCTION. Indeed, such examples can easily be constructed by considering a figure-eight-curve, a dumbbell, or a lollipop-curve, and replacing one of the circles in these graphs by a copy of the ω -cc graph-like continuum from the previous example. \square

THEOREM 12.4.3.

- (a) There are 2^{\aleph_0} many pairwise non-homeomorphic ω -cc graph-like continua.
- (b) There are 2^{\aleph_0} many pairwise non-homeomorphic ω -ac, but not 2-cc, graph-like continua.

PROOF. From Example 12.4.2 it is clear that (b) follows from (a).

Let G_1 be the graph on a single vertex with a single loop. Take any function $f \in \mathbb{N}^{\mathbb{N}}$ which is strictly increasing, and for every n we have $f(n)$ divisible by 2. Define $X_f = \varprojlim G_k^f$ where the graphs G_k^f are given recursively by:

- $G_1^f = G_1$, and
- G_{k+1}^f is obtained from G_k^f by uncontracting every vertex v of G_k^f to a $\tilde{K}_{f(k)} \supseteq K_{f(k)}$, where the edges incident with v are incident with distinct vertices of $\tilde{K}_{f(k)}$ and the remaining vertices of $K_{f(k)}$ get paired up, and get an additional parallel edge between each pair as to satisfy the even degree condition.

Note that, inductively, each G_{k+1}^f is a connected, $f(k)$ -regular graph (hence, as $f(k)$ is even, Eulerian), and this combined with the fact that f is strictly increasing and has even values ensures that G_{k+1}^f is well-defined from G_k^f .

The graphs G_k^f satisfy properties (1) and (2) noted after Example 12.4.1, from which it follows that the graph-like continuum X_f is ω -cc.

Claim 1: *If v and w are distinct vertices of G_{k+1}^f which are projected to the same vertex x of G_k^f , then $f(k) - 1 \leq k_{G_{k+1}^f}(v, w) \leq f(k)$.*

By $f(k)$ -regularity of G_{k+1}^f , the edge-connectivity is at most $f(k)$. The first inequality holds since the complete graph $K_{f(k)}$ has edge-connectivity $f(k) - 1$.

Claim 2: *If v and w are vertices of G_{k+1}^f such that their projections $v' = p_k(v)$ and $w' = p_k(w)$ are distinct in G_k^f , then $k_{G_{k+1}^f}(v, w) = k_{G_k^f}(v', w')$.*

By Lemma 12.3.8, it suffices to show $k_{G_{k+1}^f}(v, w) \geq k_{G_k^f}(v', w') = k$. To see this, note that $k_{G_k^f}(v', w') = k$ implies, by Menger's theorem [54, §3.3], that there is a collection of k -many edge-disjoint $v' - w'$ -paths in G_k^f . These paths lift, by the fact that we uncontracted vertices to complete graphs and by property (2), to a collection of k -many edge-disjoint $v - w$ -paths in G_{k+1}^f , establishing the claim.

Next, define $\mathcal{C}_f = \{k_{X_f}(v, w) : v \neq w \in V(X_f)\}$, the spectrum of all edge-connectivities between pairs of distinct vertices of X_f . From Claims 1 and 2, along with Lemma 12.3.10 we deduce:

Claim 3:

- (1) $\mathcal{C}_f \subseteq \{f(n) - 1 : n \in \mathbb{N}\} \cup \{f(n) : n \in \mathbb{N}\}$, and
- (2) for each $n \in \mathbb{N}$ we have $\{f(n) - 1, f(n)\} \cap \mathcal{C}_f \neq \emptyset$.

Now define $\mathcal{F} = \{f \in \mathbb{N}^{\mathbb{N}} : f \text{ is strictly increasing and } \forall n f(n) \text{ is even}\}$. Then $|\mathcal{F}| = 2^{\aleph_0}$. For each $f \in \mathcal{F}$ we know $X_f = \varprojlim G_k^f$ is an ω -cc graph-like continuum, and we now show these are pairwise non-homeomorphic.

Claim 4: *For distinct $f \neq g \in \mathcal{F}$, the graph-like continua X_f and X_g are non-homeomorphic.*

To see this, let $k \in \mathbb{N}$ be minimal such that $f(k) \neq g(k)$, and without loss of generality assume that $f(k) < g(k)$. Note that $k \geq 2$ (since $G_1^f = G_1^g$). As f, g are strictly increasing and have even values, we have $f(k - 1) = g(k - 1) < f(k) - 1 < f(k) < g(k) - 1 < g(k)$.

Hence, from Claim 3, one of $f(k) - 1$ and $f(k)$ is in \mathcal{C}_f but neither is in \mathcal{C}_g , so $\mathcal{C}_f \setminus \mathcal{C}_g \neq \emptyset$, and so we deduce $X_f \not\cong X_g$ by Lemma 12.3.9. \square

12.4.3. Graph-like continua which are n - but not $(n + 1)$ - ac or cc. In this section, we construct interesting graph-like continua which are n -ac but not $(n + 1)$ -ac, and others which are n -cc but not $(n + 1)$ -cc. For these, we present two fundamentally different constructions.

The first construction uses knowledge about certain closed or open Eulerian paths in finite minors of the graph-like space. In some sense, this first construction is all about controlling the edge-cuts in the space. The second construction starts with several copies of a graph-like space, in which we have a lot of control over which arcs we may use to pick up our favorite edge set. We then glue together these copies by identifying some finite set of vertices. In some sense, this second construction is all about controlling the vertex-cuts in the space.

12.4.3.1. *Technique 1: Using open and closed Eulerian paths in finite graphs.* For our next examples, we need the following auxiliary result. Recall that a *matching* in a graph is a collection of pairwise non-adjacent edges.

LEMMA 12.4.4. *For every $n \geq 2$, the complete graph on $N \geq 4n + 4$ vertices has the property that given (i) any matching M in K_N , (ii) any edges e_1, \dots, e_k of $K_N - M$ with $k \leq n$, and (iii) any two vertices v, w in K_N , there is a non-edge-repeating trail from v to w in $K_N - M$ containing the selected edges.*

PROOF. To see the claim, note that after removing the matching M , every vertex has degree at least $N - 2$ in the subgraph $H_0 = K_N - M$, and so any two vertices have at least $N - 4$ common neighbours in H_0 . Write $e_i = x_i y_i$. Since v and x_1 have a common neighbour, there is a path P_1 from v to y_1 with $e_1 \in E(P_1)$. Next, consider $H_1 = H_0 - E(P_1)$ and note that every vertex in H_1 has degree at least $N - 4$, and so any two vertices have at least $N - 8 \geq 4n - 4 > 0$ common neighbours in H_1 . If e_2 isn't yet covered by P_1 , find a path P_2 in H_1 from y_1 to y_2 containing the edge e_2 . If we continue in this manner, then in $H_k = H_0 \setminus \bigcup_{i \leq k} E(P_i)$, every vertex has degree at least $N - 2 - 2k \geq N/2$. Hence, any two vertices in H_k are either connected by an edge, or have a common neighbour. Thus, there is a path P_{k+1} in H_k from y_k to v . It is clear that $\bigcup_{i \leq k+1} P_i$ is the desired edge trail. \square

EXAMPLE 12.4.5. For each $n \geq 2$ there is a graph-like continuum which is n -ac but not $(n + 1)$ -ac.

CONSTRUCTION. Fix $n \geq 2$. We define a sequence of graphs, G_k^n , by giving the first, G_1^n , then G_2^n , and a rule defining G_{k+1}^n from G_k^n , for $k \geq 2$. This naturally gives an inverse limit $X_n = X = \varprojlim G_k^n$ which is graph-like.

Case 1: $n = 2m + 1$ is odd where $m \geq 1$. The graph G_1^n has four vertices, v_1, w_1, w_2 and v_2 . There is an edge connecting v_i to w_i for $i = 1, 2$; and $2m$ edges connecting w_1 and

w_2 . Thus G_1^n has $n + 1$ edges, two vertices of degree 1 and two of degree n . It is easy to check that G_1^n is n -oE. But G_1^n is not $(n + 1)$ -oE, and so by Proposition 12.3.4(b) X is not $(n + 1)$ -ac. Next, let $N = N(n)$ be large enough as to satisfy Lemma 12.4.4. To define G_2^n from G_1^n leave the two vertices of degree 1 alone, and uncontract the two vertices of degree n to a K_N , such that all vertices of G_2^1 are either of degree 1, $N - 1$, or N . To define G_{k+1}^n from G_k^n leave the (two) vertices of degree 1 alone, and replace all vertices of degree $N - 1$ or N by a complete graph on N new vertices. Since all vertices of G_k^1 are either of degree 1, $N - 1$ or N , inductively, the same is true for G_k^n , and then G_{k+1}^n . Hence the definition is complete.

We now show by induction on k that for all k the graph G_k^n is n -oE. Then the proof that X is n -ac then follows as in the previous examples. Fix $k \geq 2$. Let $\pi = \pi_k: G_{k+1}^n \rightarrow G_k^n$ be the bonding map. Take any subset S of G_{k+1}^n containing no more than n points. Then, inductively, in G_k^n there is an edge-disjoint trail containing $\pi(S)$. The edges in this trail pull back to an edge-disjoint sequence of (directed) edges in G_{k+1}^n so that successive edges have end and start points (respectively) mapping to the same vertex in G_k^n . We explain how to add edges in fibers of vertices of G_k^n so as to form an edge-disjoint trail in G_{k+1}^n containing the points of S .

It suffices to consider one vertex v of G_k^n , and add edges in $\pi^{-1}\{v\}$ so as to connect together successive edges in the edge-disjoint sequence while preserving edge-disjointness and ensuring that all points in S which happen to lie in $\pi^{-1}\{v\}$ are contained in the resulting trail. If $\pi^{-1}\{v\}$ is just one point then there is nothing to do. Otherwise $\pi^{-1}\{v\}$ is a complete graph on N vertices. If no edges in the edge-disjoint sequence meet $\pi^{-1}\{v\}$ there is nothing to do. List all successive pairs entering and exiting $\pi^{-1}\{v\}$ as $e_1^0, e_2^0, e_1^1, e_2^1, \dots, e_1^p, e_2^p$, where $p \geq 0$. Let f_1, \dots, f_q be the edges in $\pi^{-1}\{v\}$ containing points of S . Note $q \leq n$.

For $i = 1, \dots, p - 1$ add the edge in $\pi^{-1}v$ connecting the end of e_1^i to the start of e_2^i . By construction, this edge set is a matching M . If at this point, some of the edges f_i are yet uncovered, we may add, by Lemma 12.4.4, a trail from the end of e_1^p to the start of e_2^p disjoint from M in $\pi^{-1}v$ containing all uncovered edges of f_1, \dots, f_q . Otherwise, simply add the edge in $\pi^{-1}v$ connecting the end of e_1^p to the start of e_2^p . Now we are done.

Case 2: $n = 2m$ is even where $m \geq 1$. The graph G_1^n has four vertices, v_1, w_1, w_2, v_2 . There are $n - 1$ edges connecting w_1 and w_2 , and one edge from each of v_1 and v_2 to w_1 . Then G_1^n has $n + 1$ edges, two vertices of degree 1, one of degree $n + 1$ and one of degree $n - 1$. It is easy to check that G_1^n is n -oE but not $(n + 1)$ -oE.

Let $N = N(n + 1)$ be large enough as to satisfy Lemma 12.4.4 for $n + 1$. Define G_2^n by replacing the single vertex of degree $n + 1$ with N new vertices connected by a complete graph, but leaving the other vertices alone. To define G_{k+1}^n from G_k^n leave the two vertices of degree 1 alone, leave the vertex of degree $n - 1$ alone, and replace all vertices of degree N

or $N - 1$ with N new vertices and a complete graph connecting them. Now the argument that $X = \varprojlim G_k^n$ is as required is very similar to that given above in Case 1. \square

EXAMPLE 12.4.6. For each even n there is a graph-like continuum which is n -cc but not $(n + 1)$ -cc.

CONSTRUCTION. The argument is similar to that given above for graph-like continua which are n -ac but not $(n + 1)$ -ac. So we give a sketch only, highlighting differences.

Fix even n . Let G_1^n be the (multi-)graph with two vertices and $n + 1$ parallel edges connecting them. Note that the vertices have degree $n + 1$, and it is easy to check G_1^n is n -E (given any n points there is a closed edge-disjoint trail containing them). Pick $N = N(n + 1)$ be large enough as to satisfy Lemma 12.4.4 for $n + 1$. Recursively define G_{k+1}^n from G_k^n by uncontracting each vertex to a K_N . By induction one can check that every G_k^n is n -E.

Define $X = \varprojlim G_k^n$. Then X is a graph-like continuum, and arguing as before it can be verified to be n -cc. But picking a point from the interior of each edge easily shows G_1^n is not $(n + 1)$ -E. Hence, by Proposition 12.3.4, X is not $(n + 1)$ -cc. \square

Our strategy from above is bound to fail when trying to build an example for a graph-like continuum which is n -cc but not $(n + 1)$ -cc for odd n . Indeed, given odd n we would need graphs which are n -E but not $(n + 1)$ -E, however the second author and Knappe have shown that this is impossible – any graph which is n -E, where n is odd, is automatically $(n + 1)$ -E, see [101]. Hence, a fundamentally different approach is required to construct, for odd n , graph-like continua which are n -cc but not $(n + 1)$ -cc. This is the purpose of our next and final section.

12.4.3.2. *Technique 2: Using small vertex cuts in graph-like spaces.* Recall that in an $n + 1$ -ac graph-like continuum, deleting $n - 1$ vertices creates at most n distinct connected components, [76, Lemma 2.3.3]

A similar result holds for $(n + 1)$ -cc graphs: Recall that a connected graph, or a graph-like continuum G is called k -tough, if for any finite, non-empty set of vertices S , the number of components of $G - S$ is at most $|S|/k$. Adapting this notion slightly, let us say that a graph-like continuum G is (k, n) -tough if for any set of vertices S with $1 \leq |S| \leq n$, the number of components of $G - S$ is at most $|S|/k$.

The standard notion of toughness plays a well-known role in the theory of Hamilton cycles, as a necessary condition for a finite graph to be Hamiltonian is that it is 1-tough, [50, Prop. 2.1]. The straightforward adaptation of this result to our use case gives the following observation.

LEMMA 12.4.7. *Every $(n + 1)$ -cc graph-like continuum is $(1, n)$ -tough.*

PROOF. Suppose X is an n -cc graph-like continuum and, for a contradiction, $S \subseteq V(X)$ is a finite vertex set with $1 \leq |S| = s \leq n$ whose removal leaves strictly more than

s components. Pick $s + 1$ edges in different components of $X - S$. As $s + 1 \leq n + 1$, by assumption, there is a simple closed curve α in X picking up the edges. But then $\alpha \setminus S$ consists of at most s components. Hence, there are two edges in the same component of $\alpha \setminus S$, contradicting the fact that they lie in different components of $X - S$. \square

As our building blocks, we will use the following class of graphs.

EXAMPLE 12.4.8. For each $n \geq 2$ there is a graph-like continuum X containing vertices v_1, v_2, \dots, v_n such that (i) whenever an edge set $F \subseteq E(X)$ with $|F| \leq n$ is chosen, and (ii) any two vertices $v_i \neq v_j$ from our list are chosen, there is an $v_i - v_j$ arc α in X containing F but not v_k for all $k \neq i, j$.

PROOF. Let $n \in \mathbb{N}$ be fixed and consider $N = N(n)$ from Lemma 12.4.4. We will construct X as an inverse limit of finite graphs G_n where we start with $G_1 = K_N$, and uncontract in each step every vertex v of G_k to a new K_N . It follows recursively that every vertex of G_k has degree N or $N - 1$.

Let $p_k: X \rightarrow G_k$ denote the quotient map. Choose $v_1, \dots, v_n \in V(X)$ subject to the condition that the degree of $p_k(v_i)$ equals $N - 1$ for each $k \in \mathbb{N}$. Now pick any edge set F with $|F| \leq n$. We will demonstrate that there is an $v_1 - v_2$ arc α in X with $F \subseteq \alpha$ and $v_i \notin \alpha$ for all $i \geq 3$.

By Lemma 12.4.4, there is a $p_1(v_1) - p_1(v_2)$ -trail T_1 in G_1 containing $F \cap E(G_1)$. Recursively, using again Lemma 12.4.4, extend this to an $p_k(v_1) - p_k(v_2)$ -trail T_k in G_k containing $F \cap E(G_k)$ until $F \cap E(G_k) = F$. Next, using the fact that $p_{k+1}(v_i)$ equals $N - 1$, extend T_k to an $p_{k+1}(v_1) - p_{k+1}(v_2)$ -path T_{k+1} in G_{k+1} missing all $p_{k+1}(v_i)$ for all $i \geq 3$. Extending this path T_{k+1} recursively, it is clear that we end up with the desired $v_1 - v_2$ -arc. \square

EXAMPLE 12.4.9. For each $n \geq 2$ there is a graph-like continuum which is n -cc but not $(n + 1)$ -cc.

CONSTRUCTION. Let X be the space from Example 12.4.8 with special points v_1, \dots, v_n . Now take $n + 1$ many disjoint copies $X^{(1)}, \dots, X^{(n+1)}$ of the space X with the special points denoted by $v_1^{(i)}, \dots, v_n^{(i)} \in V(X^{(i)})$.

We claim the graph-like continuum

$$Z = (X^{(1)} \oplus \dots \oplus X^{(n+1)}) / \sim \quad \text{where } v_k^{(1)} \sim v_k^{(2)} \sim \dots \sim v_k^{(n+1)} \quad \text{for each } k,$$

is n -cc but not $(n + 1)$ -cc. Let us write $[v_k] \in Z$ for the vertex corresponding to the equivalence class of $v_k^{(1)}$. Then it is clear from the construction that deleting $S = \{[v_1], \dots, [v_n]\}$ from Z leaves $n + 1$ many components. Therefore, Z is not $(1, n)$ -tough, and hence cannot be $(n + 1)$ -cc by Lemma 12.4.7.

To see that Z is n -cc, consider any collection $F = \{e_1, e_2, \dots, e_n\}$ of n edges of Z (which is sufficient because of Lemma 12.3.1). We may assume that the edges are contained in the

first i spaces $X^{(1)} \cup \dots \cup X^{(i)}$ where $i \leq n$. By the properties guaranteed by example 12.4.8, we can find $v_j^{(j)} - v_{j+1}^{(j)}$ arcs $\alpha^{(j)}$ (where $i + 1 \equiv 1$) in $X^{(j)}$ missing all other special vertices and containing $F \cap E(X^{(j)})$. It is then clear that $\alpha := \bigcup_{j \leq i} \alpha^{(j)} \subseteq Z$ is the desired simple closed curve in Z containing F (as each α^j and α^{j+1} end and start at the same vertex $[v_{j+1}] \in Z$ respectively, and α^j and α^ℓ are disjoint for $|(j - \ell \pmod n)| \geq 2$). \square

THEOREM 12.4.10. *For every $n \geq 2$:*

(a) $_n$ *there are 2^{\aleph_0} many non-homeomorphic graph-like continua which are n -ac but not $(n + 1)$ -ac, and*

(b) $_n$ *there are 2^{\aleph_0} many non-homeomorphic graph-like continua which are n -cc but not $(n + 1)$ -cc.*

PROOF. This follows by the same method as we derived Theorem 12.4.3 (a) from Example 12.4.1 with some small adjustments that we show here.

Fix n . Both techniques to construct ‘ n -ac not $(n + 1)$ -ac’ and ‘ n -cc not $(n + 1)$ -cc’ graph-like continua used Lemma 12.4.4 to replace vertices by a big enough K_N where N depended on n .

As in Theorem 12.4.3, let $\mathcal{F} = \{f \in \mathbb{N}^{\mathbb{N}} : f \text{ is strictly increasing, } \forall n f(n) \text{ is divisible by } 4, \text{ and } f(1) \geq N\}$. Then $|\mathcal{F}| = 2^{\aleph_0}$. To define the sequence of graphs, G_k^f , at step $k + 1$ uncontract vertices in the k th step into a $K_{f(k)}$.

Then $X_f = \varprojlim G_k^f$ is a graph-like continuum with the requisite combination of strong connection properties (‘ n -ac not $(n + 1)$ -ac’ or ‘ n -cc not $(n + 1)$ -cc’). And, as in the proof of Theorem 12.4.3, for distinct f and g from \mathcal{F} the spaces X_f and X_g have different edge-connection spectra, and so are non-homeomorphic.

In all cases *except* for the construction of an n -cc not $(n + 1)$ -cc graph-like continuum where n is odd, these X_f are as needed. But for ‘odd n , n -cc not $(n + 1)$ -cc’ we require an extra step as in Example 12.4.9, to get Z_f for f from \mathcal{F} . So it remains to show that for distinct f and g from \mathcal{F} the spaces Z_f and Z_g are non-homeomorphic.

However Z_f is obtained by gluing $(n + 1)$ copies of X_f together over an n -point set, call it S_f , so this set is a vertex separator of size n in Z_f . From Proposition 12.3.7 we know that each X_f has vertex connectivity $\geq f(1) \geq N > n$, hence S_f is the *unique* vertex separator of Z_f of size n . Since this separator must be preserved by any homeomorphism we see that indeed distinct f and g yield topologically distinct Z_f and Z_g . \square

Graph-like compacta: characterizations and Eulerian loops

A *compact graph-like* space is a triple (X, V, E) where X is a compact, metrizable space, $V \subseteq X$ is a closed zero-dimensional subset, and E is an index set such that $X \setminus V \cong E \times (0, 1)$. New characterizations of compact graph-like spaces are given, connecting them to certain classes of continua, and to standard subspaces of Freudenthal compactifications of locally finite graphs. These are applied to characterize Eulerian graph-like compacta.

13.1. Introduction

Locally finite graphs can be compactified, to form the Freudenthal compactification, by adding their ends. This topological setting provides what appears to be the ‘right’ framework for studying locally finite graphs. Indeed, many classical theorems from finite graph theory that involve paths or cycles have been shown to generalize to locally finite infinite graphs in this topological setting, while failing to extend in a purely graph theoretic setting. See the survey series [53]. More recently, compact graph-like spaces were introduced by Thomassen and Vella, [153], as a natural class encompassing graphs, and in particular containing the standard subspaces of Freudenthal compactification of locally finite graphs.

A *compact graph-like* space is a triple (X, V, E) where: X is a compact, metrizable space, $V \subseteq X$ is a closed zero-dimensional subset, and E is a discrete index set such that $X \setminus V \cong E \times (0, 1)$. The sets V and E are the *vertices* and *edges* of X respectively. More generally, a topological space X is compact graph-like, if there exists $V \subseteq X$ and a set E such that (X, V, E) is a compact graph-like space. Recall that connected compact metrizable spaces are called *continua*, and so a *graph-like continuum* is a continuum which is graph-like.

Papers in which graph-like spaces have played a key role include [153] where several Menger-like results are given, and [49] where algebraic criteria for the planarity of graph-like continua are presented. In [30], aspects of the matroid theory for graphs have been generalized to infinite matroids on graph-like spaces.

In this paper we present two groups of new results. The first group consists of characterizations of compact graph-like spaces and continua. These connect graph-like continua to certain classes of continua which have been intensively studied by continua theorists. We also establish that compact graph-like spaces are not simply ‘like’ the Freudenthal

compactifications of locally finite graphs, but in fact *are* standard subspaces of the latter. Our second group of results consists of various characterizations of when a graph-like continuum is Eulerian. These naturally extend classical results for graphs.

13.1.1. The Main Theorems. In Section 13.2 we give various characterizations and representations of compact graph-like spaces, and graph-like continua, which demonstrate that graph-like continua form a class of continua which are also of considerable interest from the point of view of continua theory. These results can be summarized as follows.

THEOREM (A). *The following are equivalent for a continuum X :*

- (i) X is graph-like,
- (ii) X is regular and has a closed zero-dimensional subset V such that all points outside of V have order 2,
- (iii) X is completely regular,
- (iv) X is a countable inverse limit of finite connected multi-graphs with onto, monotone, simplicial bonding maps with non-trivial fibres at vertices only,
- (iv)' X is a countable inverse limit of finite connected multi-graphs with onto, monotone bonding maps that project vertices onto vertices, and
- (v) X is homeomorphic to a connected standard subspace of a Freudenthal compactification of a locally finite graph.

Here a continuum is *regular* if it has a base all of whose members have finite boundary, and *completely regular* if all non-trivial subcontinua have non-empty interior. A map is *monotone* if all fibres are connected, while a map between graph-like spaces is *simplicial* if it maps vertices to vertices, and edges either homeomorphically to another edge, or to a vertex. A *standard subspace* of a compact graph-like space is a closed subspace that contains all edges it intersects. The equivalence of (i) and (ii) is analogous to a well-known topological characterization of finite graphs, namely a continuum is a graph if and only if every point has finite order, and all but finitely many points have order 2, [121, Theorem 9.10 & 9.13]. The equivalence of (i) and (iv) provides a powerful tool to lift results in finite graph theory to graph-like continua. Indeed this is key to our results on Eulerian paths and loops below. It also is key to the equivalence of (i) and (v). The equivalence of (i) and (iii) yields a purely internal topological characterization of graph-like continua, without any reference to distinguished points, ‘vertices’, or subsets, ‘edges’.

We prove all of Theorem (A) taking ‘compact graph-like space’ as the basic notion. Because ‘compact graph-like’ takes a middle ground between topology and graph theory, our proofs are clean and efficient. However it is important to note that the equivalence of (i) and (iii) follows, modulo some basic lemmas, from a result of Krasinkiewicz, [104], while the implication (iii) implies (iv) is essentially shown by Nikiel in [130]. Nikiel also claimed, without proof, the converse implication. Regarding (v), Bowler et al. have

claimed, without proof, the weaker assertion that every compact graph-like space is a minor (essentially: a quotient) of a Freudenthal compactification of some locally finite graph, [30, p. 6].

In Sections 13.3 and 13.4 we extend some well-known characterizations of Eulerian graphs to graph-like continua. Let G be a (multi-)graph. A *trail* in G is an edge path with no repeated edges. It is *open* if the start and end vertices are distinct, and *closed* if they coincide. We also call closed trails *circuits*. A *segment* is a trail which does not cross itself. A *cycle* is a circuit which never crosses itself. A trail is *Eulerian* if it contains all edges of the graph. (Note that an *Eulerian circuit* is a closed Eulerian trail.) The graph G is *Eulerian* (respectively, *closed Eulerian*) if it has an open (respectively, closed) Eulerian trail; and *Eulerian* if it either open or closed Eulerian. Call a vertex v of a graph G *even* (respectively, *odd*) if the degree of v in G is even (respectively, odd).

Classical results of Euler and Veblen characterize multi-graphs with closed, and respectively, open, Eulerian trails as follows. Let G be a connected multi-graph with vertex set V , then the following are equivalent: (i) G is closed [open] Eulerian, (ii) every vertex is even [apart from precisely two vertices which are odd], (iii) [there are vertices $x \neq y$ such that] for every bi-partition of V , the number of cross edges is even [if and only if x and y lie in the same part], and (iv) the edges of G can be partitioned into edge disjoint cycles [and a non-trivial segment]. We extend these results to compact graph-like spaces, and prove the following result.

THEOREM (B). *Let X be a graph-like continuum with vertices V . The following are equivalent:*

- (i) X is closed [open] Eulerian,
- (ii) every vertex is even [apart from precisely two vertices which are odd],
- (ii)' every vertex has strongly even degree [apart from precisely two vertices which have strongly odd degree],
- (iii) [there are vertices $x \neq y$ such that] for every partition of V into two clopen pieces, the number of cross edges is even [if and only if x and y lie in the same part], and
- (iv) the edges of X can be partitioned into edge-disjoint circles [and a non-trivial arc].

Further, if X is closed [open] Eulerian then either X has continuum many distinct Eulerian loops [Eulerian paths], or has a finite number of distinct Eulerian loops [Eulerian paths], which occurs if and only if X is homeomorphic to a finite closed [open] Eulerian graph.

Let X be a compact graph-like space with set of vertices V . A subspace of X is called an *arc* if it is homeomorphic to $I = [0, 1]$, and a *circle* if it is homeomorphic to the circle, S^1 . A (standard) *path* is a continuous map $f: I \rightarrow X$ such that $f(0)$ and $f(1)$ are vertices,

f is injective on the interior of every edge and $f^{-1}(V)$ has empty interior. Note that every continuous map $f: I \rightarrow X$ with $f(0), f(1)$ as vertices is homotopy equivalent (with fixed endpoints) to a path. Also note that if X is a graph (with usual topology), then every path yields a corresponding trail, and every trail corresponds to a path. A path, f , is *open* if $f(0) \neq f(1)$, and *closed* if $f(0) = f(1)$. Closed paths are called *loops*. A path (or loop) is *Eulerian* if its image contains every edge. Note that in a graph with the usual topology there is a natural correspondence between Eulerian paths and Eulerian trails, and Eulerian circuits and Eulerian loops. We abbreviate ‘closed and open’ to ‘clopen’. A vertex v is *odd* (resp. *even*) if and only if there exists a clopen subset A of V containing v , such that for every clopen subset C of the vertices V with $v \in C \subseteq A$ the number of edges between C and $V \setminus C$ is odd (resp. even).

The equivalence of (i), (ii), (iii) and (iv) in Theorem (B) is established in Section 13.3.1. At the heart of our proof is our representation of compact graph-like spaces as inverse limits, and an induced inverse limit representation of all Eulerian loops (possibly empty, of course). The ‘further’ part of Theorem (B) follows in Section 13.3.2 from topological considerations of the space of all Eulerian loops.

Our definition of ‘even’ and ‘odd’ vertices is natural within the context of Theorem (B). An alternative approach to degree, due to Bruhn and Stein [39], leads to the notions of ‘strongly even degree’ and ‘strongly odd degree’ appearing in item (ii)’ of Theorem (B). See Section 13.4 for details and the proof that ‘(ii) implies (ii)’ and ‘(ii)’ implies (i)’.

Alternative Paths. Theorem (A) shines an unexpected light on connections between concepts from continua theory (completely regular continua and inverse limits of graphs), and concepts arising from infinite graph theory (graph-like continua, Freudenthal compactifications of graphs, and their standard subspaces). As a result we discover that the various parts of Theorem (B) generalize numerous results in the literature, and—with the considerable assistance of the machinery developed here—Theorem (B) can be derived from older work.


For Freudenthal compactification of graphs, the equivalence of (i), (iii) and (iv) is due to Diestel & Kühn, [56, Theorem 7.2], while the equivalence with (ii)’ is due to Bruhn & Stein, [39, Theorem 4]. For standard subspaces of Freudenthal compactification of graphs, the equivalence of (iii) and (iv) is due to Diestel & Kühn, [57, Theorem 5.2], the equivalence of (i) and (iv) is due to Georgakopoulos, [80, Theorem 1.3], and the equivalence (ii)’ and (iii) is due to Berger & Bruhn, [20, Theorem 5]. It should also be noted that the method of lifting Eulerian paths and loops via inverse limits, used by Georgakopoulos, was previously introduced by Bula et al. in [41, Theorem 5].

Thus an alternative path to proving Theorem (B) is as follows. Let X be a graph-like continuum. According to the equivalence of (i) and (v) in Theorem (A), which depends on the equivalence of (i) and (iv), X is homeomorphic to a standard subspace of a Freudenthal compactification of a graph. Now equivalence of (i), (ii)’, (iii) and (iv) follows from


the results cited immediately above. To add equivalence of (ii) apply Theorem 13.4.1 (which uses Theorem (A) (i) \iff (iv), and a non-trivial inverse limit argument) and Lemmas 13.3.4 and 13.4.2. Although this alternative path exists, the direct proofs given here in Sections 13.3 and 13.4, using compact graph-like spaces as the basic notion, are—in the authors' view—much shorter and more natural.

13.1.2. Examples. Before proving our results on graph-like continua, we now introduce some examples. With these examples we have three objectives. First show a little of the variety of graph-like continua. Second elucidate some of the less familiar terms in Theorem (B), in particular 'even' and 'odd' vertices. Third demonstrate the remarkable complexity of Eulerian loops and paths in graph-like continua. This complexity highlights the hidden depths of Theorem (B).


Example 1. The two-way infinite ladder with single diagonals, which is the infinite graph G shown below. Notice that all its vertices are even, but it has no Eulerian loop.



The Freudenthal compactification, γG , of G adds two ends. Then, as shown in the diagram, γG has an open Eulerian path from one end to the other.




It follows from Theorem (B) that the two ends are odd. We now demonstrate that the left end is odd from the definition. For the clopen neighborhood A take the left end along with all vertices to the left of some rung of the ladder. Now consider an arbitrary clopen C containing the left end and contained in A (depicted by the green vertices in the diagram below). Then C is the disjoint union of a C_0 which contains the end and all vertices to the left of a rung, and a C_1 which is a finite subset of $A \setminus C_0$. In the diagram we see that the number of edges from C to $V \setminus C$ is 9, which is odd. In general, if we identify C_0 (and all edges between members of C_0) to a vertex v and identify $V \setminus A$ to a vertex w , then we get a finite graph with exactly two odd vertices (namely, v and w , both of degree 3). Hence from the equivalence of (ii) and (iii) of the graph version of Theorem (B), we see that the number of edges from C to $V \setminus C$ is odd – as required for the left end to be odd.



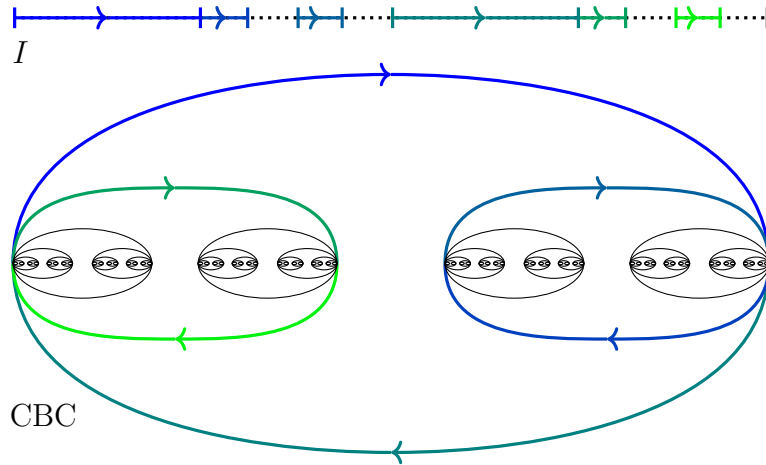
For the next two examples let C denote the standard ‘middle thirds’ Cantor subset of I .

Example 2. The Cantor bouquet of semi-circles, CBS. The vertices are $C \times \{0\}$ in the plane, along with semi-circular edges centered at the midpoint of each removed open interval. Note that CBS is not a Freudenthal compactification of graph.



The vertex $\mathbf{0} = (0, 0)$ is neither odd nor even, and hence CBS is not Eulerian. Indeed, as indicated on the diagram, there is one (odd) edge connecting all the vertices in the ‘left half’ of the vertices to its complement (the ‘right half’), but two (even) edges connecting the ‘left quarter’ to its complement. Similarly we see that every clopen neighborhood of $\mathbf{0}$ contains two clopen neighborhoods of $\mathbf{0}$ of which one has an odd number of edges to its complement, and the other an even number.

Example 3. The Cantor bouquet of circles, CBC, can be obtained from the Cantor bouquet of semi-circles by reflecting it in the real axis. One can check that all vertices are even. The diagram illustrates an Eulerian loop in CBC.



Suppose $f: I \rightarrow X$ is a standard path in a graph-like continuum X with vertices V and (open) edges $(e_n)_n$. Then $f^{-1}(V)$ is a closed nowhere dense subset of I , and its complement, $f^{-1}(\bigcup_n e_n)$ is dense and a disjoint union of open intervals. This countable family, $\{f^{-1}(e_n) : n \in \mathbb{N}\}$ inherits an order from the order on I . So to every path f we can associate a countable linear order L_f , which we informally call the *shape* of f .

To illustrate this, consider $L = L_f$ where f is the Eulerian loop in the Cantor bouquet of circles diagrammed above. Then f traverses the top edge from left to right, covers the right-hand copy of CBC, traverses the bottom edge from right to left, and then covers the left-hand copy of CBC. So L satisfies the equation $L = 1 + L + 1 + L$. It follows that L is an infinite ordinal. Thus $L = 1 + L$, and we see $L = L + L$. The first infinite ordinal which is a fixed point under addition of linear orders is the ordinal ω^ω . Hence $L = \omega^\omega$.

We now see how to construct for each countable linear order L a graph-like continuum X_L with an Eulerian loop f so that $L_f = L$. To do so recall: every countable linearly ordered set L can be realized (is order isomorphic to) a countable family of disjoint open subintervals of I , with dense union. For further material on the interaction of linear orders and graph-like compacta, see [30, §4].


Given a line segment, S , in the plane the ‘circle with diameter S ’ is the circle with center the midpoint of the line segment, and radius half the length of the segment.

Example 4. Let L be a countable linear order. Fix a family \mathcal{U} of pairwise disjoint open subintervals of I , with dense union, which is order isomorphic to L . Define X_L to be the subspace of the plane obtained by starting with $X = I \times \{0\}$, and for each U in \mathcal{U} , remove $U \times \{0\}$ from X and add the circle with diameter $U \times \{0\}$.

The Eulerian loops on X_L are naturally bijective with all functions $\varrho: L \rightarrow \{\pm 1\}$. To see this take any $\varrho: L \rightarrow \{\pm 1\}$. Since \mathcal{U} and L are isomorphic we can think that the domain of ϱ is actually \mathcal{U} . Define $g_\varrho: [0, 1] \rightarrow X_L$ by requiring (i) $g(t) = t$ on $I \setminus \bigcup \mathcal{U}$, and (ii) on U in \mathcal{U} the path g traverses the top (resp. bottom) semi-circle in X_L corresponding to U if $\varrho(U) = +1$ (resp. $\varrho(U) = -1$). Now define f_ϱ – the desired Eulerian loop – by $f_\varrho(t) = g_\varrho(2t)$ on $[0, 1/2]$ and $f_\varrho(t) = g_{-\varrho}(2 - 2t)$ on $[1/2, 1]$. Informally, on $[0, 1/2]$ the


path f_ϱ travels from left to right along X_L crossing the circles by either taking the upper or lower semi-circles depending on ϱ ; and then on $[1/2, 1]$ it travels across X_L from right to left taking the opposite upper/lower semi-circles than before. Every Eulerian loop arises in this way, and observe that they all have the same shape, L .

The following diagram depicts X_Q where Q is the linearly ordered set of dyadic rationals in $(0, 1)$. Recall that Q is order isomorphic to the rationals, \mathbb{Q} .



The graph-like continuum X_Q provides an example of the difficulties involved in naively trying to lift arguments for graphs to graph-like continua. In the standard proof of Theorem (B) for graphs one moves from (iv) ‘the edges of the graph can be decomposed into disjoint cycles’ to (i) ‘there is an Eulerian circuit’ by amalgamating the cycles, one after another to form the circuit. Notice that in X_Q there is a canonical decomposition of X_Q into edge disjoint circles – namely the circles in the definition of X_Q . But these circles are *pairwise disjoint*. Hence there is no natural method of amalgamating them into an Eulerian loop for X_Q .

Example 5. The Hawaiian earring, H , is also Eulerian. Unlike the X_L examples above, every countable linear order can be realized as the L_f of an Eulerian loop. Write H as $\mathbf{0} = (0, 0)$ (the sole vertex) and the union of circles in the plane C_n , for $n \in \mathbb{N}$, where C_n has radius $1/n$ and is tangential at $\mathbf{0}$ to the x -axis.



We can identify the Eulerian loops in the Hawaiian earring as follows. For any countable linear order L and function $\varrho: L \rightarrow \mathbb{N} \times \{\pm 1\}$ such that $\pi_1 \circ \varrho: L \rightarrow \mathbb{N}$ is a bijection, there is a naturally corresponding Eulerian loop f_ϱ of H . Indeed, given L and ϱ , let \mathcal{U} be a family of pairwise disjoint open subintervals of I , with dense union, which is order isomorphic to L (and identify them). Define f_ϱ to have value $\mathbf{0}$ on the complement of $\bigcup \mathcal{U}$, and on U in \mathcal{U} , writing $\varrho(U) = (n, i)$, it should traverse C_n clockwise (respectively, anticlockwise) if $i = +1$ (respectively, $i = -1$). One can check all Eulerian loops arise this way.

13.2. Properties and characterizations of graph-like continua

13.2.1. Basic Properties. Most of the following basic properties of graph-like spaces are well-known, see e.g. [153]. Nonetheless, it might be helpful to give a self-contained outline of the most important properties we use.

Let (X, V, E) be a compact graph-like space. We often identify the label, e , of an edge, with the subspace $e \times (0, 1)$ of X . Note that since V is zero-dimensional, for every edge e , the closure, \bar{e} , of e adds at most two vertices – the ends of the edge – and \bar{e} is homeomorphic to the circle, S^1 , or $I = [0, 1]$. With this in mind, our definition of compact graph-like space is the same as the original in [153].

A *separation* (A, B) of a graph like space X is a partition of $V(X)$ into two disjoint clopen subsets. The *cut induced by the separation* (A, B) is set of edges with one end vertex in A and the other in B , denoted by $E(A, B)$. More generally, we call a subset $F \subseteq E$ a cut if there is a separation (A, B) of X such that $F = E(A, B)$. A *multi-cut* is a partition $\mathcal{U} = \{U_1, U_2, \dots, U_n\}$ of $V(X)$ into pairwise disjoint clopen sets. For each two U_i, U_j , not necessarily different, $E(U_i, U_j)$ denotes the set of edges with one endpoint in U_i and the other endpoint in U_j . By $X[U_i]$ we denote the *induced subspace* of X , i.e. the closed graph-like subspace with vertex set U_i and edge set $E(U_i, U_i)$. Finally, a clopen subset $U \subseteq V(X)$ is called a *region* if the induced subspace $X[U]$ is connected.

LEMMA 13.2.1. *In a compact graph-like space, all cuts are finite.*

PROOF. Suppose there is an infinite cut $F = \{f_n : n \in \mathbb{N}\}$ induced by a separation (A, B) of a graph-like space X . Then A and B are disjoint closed subsets of X , so by normality there are disjoint open subsets $U \supseteq A$ and $V \supseteq B$. Since edges are connected, there exist $x_n \in f_n \setminus (U \cup V)$ for all n . It follows that $\{x_n : n \in \mathbb{N}\}$ is an infinite closed discrete subset, contradicting compactness. \square

LEMMA 13.2.2. *Let X be a compact graph-like space. For every vertex v of X and any open neighborhood U of v , there is a clopen $C \subseteq V(X)$ such that $v \in C$ and $X[C] \subseteq U$. Moreover, if X is connected, then C can be chosen to be a region.*

PROOF. Since $V(X)$ is totally disconnected we have

$$\{v\} = \bigcap \{X[A] : (A, B) \text{ a separation of } X, v \in A\}.$$

Now $\bigcap X[A] \subseteq U$ and compactness implies that there is a finite subcollection A_1, \dots, A_n such that for the clopen set $B = A_1 \cap \dots \cap A_n$ we have

$$v \in X[B] = X[A_1] \cap \dots \cap X[A_n] \subseteq U.$$

For the moreover part, since $E(B, V \setminus B)$ is finite by Lemma 13.2.1, it follows from connectedness of X that $X[B]$ consists of finitely many connected components, say $X[B] = X[C_1] \oplus \dots \oplus X[C_k]$, one of which contains the vertex v . This is our desired region C . \square

DEFINITION 13.2.3. Let X be a graph-like space and \mathcal{U} be a multi-cut of X . The *multi-graph associated with \mathcal{U}* is the quotient space $G(\mathcal{U}) = X/\{X[U]: U \in \mathcal{U}\}$. The map $\pi_{\mathcal{U}}: X \rightarrow G(\mathcal{U})$ denotes the corresponding quotient map.

We remark that $G(\mathcal{U})$ is indeed a finite multi-graph. The identified $X[U]$ form a finite collection of vertices, which are connected by finitely many edges (see Lemma 13.2.1). The degree of $\pi_{\mathcal{U}}(U_i)$ in $G(\mathcal{U})$ is given by $|E(U_i, V \setminus U_i)| < \infty$. Our next proposition gathers properties of graphs associated with multi-cuts.

PROPOSITION 13.2.4. *Let X be a graph-like compact space. Then*

- (1) *X is connected if and only if $G(\mathcal{U})$ is connected for all multi-cuts \mathcal{U} of X .*
- (2) *All cuts of X are even if and only if all vertices in $G(\mathcal{U})$ have even degrees for all multi-cuts \mathcal{U} of X .*

PROOF. (1) If X is connected, then connectedness of $G(\mathcal{U})$ follows from the fact that it is the continuous image of X . Conversely, a disconnection of X gives rise to a $G(\mathcal{U})$ which is the empty graph on two vertices.

(2) If every cut of X is even, then the above degree considerations show that every vertex in $G(\mathcal{U})$ has even degree. And conversely, any odd cut of X gives rise to a graph $G(\mathcal{U})$ on two vertices of odd degree. \square

Recall that a *standard subspace* Y of a graph-like space X is a closed subspace that contains all edges it intersects (i.e. whenever $e \cap Y \neq \emptyset$ then $e \subseteq Y$). Standard subspaces of graph-like spaces are again graph-like. Write $E(Y)$ for the collection of edges of Y .

LEMMA 13.2.5. *Let X be a graph-like space and $C \subseteq X$ a copy of a topological circle. Then C is a standard subspace.*

PROOF. Assume, by contradiction, that there exists $e \in E(X)$ such that $e \cap C \neq \emptyset$ and that $e \not\subseteq C$. Let $y \in e \setminus C$. Then there exist $x_0 \in C$ with the properties that the arc $[x_0, y]$ is a subset of e and $[x_0, y] \cap C = \{x_0\}$. Observe that $x_0 \notin V$. Let U be an open set containing x_0 such that $U \cap V = \emptyset$. Let α be the component in $[x_0, y]$ of x_0 contained in U and β be the component in C of x_0 contained in U . Then $\alpha \cup \beta$ contains a triod and $\alpha \cup \beta \subseteq X \setminus V$ which is a contradiction to the fact that $X \setminus V \cong E \times (0, 1)$ contains no triods. \square

LEMMA 13.2.6. *Let X be a graph-like space and $C \subseteq X$ a copy of a topological circle. Then $E(C) \cap F$ is finite and even for all cuts $F = E(A, B)$ of X .*

PROOF. By Lemma 13.2.5 we may assume $X = C$. Let $F = E(A, B)$. That F is finite is immediate from Lemma 13.2.1, so we only need to prove that $|F|$ is even.

Let $C[A]$ (resp. $C[B]$) be the standard subspace containing A (resp. B) and all edges with both endpoints in A (resp. B). Observe that

- (a) $C = C[A] \cup F \cup C[B]$, and
- (b) $C[A]$ and $C[B]$ have finitely many components.

Let A_1, \dots, A_r be the components of $C[A]$ and B_1, \dots, B_s be the components of $C[B]$. These components induce a multi-cut, $\mathcal{U} = \{U_{A_1}, \dots, U_{A_r}, U_{B_1}, \dots, U_{B_s}\}$, of the vertices of C where U_{A_i} (resp. U_{B_i}) consists of all vertices contained in A_i (resp. B_i). Then $G(\mathcal{U})$, the multi-graph associated with \mathcal{U} , is a cycle whose edges are the elements of F and whose vertices are the equivalence classes containing the sets U_{A_1}, \dots, U_{B_s} . Observe that the sets $\mathcal{A} = \{U_{A_1}, \dots, U_{A_r}\}$ and $\mathcal{B} = \{U_{B_1}, \dots, U_{B_s}\}$ give a 2-coloring of the vertices of $G(\mathcal{U})$. Hence $G(\mathcal{U})$ has an even number of edges, i.e. $|F|$ is even. \square

13.2.2. Characterizations and Representations. In this section we prove Theorem (A). The equivalence of (i) and (iii) is given by Proposition 13.2.10, the equivalence of (i) and (ii) is Theorem 13.2.11, while the equivalence of (i), (iv) and (iv)' follows from Theorems 13.2.13 and 13.2.14. Compact graph-like spaces were explicitly defined to include standard subspaces of the Freudenthal compactification of locally finite graphs. Theorem 13.2.15 provides the converse, establishing equivalence of (i) and (v).

Recall that a continuum X is *regular* if it has a basis of open sets, each with finite boundary, and it is called *completely regular* if each non-degenerate subcontinuum of X has non-empty interior in X , see [41, Page 1176]. A continuum is *hereditarily locally connected (hlc)* if every subcontinuum is locally connected, and *finitely Souslian* if each sequence of pairwise disjoint subcontinua forms a null-sequence, i.e. the diameters of the subcontinua converge to zero. It is known that for continua

$$(\ddagger) \text{ completely regular} \Rightarrow \text{regular} \Rightarrow \text{finitely Souslian} \Rightarrow \text{hlc} \Rightarrow \text{arc-connected}.$$

For the first three implications, see [104, Proposition 1.1].

LEMMA 13.2.7. *Every compact graph-like space is regular.*

PROOF. Let X be a compact graph-like space, $p \in X$, and U be an open of X set such that $p \in U$. We will show that there is an open set O with finite boundary such that $p \in O \subseteq U$.

The case when p is in the interior of an edge follows from the fact that the set of edges is discrete. So we may assume $p \in V$. For this case let $X[B]$ as in the proof of Lemma 13.2.2, then $p \in X[B] \subseteq U$. Now for each $e \in E(B, V \setminus B)$, let (v, x_e) be a subarc of e such that $(v, x_e) \subseteq U$ and such that $v \in B$. Since cuts are finite, then there are only finitely many of these arcs. The desired open set O is then $X[B] \cup \{(v, x_e) : e \in E(B, V \setminus B)\}$ as its boundary is the set $\{x_e : e \in E(B, V \setminus B)\}$. \square

COROLLARY 13.2.8. *Every graph-like continuum is finitely Souslian, hereditarily locally connected and arc-connected.*

PROOF. By Lemma 13.2.7 and (\ddagger) , this is a consequence of regular.

For a direct proof that graph-like continua are finitely Souslian, suppose for a contradiction that $\{A_i : i \in \mathbb{N}\}$ forms a sequence of disjoint subcontinua of X with non-vanishing diameter. It follows from the sequential compactness of the hyperspace of subcontinua, [121, Corollary 4.18], that there is a subsequence A_{i_j} such that $A = \lim_{j \rightarrow \infty} A_{i_j} = \overline{\bigcup_j A_{i_j}} \setminus \bigcup_j A_{i_j}$ is a non-trivial subcontinuum of X . But since edges are open, we also have that $A \subseteq V(X)$, so is totally disconnected, a contradiction.

For a direct proof that graph-like continua are hlc, see Lemma 13.2.2. \square

In particular, noting that a compact graph-like space has at most countably many edges (as they form a collection of disjoint open subsets), it follows that the edges of X form a null-sequence, i.e. $\lim_{n \rightarrow \infty} \text{diam}(e_n) = 0$. Here, for a subset A of a metric space, we denote by $\text{diam}(A)$ the diameter of A .

In the next theorem we use the following notation. For a subspace $A \subseteq X$ we denote by $\text{Bd}(A)$ its boundary. A subarc $A \subseteq X$ is called *free* if A removed its endpoints is open in X .

THEOREM 13.2.9 ([104, Theorem 1.3]). *A continuum X is completely regular if and only if there exists a 0-dimensional compact subset F of X and a finite or countable null sequence of free arcs A_1, A_2, \dots in X such that*

$$X = F \cup \left(\bigcup \{A_n : n \geq 1\} \right) \text{ and } A_j \cap F = \text{Bd}(A_j)$$

for each $j \geq 1$

Observe that Theorem 13.2.9 implies that every completely regular continuum is a graph-like space. Conversely, if X is a graph-like continuum, then the set of vertices V is zero-dimensional. Also by Corollary 13.2.8, $E(X)$ forms a null sequence. By Theorem 13.2.9, X is a completely regular continuum.

PROPOSITION 13.2.10. *Let X be a continuum. Then X is completely regular if and only if X is a graph-like space.*

Recall that a graph can be characterized in terms of order: a continuum is a graph if and only if every point has finite order, and all but finitely many points have order 2, [121, Theorem 9.10 & 9.13].

THEOREM 13.2.11 (Graph-like Characterization). *A continuum is graph-like if and only if it is regular and has a closed zero-dimensional subset V such that all points outside of V have order 2.*

PROOF. Sufficiency follows from the definition of graph-like and Lemma 13.2.7.

For the necessity, first observe that regular implies local connectedness. Let $V \subseteq X$ be a closed zero-dimensional collection of points in X such that all points outside of V have order 2. By local connectedness, all components of $X \setminus V$ are open subsets of X .

In particular, we have at most countably many components, and each component is non-trivial, non-compact, and consists exclusively of points of order 2. So each component is homeomorphic to an open interval. So all that remains to show for graph-like is that the closure of each edge is compact, which is automatic. \square

COROLLARY 13.2.12 (Canonical Representation of Graph-like Spaces). *Let $X \not\cong S^1$ be a graph-like continuum. Then there is a unique minimal set $V \subseteq X$ which witnesses that X is a graph-like space. We call (X, V, E) the standard representation of X .*

PROOF. Let $\{V_s : s \in S\}$ be the collection of all subsets of X which witness that X is graph-like. We claim that $V = \bigcap_{s \in S} V_s$ is also a vertex set.

Clearly, V is closed and zero-dimensional. Further, if $x \notin V$, then $x \notin V_s$ for some $s \in S$, so x has order 2. So either V is empty, in which case $X \cong S^1$; or V is non-empty, in which case every component of $X \setminus V$ is non-compact, open, and consists of points of order 2, so is homeomorphic to an open interval. \square

Our next theorem has been proved, for completely regular continua, by Nikiel, [130, 3.8]. We reprove his theorem here (and extend it to graph-like compacta), phrased for convenience in the language of graph-like continua.

THEOREM 13.2.13 (Inverse Limit Representation). *Every graph-like compact space X can be represented as an inverse limit of multi-graphs G_n ($n \in \mathbb{N}$) with onto simplicial bonding maps that have non-trivial fibres at vertices only, such that*

- (1) X connected if and only if G_n is connected for all n , and
- (2) all cuts $E(A, B)$ in X are even \Leftrightarrow all vertices in G_n are even for all n .

Moreover, if X is connected, then the bonding maps can additionally be chosen monotone.

PROOF. Let X be a graph-like compactum with vertex set V and edge set E . Without loss of generality, X contains no loops, as otherwise we can subdivide each edge once (this does not change the homeomorphism type of X , and the new edge set is still a dense open subset, so the new vertex set is a compact, zero-dimensional subspace as required).

Since V is a compact, zero-dimensional metrizable space, we can find, as in Lemma 13.2.2, a sequence of multi-cuts $\{\mathcal{U}_n : n \in \mathbb{N}\}$ such that

- (a) \mathcal{U}_{n+1} is a refinement of \mathcal{U}_n ,
- (b) $\bigcup_{n \in \mathbb{N}} \mathcal{U}_n$ forms a basis for $V(X)$, and

Writing $\mathcal{U}_n = \{U_1^n, U_2^n, \dots, U_{i(n)}^n\}$ we observe that every $v \in V$ has a unique description in terms of $\{v\} = \bigcap_{n \in \mathbb{N}} U_{l(v)}^n$ and that conversely, for every nested sequence of cut elements, there is precisely one vertex in $\bigcap_{n \in \mathbb{N}} U_{l_n}^n$ by compactness and (b).

The inverse system: Let $\{\mathcal{U}_n : n \in \mathbb{N}\}$ be as above. To simplify notation, let q_n stand for $\pi_{\mathcal{U}_n}$. For each $n \in \mathbb{N}$ let $f_n : G(\mathcal{U}_{n+1}) \rightarrow G(\mathcal{U}_n)$ be defined as

$$f_n(x) = q_n(q_{n+1}^{-1}(x)) \text{ for all } x \in G(\mathcal{U}_{n+1}).$$

Observe that if $U_i^{n+1}, U_j^{n+1} \subseteq U_s^n$,

- (i) then $f(q_{n+1}(U_i^{n+1})) = f(q_{n+1}(U_j^{n+1})) = q_n(U_s^n)$;
- (ii) and if $e \in E(U_i^{n+1}, U_j^{n+1})$; in particular $e \in E(U_s^n, U_s^n)$, then $f_n(e) = q_n(U_s^n)$.

In particular, each f_n is an onto simplicial map with non-trivial fibres only at vertices of $G(\mathcal{U}_n)$. Then $\{G(\mathcal{U}_n), f_n\}_{n \in \mathbb{N}}$ is an inverse sequence of multi-graphs. Hence, its inverse limit is compact and nonempty. We will show that there is a continuous bijection

$$f: X \rightarrow \varprojlim_{n \in \mathbb{N}} G(\mathcal{U}_n).$$

For $x \in X$, we define $f(x) = (q_1(x), q_2(x), \dots)$. By the product topology, this is a continuous map into the product $\prod_n G(\mathcal{U}_n)$, as all coordinate maps q_n are continuous. Moreover, it is straightforward from the definition of f_n to check that $f(x) \in \varprojlim G(\mathcal{U}_n)$. That the map f is surjective follows from the fact that each q_n is continuous and X is compact (see [121, 2.22]). Finally, f is injective because of the neighborhood bases requirement (b) on \mathcal{U}_n . Since X is compact and $\varprojlim G(\mathcal{U}_n)$ is Hausdorff, it follows that f is a homeomorphism as desired, and properties (1) and (2) now follow from Proposition 13.2.4.

For the moreover part, simply require that besides (a) and (b), our sequence of multi-cuts $\{\mathcal{U}_n: n \in \mathbb{N}\}$ also satisfies

- (c) every multi-cut \mathcal{U}_n partitions $V(X)$ into regions.

That this is possible follows from Lemma 13.2.2; and clearly, property (c) implies that each f_n as defined above will be a monotone map. \square

In fact, a converse of the above theorem holds. This has been mentioned, for completely regular continua, by Nikiel, [130, 3.10(i)], though without proof. We provide the proof in the language of graph-like continua.

THEOREM 13.2.14. *Let X be a countable inverse limit of connected multi-graphs X_n with finite vertex sets $V(X_n)$ and onto monotone bonding maps $f_n: X_{n+1} \rightarrow X_n$ satisfying: (+) $f_n(V(X_{n+1})) \subseteq V(X_n)$. Then X is a graph-like continuum.*

PROOF. By Theorem 13.2.11, every regular continuum with the property that all but a closed zero-dimensional subset of points are of order 2 is a graph-like continuum.

That X is regular follows from [130, 3.6]. For sake of completeness, we provide the argument. Let $\pi_n: X \rightarrow X_n$ denote the projection maps, and for $m \geq n$ write $f_{m,n} = f_n \circ f_{n+1} \circ \dots \circ f_{m-1} \circ f_m: X_{m+1} \rightarrow X_n$.

Claim: For every $n \in \mathbb{N}$, the set $P_n = \{y \in X_n: |\pi_n^{-1}(y)| > 1\}$ is countable.

This holds, because for every $m \geq n$, the set $Q_m = \{y \in X_n: |f_{m,n}^{-1}(y)| > 1\}$ is countable: By assumption, all bonding maps f_m are monotone, and hence so is $f_{m,n}$. Thus, the collection of non-degenerate $f_{m,n}^{-1}(y)$ from a disjoint collections of subcontinua of X_m ,

all with non-empty interior. It follows that $P_n = \bigcup_{m \geq n} Q_m$ is countable, completing the proof of the claim.


To conclude that X is regular, let $x \in X$ and let U be an open neighborhood of $x \in X$. Then there is $k \in \mathbb{N}$ and an open subset $W \subseteq X_k$ with $x \in \pi_k^{-1}(W) \subseteq U$. Note that since X_k is a graph, and P_k is countable by the claim, we may choose W with finite boundary such that $\text{Bd}(W) \cap P_k = \emptyset$. It follows that $\pi_k^{-1}(W)$ has finite boundary, as well.

Our candidate for the vertex set of X is $V(X) = \bigcap_{n \in \mathbb{N}} \pi_n^{-1}(V(X_n))$. By (+), the family $\{(V(X_n), f_n) : n \in \mathbb{N}\}$ gives a well-defined inverse limit, which is identical with our vertex set, i.e. $V(X) = \lim_{\leftarrow} \{V(X_n), f_n\}$. Since all $V(X_n)$ are finite discrete sets, it follows that $V(X)$ is a compact zero-dimensional metric space, as desired.

To see that elements $y \in X \setminus V(X)$ have order 2, note that $y \notin V(X)$ means there is an index $N \in \mathbb{N}$ such that $\pi_n(y)$ is an interior point of an edge of X_n for all $n \geq N$. Consider an open neighborhood U with $y \in U \subseteq X$. As before, there is an index $k > N$ and an open subset $W \subseteq X_k$ with $y \in \pi_k^{-1}(W) \subseteq U$. Since $\pi_k(y) \in X_k$ is an interior point of an edge, and P_k is countable by the claim, we may assume that W has 2-point boundary with $\text{Bd}(W) \cap P_k = \emptyset$. It follows that $\pi_k^{-1}(W)$ has a 2-point boundary, as well. \square

In fact, the class of continua, which can be represented as countable monotone inverse limits of finite connected multi-graphs are precisely the so-called *totally regular continua*, [42] – for each countable $P \subseteq X$, there is a basis \mathcal{B} of open sets for X so that for each $B \in \mathcal{B}$, $P \cap \text{Bd}(B) = \emptyset$ and B has finite boundary. These continua have also been studied under the name *continua of finite degree*. The class of totally regular continua is strictly larger than the class of completely regular continua. In particular, the condition in Theorem 13.2.14 on f_n having nontrivial fibers only at vertices cannot be omitted. For example, the universal dendrite D_n of order n can be obtained as the inverse limit of finite connected graphs, see [47, Section 3], and D_n has a dense set of points of order $\neq 2$.

In [49] the graph-like continuum depicted on the left side of the diagram below served to show that graph-like continua form a wider class than Freudenthal compactifications of locally finite graphs. Note that the two black nodes simultaneously act as ends for the blue double ladder, and as vertices for the red edge.



However, after subdividing the red edge appropriately – turning it into a double ray – we see from the right side that it can be realized as a standard subspace of the Freudenthal compactification of the triple ladder. We now show that every graph-like continuum has the same property.

THEOREM 13.2.15. *Every graph-like continuum can be embedded as a standard subspace of a Freudenthal compactification of a locally finite graph.*

We remark that Theorem 13.2.15 can be rephrased as saying that *every graph-like continuum has a subdivision, turning each edge into a double ray, which is a standard subspace of a Freudenthal compactification of a locally finite graph.*

In the proof of Theorem 13.2.15, we use the following notation. Let G be a finite, connected graph with vertex set V , and let $L(G)$ be its (connected) line graph, both considered as 1-complexes. For every edge $e \subseteq G$, let $m_e \in e$ be the mid-point of that edge. Then by G^{\otimes} we denote the graph

$$G^{\otimes} = (G \oplus L(G)) / \sim, \quad \text{where } m_e \sim e \text{ for } m_e \in G \text{ and } e \in V(L).$$

Geometrically, we subdivide each edge of G in its mid-point, and connect two new such vertices if and only if their underlying edges share a common vertex.

PROOF OF THEOREM 13.2.15. Let X be a graph-like continuum. Represent X as a monotone inverse limit of finite multi-graphs G_n with onto, monotone simplicial bonding maps $f_n: G_{n+1} \rightarrow G_n$ having non-trivial fibres at vertices only.

Recall first that the Freudenthal compactification of a locally finite graph can be realized as an inverse limit: Let L be a locally finite graph with vertex set $V(L) = \{v_k: k \in \mathbb{N}\}$ say. Let k_n be an increasing sequence of integers, and consider for each n the induced subgraph $L_n = L[v_0, \dots, v_{k_n}]$. Let L^n denote the multi-graph quotient of L where we contract every connected component of the induced subgraph $L[V(L) \setminus V(L_n)]$, deleting all arising loops. Since L was locally finite, it is easy to check that L^n is a finite multi-graph. Then $\{L^n: n \in \mathbb{N}\}$ forms an inverse system under that natural projection maps $g_n: L^{n+1} \rightarrow L^n$, such that the resulting inverse limit $\lim_{\leftarrow} L^n \cong \gamma L$ is the Freudenthal compactification of L ; moreover, this holds independently of the sequence k_n .

Now our proof strategy is as follows. We plan to find a locally finite graph L as above such that there are subgraphs $T_n \subseteq L^n$ such that

- (i) $\hat{g}_n = g_n \upharpoonright V(T_{n+1}) \rightarrow V(T_n)$ restricts to a surjection (so that the T_n form a subsystem of the inverse limit with bonding maps \hat{g}_n), and
- (ii) for each $n \in \mathbb{N}$, the graph T_n witnesses that G_n is a topological minor of L^n , meaning there are homeomorphisms $h_n: G_n \rightarrow T_n$ of the underlying 1-complexes which map $V(G_n) \hookrightarrow V(T_n)$, and map distinct edges vw and xy of G_n to independent $h_n(v)h_n(w)$ - and $h_n(x)h_n(y)$ -paths in T_n , and
- (iii) we have $\hat{g}_n \circ h_{n+1} = h_n \circ f_n$ for all n , i.e. the following diagram commutes:

$$\begin{array}{ccc}
 T_n & \xleftarrow{\hat{g}_n} & T_{n+1} \\
 h_n \uparrow & & \uparrow h_{n+1} \\
 G_n & \xleftarrow{f_n} & G_{n+1}
 \end{array}$$

Under these assumptions, it follows that $X = \lim_{\leftarrow} G_n$ is homeomorphic to the inverse limit $\lim_{\leftarrow} T_n$, which in turn, as it was constructed as a subsystem, embeds into the inverse limit $\lim_{\leftarrow} L^n = \gamma L$, which equals the Freudenthal compactification of L by the foregoing discussion. Thus, it remains to find a locally finite graph L subject to requirements (i)–(iii).

We will build this locally finite graph L by geometric considerations as a direct limit of finite connected graphs F_n , so that $F_n = L[V(F_n)] = L_n$. More precisely, we will define finite connected 1-complexes F_n such that

- (1) $F_0 \hookrightarrow F_1 \hookrightarrow F_2 \hookrightarrow \dots$ forms a direct limit such that for all $n > 0$, no vertex of $F_{n+1} \setminus F_n$ is incident with a vertex of F_{n-1} , and
- (2) F_n is embedded together with G_n in some ambient 1-complex $H_n = F_n \cup G_n$ such that
 - (a) no vertex of G_n lies in F_n ,
 - (b) every vertex of F_n lies on an edge of G_n ,
 - (c) every open edge of F_n is either disjoint from G_n , or completely contained in an edge of G_n , and
 - (d) every edge of G_n intersects with F_n in a non-trivial path $P \subseteq F_n$ such that the end-vertices of P are vertices of $V(F_n) \setminus V(F_{n-1})$.

To begin, put $H_0 = G_0^*$, and let F_0 denote the subgraph $L(G_0) \subseteq G_0^*$. Then (2) is satisfied since vertices of F_0 are mid-points of edges of G_0 , and every open edge of F_0 is disjoint from G_0 ; and (1) is trivially true.




FIGURE 13.1. Depicts the first bonding map f_0 between graphs G_1 and G_0 in black, where $f(\{v_1, v_2, v_3\}) = \{v\}$. Further, the figure on the left shows $F_0 \subseteq G_0^*$ in red, and on the right $F_1 \subseteq H_1$ as the union of \tilde{F}_0 in red, $\bigcup_v L_v$ in blue, and edges induced by \tilde{F}_0 and $\bigcup_v L_v$ in green.

Now inductively, suppose we have already defined $H_n = G_n \cup F_n$ for some $n \in \mathbb{N}$ according to (1) and (2). First, consider the natural pull-back $\tilde{F}_n \subseteq G_{n+1}$ of F_n under f_n . More precisely, by (2), the preimage $f_n^{-1}(F_n) \subseteq G_{n+1}$ is isomorphic to a subgraph of F_n . Let \tilde{F}_n be an isomorphic copy of F_n on the vertex set $f_n^{-1}(V(F_n))$ obtained by adding all edges missing from $f_n^{-1}(F_n)$ so that they are disjoint from G_{n+1} .

For every component C_v of the topological subspace $H_n \setminus F_n$ (which by (2)(a) and (b) will be a vertex v of G_n incident with finitely many half-open edges), consider the subcontinuum $K_v = \overline{f_n^{-1}(C_v)} \subseteq G_{n+1}$. Then K_v is a finite connected graph. For each v , consider K_v^\circledast , and $L_v = L(K_v) \subseteq K_v^\circledast$, and define F_{n+1} to be the induced subgraph $F_{n+1} = \tilde{F}_n \cup \bigcup \{L_v : v \in V(G_n)\}$.

Claim 1: F_{n+1} is a connected graph.

By induction on n . If F_n is connected, then so is its isomorphic copy \tilde{F}_n . As line graphs of connected graphs, every L_v is connected. Since by construction, every L_v is connected via an induced edge to \tilde{F}_n , it follows that F_{n+1} is connected.

Claim 2: Property (2) holds for F_{n+1} and G_{n+1} .

(a) No vertex v of G_{n+1} lies on \tilde{F}_n , as otherwise $f_n(v)$ would be a vertex of G_n on F_n . Also, since all L_v are partial line graphs of G_{n+1} , we see that (a) holds at step $n + 1$.

(b) Similar.

(c) By construction, this holds for all edges of \tilde{F}_n . Further, all edges of L_v are disjoint from G_{n+1} , and all edges of F_{n+1} induced \tilde{F}_n and L_v are completely contained in one edge of G_{n+1} by definition.

(d) Let $e = vw$ be an edge of G_{n+1} . If $e \notin E(G_n)$ then $F_{n+1} \cap e = L_v \cap e$ is a trivial path consisting of one new vertex. Otherwise, if $e \in E(G_n)$, then by construction and induction assumption, \tilde{F}_n intersects e in a non-trivial path $P \subseteq \tilde{F}_n$ such that the end-vertices of P have been added only at the previous step. But now, we see that $F_{n+1} \cap G$ is a path P' which is one edge longer on either side than P , because we added two edges induced by L_v and L_w . In particular, the end vertices of P' are vertices of L_v and L_w , and so have only been added at this step.

Claim 3: Property (1) holds.

Since $F_n \cong \tilde{F}_n \subseteq F_{n+1}$ it is clear how to choose the embedding $F_n \hookrightarrow F_{n+1}$. The second part of the claim now follows from (2)(d) as follows: Every vertex of $F_{n+1} \setminus F_n$ is a vertex of some L_v . By construction, any such vertex is connected at most to one of the end vertices on some path P , which is, by (2)(d), a vertex of $F_n \setminus F_{n-1}$.

This completes the recursive construction. As indicated above, the graphs $F_0 \hookrightarrow F_1 \hookrightarrow F_2 \hookrightarrow \dots$ give rise to a direct limit, which we call L . Let $V(L) = \{v_k : k \in \mathbb{N}\}$ be an enumeration of the vertices of L , first listing all vertices of F_0 , then all (remaining) vertices of F_1 etc. It is clear that there is an increasing sequence of integers k_n such that $L_n = L[\{v_0, \dots, v_{k_n}\}] = F_n$.

Claim 4: L is a locally finite connected graph.

To see that L is locally finite, note that any vertex $v \in L$ is contained in some F_n for some n , and then (1) implies that $\deg_L(v) = \deg_{F_{n+1}}(v) < \infty$. And since every L_n is connected, so is L .

Claim 5: *There are isomorphisms $\varphi_n: H_n \rightarrow L^n$. It suffices to show that $L^n \cong F_{n+1}/\{L_v: v \in V(G_n)\}$. Indeed, (1) implies that the connected components of $L \setminus F_n$ correspond bijectively to the connected components of $F_{n+1} \setminus F_n$, which are, by construction, precisely the L_v indexed by the different $v \in V(G_n)$. In particular, φ_n is a bijection between $V(G_n)$ and the dummy vertices of L^n that commutes with the respective bonding maps, i.e.*

$$(\dagger) \quad g_{n-1} \circ \varphi_n(v) \upharpoonright V(G_n) = \varphi_{n-1} \circ f_{n-1}(v) \upharpoonright V(G_n).$$

Claim 6: *For $T_n = \varphi_n(G_n) \subseteq L^n$ the subgraph of L^n which is the image of 1-complex $G_n \subseteq H_n$ subdivided by the vertices of L_n , satisfies (i)–(iii). Everything is essentially set up by construction; (iii) follows by (\dagger) with $h_n = \varphi_n \upharpoonright G_n$. \square*

Note that our embedding of X into γL has the property that every vertex of the graph-like continuum X is represented by a compactification point (an end) of γL . By exercising some extra care in the above construction, one could arrange for isolated vertices of $V(X)$ to be mapped to vertices of L .

Remark. Theorem 13.2.15 has the following notable consequence. Diestel asked in [51] whether every connected subspace of the Freudenthal compactification of a locally finite graph is automatically arc-connected. In 2007, Georgakopoulos gave a negative answer, [79]. However, the analogous problem for arbitrary continua is a well-studied problem. Indeed, a continuum is said to be *in class A* if every connected subset is arc-wise connected. Continua in class A have been characterized by Tymchatyn in 1976, [157]. Even earlier, in 1933, Whyburn gave an example of a completely regular continuum which is not in class A, [165, Example 4]. Applying Theorem 13.2.15, Whyburn's example shows at once that Freudenthal compactifications of locally finite graphs are not necessarily in class A.

13.3. Eulerian graph-like continua

13.3.1. Characterizing Eulerian Graph-Like Continua. We now prove the equivalence of (i), (ii), (iii) and (iv) of Theorem (B) in the case of *closed* paths, and then deduce the same equivalences in Theorem (B) for *open* paths. To start note that (iv) \Rightarrow (iii) and (i) \Rightarrow (iii) of Theorem (B) follow from Lemma 13.2.6. The next lemma takes care of (iii) \Rightarrow (iv).

LEMMA 13.3.1. *A graph-like continuum such that every topological cut is even can be decomposed into edge-disjoint topological cycles.*

PROOF. We adapt the proof from [123] as follows. Let G be a graph-like continuum, and $E(G) = \{e_0, e_1, \dots\}$ an enumeration of its edges. Note that $G - e_0$ is not disconnected:

If $G - e_0 = A \oplus B$ then (A, B) would be a separation in G with $E(A, B) = \{e\}$, so odd, a contradiction. Since G is arc-connected by Corollary 13.2.8, there is an arc in $G - e_0$ connecting x and y . Together with e_0 that gives a topological circle C_0 .

Now let $e_i = x_i y_i$ be the first edge not on C_0 . We claim that there is a path connecting x_i to y_i in $G \setminus (E(C_0) \cup \{e_i\})$. Otherwise, there is a cut (A, B) of $G' = G \setminus E(C_0)$ such that $E_{G'}(A, B) = \{e_i\}$. But then the same cut viewed in G would be odd by Lemma 13.2.6. A contradiction.

It is clear that we can continue in this fashion until all edges are covered. \square

To establish the equivalence of clauses (i), (iii) and (iv) of Theorem (B), it remains to show (iii) implies (i), which is established by the next result.

PROPOSITION 13.3.2. *Let X be a graph-like continuum. If all topological cuts of X have even size then X has an Eulerian loop.*

PROOF. By Theorem 13.2.13 (2), X can be written as an inverse limit of graphs G_n , which are all closed Eulerian. Let f_n denote the bonding map $f_n: G_{n+1} \rightarrow G_n$.

For each n , let \mathcal{E}_n be the collection of all Euler cycles of G_n . Since G_n is finite, so is \mathcal{E}_n . For each $n \in \mathbb{N}$, let $\hat{f}_n: \mathcal{E}_{n+1} \rightarrow \mathcal{E}_n$ be the map induced by f_n . That is, if $E = (v_0 e_0 v_1 e_1 v_2 e_2 \cdots v_k e_k v_0)$, then

$$\hat{f}_n(E) = (f_n(v_0) f_n(e_0) f_n(v_1) f_n(e_1) \cdots f_n(e_k) f_n(v_0)).$$

Observe that from the proof of Theorem 13.2.13 some of the edges in E get contracted to a vertex. So $\hat{f}_n(E)$ is an Eulerian circuit in G_n . Now, $\{\mathcal{E}_n, \hat{f}_n\}_{n \in \mathbb{N}}$ forms an inverse system, and since each \mathcal{E}_n is compact, we see $\lim_{\leftarrow} \mathcal{E}_n \neq \emptyset$.

Let $(E_n) \in \lim_{\leftarrow} \mathcal{E}_n$. For each $n \in \mathbb{N}$, fix an Eulerian loop $\varphi_n: S^1 \rightarrow G_n$ following the pattern given by E_n . Now observe that since the $(E_n)_{n \in \mathbb{N}}$ are compatible, there are monotone continuous maps $g_n: S^1 \rightarrow S^1$ ($n \in \mathbb{N}$) such that the diagram

$$\begin{array}{ccccccc} G_1 & \xleftarrow{f_1} & G_2 & \xleftarrow{f_2} & G_3 & \xleftarrow{f_3} & \cdots \\ \uparrow \varphi_1 & & \uparrow \varphi_2 & & \uparrow \varphi_3 & & \\ S^1 & \xleftarrow{g_1} & S^1 & \xleftarrow{g_2} & S^1 & \xleftarrow{g_3} & \cdots \end{array}$$

commutes. As an inverse limit of circles under monotone bonding maps, we have $\lim_{\leftarrow} S^1 \cong S^1$, [44, 4.11], and so the map

$$g: \lim_{\leftarrow} S^1 \rightarrow \lim_{\leftarrow} G_n, (x_n)_{n \in \mathbb{N}} \mapsto (\varphi_n(x_n))_{n \in \mathbb{N}}$$

is our desired Eulerian loop. \square

The proof of the equivalence of (i), (ii), (iii) and (iv) in Theorem (B), for *closed* loops, is completed by Lemma 13.3.4 showing the equivalence of (ii) and (iii). A preliminary lemma is needed.

LEMMA 13.3.3. *Let X be a graph-like continuum, (A, B) be a separation of V , and $\mathcal{U} = \{A_1, \dots, A_n\}$ be a multi-cut of A . If the cut $E(A, B)$ is odd, then $E(A_j, V \setminus A_j)$ is odd for some $1 \leq j \leq n$.*

PROOF. Consider the contraction graph induced by the multi-cut (B, A_1, \dots, A_n) .

By assumption, the vertex $\{B\}$ has odd degree. Since by the Handshaking Lemma, the number of odd-degree vertices in a finite graph is even, there must be some further vertex $\{A_j\}$ with odd degree, so $E(A_j, V \setminus A_j)$ is odd. \square

LEMMA 13.3.4. *Let X be a graph-like continuum. All topological cuts of X are even if and only if every vertex of X is even.*

PROOF. If all cuts are even, then from the definition every vertex is even. We prove the converse by contrapositive. Assume there exists a separation (A_0, B_0) of V such that $E(A_0, B_0)$ is odd. Let $\mathcal{U}_0 = \{U_{1_0}, \dots, U_{n_0}\}$ be a separation of A_0 into sets with diameter $< \frac{1}{2} \text{diam}(A_0)$. By Lemma 13.3.3 there exists $1 \leq j_0 \leq n_0$ such that $E(U_{j_0}, V \setminus U_{j_0})$ is odd. Denote U_{j_0} by A_1 and $V \setminus U_{j_0}$ by B_1 . Let $\mathcal{U}_1 = \{U_{1_1}, \dots, U_{n_1}\}$ be a separation of A_1 into sets with diameter $< \frac{1}{2} \text{diam}(A_1)$. Again by Lemma 13.3.3 there exists $1 \leq j_1 \leq n_1$ such that $E(U_{j_1}, V \setminus U_{j_1})$ is odd. Denote U_{j_1} by A_2 and $V \setminus U_{j_1}$ by B_2 . Continuing with this procedure we obtain a nested sequence of nonempty cut elements $\{A_i\}_{i \in \mathbb{N}}$. By construction $\bigcap_{i \in \mathbb{N}} A_i = \{v\} \in V$ and $E(A_i, B_i)$ is odd for every $i \in \mathbb{N}$, hence v is not even. \square

It remains to deduce the equivalence of (i), (ii), (iii) and (iv) in Theorem (B) for the case of *open* paths from that of *closed* paths. This can be achieved with a simple trick.

Suppose, to start, that item (i) for open paths of Theorem (B), holds for a graph-like continuum X . So in X there is an open Eulerian path starting at a vertex v and ending at another vertex w . Create a new graph-like continuum Z by adding one edge to X with endpoints at v and w . Then Z is a graph-like continuum with an Eulerian loop. So, by Theorem (B) applied to Z , each of (ii)-(iv) (for closed paths) of that theorem hold for Z . But now it easily follows from the definitions that each of (ii)-(iv) (for open paths) of Theorem (B) hold for X .

Now let X be a graph-like continuum for which one of items (ii)-(iv) for open paths in Theorem (B) holds. To complete the deduction we show (i) holds for open paths. Each of these items highlights two distinct vertices (the two odd vertices in (ii) and the ends of the arc in (iv)). Call them v and w . Create a new graph-like continuum Z by adding one edge to X with endpoints at v and w . Then Z is a graph-like continuum and it is easily verified from the definitions that it satisfies one of (ii)-(iv) for closed paths in Theorem (B). Hence (i) for closed paths of Theorem (B) holds, and there is a closed Eulerian path in Z . Removing the added edge yields an open Eulerian path in X .

13.3.2. Counting All Eulerian Loops and Paths. In this section we aim to count the number of distinct Eulerian loops and paths in a given graph-like continuum. To do

so we must decide what it means for two paths to be equivalent. This is a well-studied problem in combinatorial group theory, and we adopt the approach taken there. Two maps $f, g: I \rightarrow X$ are *equivalent* if $v = f(0) = g(0)$, $w = f(1) = g(1)$, v and w are vertices, and f is homotopy equivalent to g relative to v, w . As noted in the Introduction, every map $f: I \rightarrow X$ with vertices for endpoints is equivalent to a standard path.

Let X be a graph-like continuum. By Theorem 13.2.13 (2), X can be written as an inverse limit of graphs G_n , via bonding maps $f_n: G_{n+1} \rightarrow G_n$. As in Proposition 13.3.2, for each n , let \mathcal{E}_n be the collection of all Eulerian cycles in G_n , and let $\hat{f}_n: \mathcal{E}_{n+1} \rightarrow \mathcal{E}_n$ be the map induced by f_n . Recall, $(\mathcal{E}_n, \hat{f}_n)_n$ forms an inverse system, and set $\mathcal{E} = \mathcal{E}(X) = \lim_{\leftarrow} \mathcal{E}_n$. As in Proposition 13.3.2, every $(E_n)_n$ in $\mathcal{E}(X)$ gives rise to an Eulerian loop in X . It is straightforward to check that distinct members of $\mathcal{E}(X)$ gives rise to inequivalent Eulerian loops. The converse is also true, although we do not need that for our counting result. In any case we consider $\mathcal{E}(X)$ to be the *space of Eulerian loops* in X .

THEOREM 13.3.5. *A closed Eulerian graph-like continuum has either finitely many distinct Eulerian loops in which case it is a graph, or it has continuum many Eulerian loops.*

PROOF. Since every $\mathcal{E}(G_n)$ is finite discrete, the inverse limit is a compact subspace of a Cantor set. As compact subspaces of a Cantor set without isolated points have size continuum, the result follows from the next claim.

Claim: If $\mathcal{E}(X)$ contains an isolated point, then X is homeomorphic to a graph.

Fix an isolated element $(E_n)_n$ in $\mathcal{E}(X)$. Fix an Eulerian loop $f: I \rightarrow X$ of X corresponding to $(E_n)_n$ (as in Proposition 13.3.2). To witness that f is isolated, find coordinate graph G_n induced by a multi-cut $\mathcal{U} = (U_1, \dots, U_n)$ of X such that the the quotient map $q: X \rightarrow G$ acting on the set of (distinct) Euler cycles $\mathcal{E}(X) \rightarrow \mathcal{E}(G_n)$ satisfies $q^{-1}(q(f)) = \{f\}$. We claim that every $X[U_i]$ (the subspace of X induced by the vertex set U_i) is a graph. This would show that X itself is also a graph.

Without loss of generality, $f(0) \notin X[U_i]$. The map f induces a linear order on $E(U_i, V \setminus U_i)$, say (e_0, \dots, e_{2k-1}) . For all $0 \leq l < 2k$ write x_l for the end vertex of e_l in U_i (of course, the x_l need not be distinct). Let f_m be the arc between x_{2m} and x_{2m+1} induced by f . We claim that the arcs $\{f_m: 0 \leq m < k\}$ witness that $X[U_i]$ is a graph.

First of all, $X[U_i] = \bigcup_{m < k} f_m$ since $f(0) \notin X[U_i]$ implies $f_m \subseteq X[U_i]$, and f Eulerian implies that all edges in $E(U_i, U_i)$ are hit. As the edges are dense, all of $X[U_i]$ is covered.

To complete the proof, it remains to show that our arcs intersect pairwise only finitely. Indeed, we claim that $|f_m \cap f_p| \leq 1$. Otherwise, suppose that $y \neq z$ are two vertices lying in the interior of both arcs. Denote by $e_m = f_m \upharpoonright [y, z]$ and $e_p = f_p \upharpoonright [y, z]$ (or $e_m = f_m \upharpoonright [z, y]$ depending on which vertex comes first). Since f_m, f_p are edge disjoint, $e_m \neq e_p$. Then replace

- f_m by $f_m \upharpoonright [x_{2m}, y] \cup e_p \cup f_m \upharpoonright [y, x_{2m+1}]$, and

- f_p by $f_p \upharpoonright [x_{2p}, y] \cup e_m \cup f_p \upharpoonright [y, x_{2p+1}]$.

This change gives rise to an Eulerian loop f' of X distinct from f , with $q^{-1}(q(f)) \supset \{f, f'\}$, a contradiction. \square

We can deduce the analogous result for the number of open Eulerian paths by the same trick used to derive the open version of Theorem (B) from the closed version. Let X be an open Eulerian graph-like continuum, and let v, w be the two odd vertices of X . Add an edge connecting v and w , to get a closed Eulerian graph-like continuum Z . Apply the preceding result to deduce Z has either finitely many distinct Eulerian loops in which case it is a graph, or it has continuum many Eulerian loops. Removing the added edge yields either that X is a graph or has continuum many open Eulerian paths.

THEOREM 13.3.6. *An open Eulerian graph-like continuum has either finitely many distinct open Eulerian paths in which case it is a graph, or it has continuum many open Eulerian paths.*

13.4. Bruhn & Stein Parity

Let X be a graph-like continuum with vertex set V . Let v be a vertex of X . Then we say that v has *strongly even degree* (respectively, *strongly odd degree*) if there is a clopen neighborhood C of v such that for every clopen neighborhood A of v contained in C the maximal number of edge-disjoint arcs from $V \setminus A$ to v is even (respectively, odd). By Lemma 13.2.1, this is well-defined. We further say that v has *weakly even degree* (resp., *weakly odd degree*) if v does *not* have strongly odd (resp. even) degree. Equivalently, v has weakly even degree if v has a neighborhood base of clopen sets, C , so that the maximal number of edge-disjoint arcs from $V \setminus C$ to v is even. And similarly for weakly odd degree. Bruhn & Stein [39] use the same terminology for ‘strongly odd’ and ‘weakly even’ degrees, but use ‘even’ for our ‘strongly even’ and ‘odd’ for our ‘weakly odd’.

Note that isolated vertices have finite degree by Lemma 13.2.1, so for them being even and having strongly even degree coincide (and similarly for odd). In general, our notion of ‘even’ and ‘odd’ vertices implies those of Bruhn & Stein. To see this, we shall need a version of Menger’s theorem in the edge-disjoint version. That Menger-like theorems hold for graph-like continua is not surprising, and vertex-disjoint versions of Menger have been proved in [153]. We complement their results by the following theorem. Note that in finite graph theory, the edge disjoint version follows from the vertex disjoint version by applying the latter theorem to the line graph. As it is unclear, what a line-graph for graph-like spaces should be, we need a different proof.

THEOREM 13.4.1 (Menger for Graph-like Continua—Edge Disjoint Version). *Let X be a graph-like continuum. For disjoint closed sets A and B of vertices of X , the maximum number of edge-disjoint $A - B$ paths equals the minimum cut separating A from B .*

PROOF. Let k be the size of a smallest cut separating A from B . Note that since A and B are closed disjoint, it follows from compactness that such a cut exists, and hence k is finite by Lemma 13.2.1. It is clear that the maximum number of edge-disjoint $A - B$ paths is bounded by k .


Conversely, write X as an inverse limit $X = \lim_{\leftarrow} G_n$ with simplicial bonding maps $f_n: G_{n+1} \rightarrow G_n$ and simplicial projection maps $\pi_n: X \rightarrow G_n$. Without loss of generality, $\pi_n(A) \cap \pi_n(B) = \emptyset$ for all n . Let \mathcal{T}_n be the (finite) space of all k -tuples of edge-disjoint connected subgraphs of G_n that intersect both $\pi_n(A)$ and $\pi_n(B)$. By Menger's theorem for finite graphs, $\mathcal{T}_n \neq \emptyset$ for all n , so \mathcal{T}_n with natural bonding maps \hat{f}_n form their own inverse system, which is non-empty. Taking the inverse limit in each coordinate, we obtain k edge-disjoint subcontinua of X each intersecting both A and B . By Corollary 13.2.8, we can find $A - B$ paths inside each subcontinuum, which are then edge-disjoint by construction. \square

LEMMA 13.4.2. *Let X be a graph-like continuum and v an even (resp. odd) vertex in X . Then v has strongly even (resp. odd) degree.*

PROOF. Let v be an even vertex and let C be a clopen neighborhood of v such that A is a clopen neighborhood of v contained in C , then $E(V(X) \setminus A, A)$ is even. Observe that $E(V(X) \setminus A, A)$ is the minimum cut separating $V(X) \setminus A$ from v . Hence by Theorem 13.4.1, the maximum number of edge-disjoint paths from $V(X) \setminus A$ to v is equal $|E(V(X) \setminus A, A)|$ which is even. This shows that v is strongly even. \square

However, in general, strongly even degree vertices need not be even.

Example. The right hand vertex in the graph-like continuum illustrated below is neither even nor odd but has strongly even degree.



If each simple circle, \bigcirc , in the above example is replaced with a copy of \ominus , then in the resulting graph-like continuum the right hand vertex has strongly odd degree.

Our aim is to prove the following theorem, generalizing corresponding results of Bruhn & Stein [39] and Berger & Bruhn [20] for Freudenthal compactifications of graphs, and their Eulerian subspaces. Observe that this theorem can be rephrased as saying that although not every vertex of strongly even degree must be even, if *all* vertices of a graph-like continuum have strongly even degree then they are *all* even.

THEOREM 13.4.3. *A graph-like continuum is closed Eulerian if and only if all its vertices have strongly even degree.*

It is an interesting open problem, whether the same conclusion holds under the assumption that all vertices have weakly even degree. The forward implication of Theorem 13.4.3 follows from Lemma 13.4.2, Lemma 13.3.4 and Proposition 13.3.2. Theorem 13.4.12 establishes the converse. The plan for the proof of Theorem 13.4.12 is to establish the contrapositive: if X is a graph-like continuum which is not closed Eulerian then it contains a vertex without strongly even degree (i.e. of weakly odd degree). Lemma 13.4.5 shows how a certain sequence of regions leads to such a vertex. Now if X is a graph-like continuum which is not closed Eulerian, then by Theorem (B) (iii) \implies (i), there must be an odd cut in X . This provides the starting point for the sequence needed to apply Lemma 13.4.5. Theorem 13.4.6 then provides the ‘Contraction Machine’ required to create the remaining elements of the sequence.

If v and w are distinct vertices in a graph-like continuum X , and they both have strongly odd degree, then after connecting them with a new edge they will both have strongly even degree. Conversely if they both have strongly even degree, then after removing an edge connecting them, they will have strongly odd degree. Hence, as we deduced the open version of Theorem (B) from the closed version, we now derive the following characterization of open Eulerian graph-like continua.

THEOREM 13.4.4. *A graph-like continuum is open Eulerian if and only if it has exactly two strongly odd degree vertices, and the rest have strongly even degree.*

13.4.1. The odd-end lemma. For a clopen subset $U \subseteq V(X)$, consider the induced graph-like space $X[U]$. We say that a clopen subset $U \subseteq V(X)$ is a *region* if $X[U]$ is connected. By $\partial U \subseteq E(X)$ we denote the set of edges between the separation $(U, V \setminus U)$. This set is finite for regions U . Let us call a region U of X a k -region if $|\partial U| = k$, and an *even* or an *odd* region depending on whether k is even or odd.

The following lemma generalizes the corresponding lemma of Bruhn & Stein for locally finite graphs, [39, p.7f], to graph-like continua.

LEMMA 13.4.5. *Let X be a graph-like continuum, and let $E(X) = \{e_0, e_1, \dots\}$ be an enumeration of its edges. Assume there exists a sequence of regions U_0, U_1, \dots of X with the following properties:*

- (1) $|\partial U_n|$ is odd for all $n \in \mathbb{N}$,
- (2) $U_n \supset U_{n+1}$,
- (3) if D is a region of X with $U_n \supset D \supset U_{n+1}$ then $|\partial U_n| \leq |\partial D|$ for all $n \in \mathbb{N}$, and
- (4) $e_n \notin E[U_{n+1}]$.

Then X has a vertex which has weakly odd degree.

PROOF. Since $A = \bigcap_{n \in \mathbb{N}} X[U_n]$ is a nested intersection of continua by (2), it is non-empty and connected. It follows from (4) that $A \subseteq V(X)$, so $A = \{v\}$ for some vertex v ,

since $V(X)$ is totally disconnected. Furthermore, compactness implies that $\{U_n : n \in \mathbb{N}\}$ is a neighborhood base for v in $V(X)$.

Property (3) together with Theorem 13.4.1 shows that for all U_n the maximal number of edge disjoint arcs from $V \setminus U_n$ to v equals $|\partial U_n|$, so is odd by (1). Since the U_n form a neighborhood base, it follows that v has weakly odd degree. \square

13.4.2. The contraction machine. Suppose we have an odd region U_0 . We want to construct a sequence as in Lemma 13.4.5. If we recursively choose an odd region U_{n+1} of minimal $|\partial U_{n+1}|$ amongst all odd regions contained in U_n , then (1) and (2) are fine, and property (3) is satisfied at least for all odd regions D nested between U_n and U_{n+1} . Following Bruhn & Stein's idea [39], our plan for evading all even regions D with $|\partial D| < |\partial U_n|$ nested between U_n and U_{n+1} is roughly as follows: first, we contract all even regions $D \subseteq U_n$ with boundary smaller than $|\partial U_n|$ to single points. Only then do we pick our region U_{n+1} . After uncontracting, this means that every small even region lies either behind U_{n+1} , or is completely disjoint from U_{n+1} .

The next result formalizes this idea for contracting regions.

THEOREM 13.4.6 (Contraction Theorem). *Let X be a graph-like continuum such that all isolated vertices are even. Suppose further that $U \subseteq X$ is an odd region of X such that for some even $m > 0$, there is no infinite k -region with $k < m$ of X contained in U .*

Then there is a collection \mathcal{M} of disjoint regions of U such that after contracting every element of \mathcal{M} to a single point, the graph-like continuum X/\mathcal{M} , with associated (monotone) quotient map $\pi : X \rightarrow X/\mathcal{M}$, has the property that

- (i) *all isolated vertices of X/\mathcal{M} are even,*
- (ii) *there are no infinite k regions with $k \leq m$ contained in the region $\pi(U) \subseteq X/\mathcal{M}$,*
and
- (iii) *if $D \subseteq U$ is an ℓ -region of X , then there is an $\leq \ell$ -region $D' \subseteq \pi(U)$ such that $|\pi(D) \setminus D'| < \infty$.*

We divide the proof into a sequence of lemmas. For two subsets $A, B \subseteq X$, say that A splits B , or B is split by A , if $A \cap B \neq \emptyset \neq B \setminus A$.

LEMMA 13.4.7. *Let X be a graph-like continuum, and $U \subseteq X$ a region. Let R, S_1, \dots, S_n be infinite m -regions contained in U , where S_1, \dots, S_n are pairwise disjoint and $|R \setminus \bigcup_{i \leq n} S_i| = \infty$.*

If there is no infinite k -region with $k < m$ of X contained in U , then there is an m -region \tilde{R} which doesn't split any S_i such that $|R \setminus (\bigcup_{i \leq n} S_i \cup \tilde{R})| < \infty$.

For the proof we need the following lemma, which can be proven, as for graphs, by a simple double-counting argument.

LEMMA 13.4.8. *Let X be a graph-like space, and $Y, Z \subseteq V(X)$ clopen subsets. Then*

$$|\partial Y| + |\partial Z| \geq \max \{ |\partial(Y \cap Z)| + |\partial(Y \cup Z)|, |\partial(Y \setminus Z)| + |\partial(Z \setminus Y)| \}.$$

PROOF OF LEMMA 13.4.7. Without loss of generality, assume that S_1 is split by R , i.e. that $R \cap S_1 \neq \emptyset \neq S_1 \setminus R$. We claim that one of $S_1 \cup R$ or $R \setminus S_1$ is an m -region. They are clearly clopen subsets of vertices of X .

Otherwise, since $S_1 \cup R$ and $R \setminus S_1$ are infinite, we have $|\partial S_1 \cup R| > m$ and $|\partial R \setminus S_1| > m$. Thus, Lemma 13.4.8 implies that $|S_1 \setminus R| < m$ and $|S_1 \cap R| < m$, so both regions are finite, contradicting that S_1 is infinite.

Hence, one of $S_1 \cup R$ or $R \setminus S_1$ has a boundary of size m , and they can't be disconnected, as otherwise their components had to be finite. Now put R' to be either one of them, whichever was the m -region. Then R' splits strictly fewer S_i than R , but covers the same set together with the S_i . Thus, we may pick \tilde{R} to be such that it splits the fewest number of S_i , subject to the condition that $|R \setminus (\bigcup_{i \leq n} S_i \cup \tilde{R})| < \infty$. By the preceding argument, it follows that \tilde{R} does not split any of the S_i . \square

Let X be a graph-like continuum, and $U \subseteq X$ a region. Assume there is no infinite k -region with $k < m$ of G contained in U . Let $\mathcal{R} = \{R_n : n \in \mathbb{N}\}$ be an enumeration of all infinite m -regions of G contained in U . Since each R_i is faithfully represented by the finite cut $\partial R_i \subseteq E$, and E is countable, there are indeed at most countably many such regions. Below we write $\mathcal{S} \preceq \mathcal{S}'$ if \mathcal{S} is a refinement of \mathcal{S}' , i.e. for all $S \in \mathcal{S}$ there is $S' \in \mathcal{S}'$ such that $S \subseteq S'$.

LEMMA 13.4.9. *For every $n \in \mathbb{N}$ there are finite collections $\mathcal{S}_n \subseteq \mathcal{R}$ of disjoint m -regions of U such that*

- (1) for all R_j with $j \leq n$ we have $|R_j \setminus \bigcup \mathcal{S}_n| < \infty$, and (2) $\mathcal{S}_n \preceq \mathcal{S}_{n+1}$.

CONSTRUCTION. We begin with $\mathcal{S}_0 = \{R_0\}$. Suppose $\mathcal{S}_n \subseteq \mathcal{R}$ has been found satisfying the above properties. Applying Lemma 13.4.7 with R_{n+1} and the collection \mathcal{S}_n , we obtain an infinite m -region \tilde{R}_{n+1} . We claim that $\mathcal{S}_{n+1} = \{\tilde{R}_{n+1}\} \cup \{S \in \mathcal{S}_n : S \cap \tilde{R}_{n+1} = \emptyset\}$ is as desired. Indeed, by construction, \mathcal{S}_{n+1} covers R_{n+1} up to finitely many vertices; and $\bigcup \mathcal{S}_n \subseteq \bigcup \mathcal{S}_{n+1}$, so we preserved the covering properties of earlier stages. \square

We would like to contract the ‘maximal’ m -regions (with respect to inclusion) contained in $\mathcal{S} = \bigcup \mathcal{S}_n$. However, for graph-like continua, there can be infinite non-trivial chains in \mathcal{S} . Still, for any such chain $S_0 \subsetneq S_1 \subsetneq S_2 \subsetneq \dots$ of m -regions, we can contract a suitable collection of disjoint even regions such that after contraction, all S_n are finite. Our plan is to contract S_0 , and each component of $S_{n+1} \setminus S_n$, to a single point for all $n \in \mathbb{N}$. Our next lemma provides the details for the second case.

LEMMA 13.4.10. *Let X be a graph-like continuum, and $U \subseteq X$ a region. Assume there is no infinite k -region with $k < m$ of G contained in U .*

If $S \subsetneq R$ are infinite m -regions contained in U , then $X[R \setminus S]$ has at most m connected components, and every such component is an even region of X .

PROOF. Note that since $X[R]$ is path-connected, it follows that every component of $X[R \setminus S]$ has to limit onto an end vertex of some $e \in \partial S$. Thus, $X[R \setminus S]$ has at most $|\partial S| = m$ components. In particular, every component is clopen in $X[R \setminus S]$, and hence a region of X .

To see that $\partial(R \setminus S)$ is even, consider the graph induced by the multi-cut $(S, R \setminus S, V \setminus R)$. This graph has two even vertices, namely $\{S\}$ and $\{V \setminus R\}$. So by the Handshaking Lemma, also the last vertex is even, i.e. $R \setminus S$ induces an even cut. Moreover, since in the contraction graph, both $\{S\}$ and $\{V \setminus R\}$ have degree m , it follows that the third vertex has the same number of edges to $\{S\}$ and to $\{V \setminus R\}$. In other words, we have $|\partial(R \setminus S) \cap \partial R| = |\partial(R \setminus S) \cap \partial S|$.

Let C denote the vertex set of one such component. It follows that in order to establish that C is an even region, it suffices to show that

$$(31) \quad |\partial C \cap \partial R| \geq |\partial C \cap \partial S|.$$

Indeed, once we know that (31) holds for every component C , then $|\partial R| = m = |\partial S|$ gives equality in (31). To see that (31) holds, note that if $|\partial C \cap \partial R| < |\partial C \cap \partial S|$, then we see that $|\partial(S \cup C)| < m$, so this is a finite region, contradicting that S was infinite. \square

We now collapse all maximal m -regions in $\mathcal{S} = \bigcup \mathcal{S}_n$, and for every infinite proper chain in \mathcal{S} we perform the above contractions. Write \mathcal{M} for the disjoint collection of even regions we contract. Write $q_{\mathcal{M}}: V(X) \rightarrow V(X/\mathcal{M})$, which extends to a continuous (monotone) quotient map on $X \rightarrow X/\mathcal{M}$ (where we also contract all potential loops), which we also call $q_{\mathcal{M}}$. Note that since we contracted regions of a compact space, the map $q_{\mathcal{M}}: X \rightarrow X/\mathcal{M}$ is a closed, monotone map. In particular, this implies that preimages of regions are regions, see Theorem 9 of [107].

PROOF OF THEOREM 13.4.6. First, to see that X/\mathcal{M} is still a graph-like continuum, note that our countable family \mathcal{M} forms a null-sequence of clopen sets by Corollary 13.2.8. It follows from the fact that if X is separable metrizable, and $\mathcal{A} = \{A_n: n \in \mathbb{N}\}$ a null-sequence of non-empty compact subsets of X , then X/\mathcal{A} is separable metrizable, [158, A.11.6], that X/\mathcal{M} is a continuum. Further, it is graph-like, because its vertex set $V(X)/\mathcal{M}$ is totally disconnected: If there was any non-trivial connected set $C \subseteq V(X/\mathcal{M})$, then C cannot contain contracted vertices (they are isolated), so $C \subseteq V(X)$ is non-trivial connected, contradiction.

Item (i), that every isolated vertex of X/\mathcal{M} is even, follows from Lemma 13.4.10, as we only contracted even regions.

For (ii), that all m -regions of X/\mathcal{M} contained in $\pi(U)$ are finite, note that for any such m -region D of X/\mathcal{M} , the clopen vertex set $D' = \pi^{-1}(D)$ is an m -region of X . If D'

was infinite, then D' appears in our list, so is covered by some finite \mathcal{S}_n . Consider $S \in \mathcal{S}_n$. Note that S either gets contracted to a single point, or S appears in an infinite chain with at most n predecessors, in which case we contract S to at most $(m \cdot n + 1)$ -many points. It follows that D' gets contracted to finitely many points, i.e. D is finite.

For (iii), let D be an ℓ -region of X . There are at most ℓ many elements $M_1, \dots, M_\ell \in \mathcal{M}$ such that $\partial D \cap E[M_i] \neq \emptyset$. Now if $D \subseteq M_i$ for some i then it is clear that $\pi(D)$ is finite. Otherwise, choose disjoint m -regions $S_i \supset M_i$ in \mathcal{S} . We claim that either $\tilde{D} = D \cup S_1$ or $\tilde{D} = D \setminus S_1$ is an $\leq \ell$ -region. Otherwise, it follows from Lemma 13.4.8 that $|\partial(D \cap S_1)| < m$ and $|\partial(S_1 \setminus D)| < m$. So S_1 is finite, a contradiction. Continue with the other S_i . This gives us an $\leq \ell$ -region D' , which differs from D by finitely many $S \in \mathcal{S}$. \square

13.4.3. Chasing odd regions. After having established Theorem 13.4.6, the proof now proceeds essentially as in [39]. We need one more simple lemma.

LEMMA 13.4.11. *A graph-like continuum in which all isolated vertices are even does not contain finite odd regions.*

PROOF. If $A = \{v_1, \dots, v_n\} \subseteq V(X)$ is a finite region, consider the finite graph induced by the multi-cut $(V \setminus A, \{v_1\}, \dots, \{v_n\})$. Since all vertices v_i are even, it follows from the Handshaking Lemma that also $\{V \setminus A\}$ must be even. \square

THEOREM 13.4.12. *A graph-like continuum is Eulerian if all its vertices have strongly even degree.*

PROOF. Assume X is not Eulerian. To prove the contrapositive we show X contains a vertex without strongly even degree. If some isolated vertex does not have (strongly) even degree then we are done. So assume all isolated vertices of X are even. We construct a sequence of graph-like continua $X = X_0, X_1, \dots$ such that

- (a) $X_0 \xrightarrow{\pi_1} X_1 \xrightarrow{\pi_2} X_2 \xrightarrow{\pi_3} \dots$ are successive quotients with monotone open quotient maps π_n , and write $f_n = \pi_n \circ \pi_{n-1} \circ \dots \circ \pi_1$,
- (b) all X_n have the property that all isolated vertices are even,
- (c) there are regions $V_n \subseteq X_n$ such that
 - (1)' $|\partial V_n|$ is odd for all $n \in \mathbb{N}$,
 - (2)' $\pi_{n+1}(V_n) \supset V_{n+1}$,
 - (3)' any ℓ -region of X gets contracted to a $\leq \ell$ -region of X_n modulo finitely many isolated vertices; and any k -region of X_n contained in V_n with $k < |\partial V_n|$ gets contracted to finitely many vertices in X_{n+1} ,
 - (4)' $e_n \notin E[V_{n+1}]$.

Before describing the construction, let us see that that $U_n = f_n^{-1}(V_n)$ defines regions satisfying the requirements of Lemma 13.4.5, and so X has a vertex which does not have strongly even degree, as desired.

Indeed, as inverse images under monotone closed maps, they are connected, and hence regions in X . Next, it is easy to check that (1)' \Rightarrow (1), (2)' \Rightarrow (2) and (4)' \Rightarrow (4). Finally, to see (3), i.e. that U_{n+1} does not lie behind some region D of U_n with small $|\partial D|$, note that by (3)', this region D would have been contracted to finitely many points in X_{n+1} , and hence V_{n+1} would be finite, which is a contradiction by (b) and Lemma 13.4.11.

Now towards the construction of our sequence X_0, X_1, \dots with (a)–(c). First, since $X = X_0$ is not Eulerian, it has an odd cut. By choosing $V_0 = U_0$ to be an odd region of X such that $|\partial U_0|$ is minimal, we see that V_0 is as desired. Now suppose we have constructed $V_n \subseteq X_n$ according to (a)–(c). Put $m_{n+1} = |\partial V_n| - 1$.

Recursively, apply Theorem 13.4.6 with graph-like continuum $X^{(k)}$ and region $q_k \circ \dots \circ q_1(U_n)$ to obtain graph-like continua $X_n = X^{(m_n)} \succ X^{(m_n+1)} \succ \dots \succ X^{(m_{n+1})} = X_{n+1}$ with corresponding monotone quotient maps $q_k: X^{(k-1)} \rightarrow X^{(k)}$ for all even $0 < k \leq m$. Define $\pi_{n+1} = q_{m_{n+1}} \circ \dots \circ q_{m_n+1}: X_n \rightarrow X_{n+1}$.

Note that Theorem 13.4.6(i) implies (b), and (ii) and (iii) imply (c)(3)'. We now want to find an odd cut $V \subseteq \pi_{n+1}(V_n)$ such that $e_n \notin E(V)$. Towards this, note that $f_{n+1}(e_n)$ is either an isolated vertex v of X_{n+1} , or $f_{n+1}(e_n)$ is an edge with end vertices say x and y in X_{n+1} . Find a multi-cut \mathcal{V} of $\pi_{n+1}(V_n)$ into regions which either displays v as a singleton, or contains x and y in different partition elements. By Lemma 13.3.3, there is an odd region $V \in \mathcal{V}$. Since isolated vertices of X_{n+1} are even, V is not the singleton $\{v\}$. In the other situation, note that in the induced graph $G(\mathcal{V})$, the edge $f_{n+1}(e_n)$ is displayed as cross edge. In either case, we have $e_n \notin E(V)$.

Finally, amongst all odd regions of X_n contained in V pick any odd region $V_{n+1} \subseteq V$ such that $|\partial V_{n+1}|$ is minimal. This choice satisfies items (1)', (2)' and (4)'. \square

Eulerian spaces

We develop a unified theory of Eulerian spaces by combining the combinatorial theory of infinite, locally finite Eulerian graphs as introduced by Diestel and Kühn with the topological theory of Eulerian continua defined as irreducible images of the circle, as proposed by Bula, Nikiel and Tymchatyn.

First, we clarify the notion of an *Eulerian* space and establish that all competing definitions in the literature are in fact equivalent. Next, responding to an unsolved problem of Treybig and Ward from 1981, we formulate a combinatorial conjecture for characterising the Eulerian spaces, in a manner that naturally extends the characterisation for finite Eulerian graphs. Finally, we present far-reaching results in support of our conjecture which together subsume and extend all known results about the Eulerianity of infinite graphs and continua to date. In particular, we characterise all one-dimensional Eulerian spaces.

14.1. Introduction

14.1.1. The Eulerian Problem. An old, well-known quest in graph theory is to find a natural generalisation for the concept of Eulerian walks to infinite graphs. An equally old problem in topology is to find a theory that allows additional control over space-filling curves from the circle in the form of *strongly irreducible maps*. We show in this paper that these seemingly unrelated strands of research represent two sides of the same coin, and develop a general theory of Eulerian spaces that combines these combinatorial and topological research efforts into a single, unified framework.

There are two main motivations for investigating generalised Eulerian spaces. First, the combinatorial one: recall that a finite multi-graph is *Eulerian* if it admits a *combinatorial Euler tour*, a closed walk that contains every edge of the graph precisely once. Euler showed, in what is commonly considered the first theorem of graph theory and foreshadowing topology, that a finite connected multi-graph is Eulerian if and only if every vertex has even degree. See [23] for a historical account of Euler's work on this problem. An equivalent characterisation of connected Eulerian graphs, the importance of which was first realised by Nash-Williams [123], is that every edge cut is even. An *edge cut* of a graph $G = (V, E)$ is a set of edges $F \subseteq E$ crossing a bipartition $(A, V \setminus A)$ of the vertices V , in other words, the set of edges with one endvertex in A and the other outside A .

There have been numerous attempts to generalise these results to infinite graphs, see for example [67, 123, 124, 140, 139, 112]. Since combinatorial Euler tours are inherently finite objects, these attempts focused rather on constructing decompositions of such graphs into cycles or collections of two-way infinite walks, sacrificing the intuitive appeal that an Euler tour should return to its start vertex. However, for locally finite graphs, an alternative solution has recently been found by Diestel & Kühn in 2004 [56] which elegantly restores this intuitive appeal: recall that every graph G naturally turns into a topological space by interpreting each edge as an arc between its endpoints, and each combinatorial Euler tour corresponds naturally to a continuous surjection from the circle S^1 to the space G which continuously traverses through the edge-arcs in the order prescribed by the combinatorial walk, henceforth called an *edge-wise Eulerian* map. Diestel and Kühn now call an infinite, locally finite (multi-)graph *Eulerian*, if there is such an edge-wise Eulerian surjection from S^1 onto the Freudenthal compactification of the graph (formalising the idea that if the Euler tour disappears in some direction towards infinity, then it should again return from that very direction). In this setting, they were able to show that a connected multi-graph is Eulerian if and only if each of its finite edge cuts is even, thus generalising the second of the characterising conditions from the finite case to infinite, locally finite graphs.

Looking at this result, it seems natural to wonder about Eulerianity in other naturally occurring compactifications of locally finite graphs, which give a more refined meaning for a ‘direction towards infinity’, for example Gromov compactifications of locally finite hyperbolic graphs, or metric completions of edge-length graphs [81], and the work presented here started out investigating whether for instance compactifications of locally finite graphs with a circle as boundary at infinity are Eulerian in this sense.




FIGURE 14.1. Three hyperbolic Eulerian structures.

Here we meet our second, topological motivation: by the Hahn-Mazurkiewicz Theorem, a space is the continuous image of the circle if and only if it is a *Peano continuum* – a compact, metrisable, connected and locally connected space. Originating with Hilbert’s observation (1891) [96] that the square is a continuous image of the circle so that each

point is visited at most three times, the natural question arises which properties beyond ‘Peano’ are needed to guarantee the existence of well-behaved such continuous surjections. Achieving additional control over the surjections from the circle, however, is a notorious open problem in continuum theory discussed, for example, in Nöbling (1933) [132], Harold (1940, 1942) [92, 93], Ward (1977) [162], Treybig & Ward (1981) [155, §4], Treybig (1983) [154], and Bula, Nikiel & Tymchatyn (1994) [41]. The latter six authors were particularly interested in the existence of *strongly irreducible* maps from the circle, continuous surjections $g: S^1 \rightarrow X$ such that for any proper closed subset $A \subsetneq S^1$ we have $g(A) \subsetneq g(S^1)$. It may not be immediately clear how the property of being strongly irreducible is related to Eulerianity. But using the intermediate value theorem, it is an easy exercise to verify that a strongly irreducible map from S^1 onto a finite multi-graph G must sweep through each edge of the graph precisely once without stopping or turning. Hence, a finite graph is Eulerian if and only if it is a strongly irreducible image of the circle. This suggests a second natural candidate for calling an arbitrary Peano continuum Eulerian, namely if it is the strongly irreducible image of the circle.

In this paper we achieve the following goals:

- (1) formalise the notion of an *Eulerian* continuum – all competing definitions in the literature are fortunately shown to be equivalent;
- (2) formulate a conjecture for characterising the Eulerian Peano continua, in a manner that naturally extends Nash-Williams’s condition, and which can be extended to a characterisation in the spirit of Euler; and
- (3) present far-reaching results in support of our conjecture, confirming it in particular for all one-dimensional Peano continua.

14.1.1.1. *Eulerianity.* Taking our cue from Bula, Nikiel and Tymchatyn [41], we say a space X is *Eulerian* if it is a strongly irreducible image of the circle, so there is a continuous surjection $g: S^1 \rightarrow X$ such that for any proper closed subset $A \subsetneq S^1$, we have $g(A) \subsetneq g(S^1) = X$. We also refer to such a map as an *Eulerian map*.

Extending Diestel & Kühn’s definition [56], let us say a space X is *edge-wise Eulerian* if there is a continuous map of S^1 onto X which sweeps through each free arc of X exactly once. Here a *free arc* is any inclusion-maximal open subset homeomorphic to $(0, 1)$, and by ‘sweeps once through a free arc’ we mean a map such that the preimage of every point in a free arc is a singleton. We also refer to such a map as an *edge-wise Eulerian map*.

As remarked earlier, every Eulerian map from S^1 onto a space X is edge-wise Eulerian. The converse, however, does not hold on the level of individual functions. Still, as our main result in Chapter 14.2, we establish that a space is edge-wise Eulerian if and only if it is Eulerian. The added flexibility of edge-wise Eulerian over Eulerian maps is convenient for constructions, and Chapter 14.3 continues with the development of a versatile framework to establish their existence, which we call *approximating sequences of Eulerian*

decompositions. Overall, our main results on the different concepts of Eulerian spaces can be summarised as follows.

THEOREM 14.1.1. *For a Peano continuum X , the following are equivalent:*

- (i) X is Eulerian,
- (ii) X is edge-wise Eulerian, and
- (iii) X admits an approximating sequence of Eulerian decompositions.

The first equivalence (i) \Leftrightarrow (ii) is the topic of Chapter 14.2, and relies on a function space Baire category argument. The second equivalence (ii) \Leftrightarrow (iii) is the topic of Chapter 14.3, and combines the classical strategy of the Hahn-Mazurkiewicz Theorem with inverse limit methods developed by Espinoza and the authors in [70].

14.1.1.2. *The conjecture.* Let X be a Peano continuum. As above a *free arc* is an inclusion-maximal open subset of X homeomorphic to $(0, 1)$. We think of free arcs as being the ‘edges’ of X . Write $E = E(X)$ for the collection of edges of X . For a subset $F \subseteq E$, we write for brevity $X - F := X \setminus \bigcup F$. The *ground-space* of X is the (compact metrisable) space $\mathfrak{G}(X) := X - E$. Every edge of a Peano continuum has two *end points*, which may agree, in which case the edge is a *loop*. An *edge cut* of a Peano continuum X is a non-empty set $F \subseteq E(X)$ of edges crossing a partition $A \oplus B$ of $\mathfrak{G}(X)$ into two disjoint clopen subsets A and B . In this case, we also write $F = E(A, B)$. Every edge cut of a Peano continuum is finite. (See Section 14.1.3.1 for a record of basic results on edge cuts.) With this set-up, we conjecture that Nash-Williams’s edge cut characterisation of finite Eulerian graphs extends to all Peano continua:

CONJECTURE 14.1.2 (The Eulerianity Conjecture).

A Peano continuum X is Eulerian if and only if every edge cut of X is even.

We also say that X satisfies the *even-cut condition* or has the *even-cut property*. The condition that an Eulerian continuum has the even-cut property is clearly necessary: if g is an (edge-wise) Eulerian map for X , and F is the set of edges crossing a disconnection $A \oplus B$ of $\mathfrak{G}(X)$, then consider g as a ‘path’ with start and end point in A , and observe that g must sweep through the edges of F in pairs, from A to B and then back. Also note that an affirmative answer to the conjecture implies the truth of (i) \Leftrightarrow (ii) in Theorem 14.1.1.

When X is the space underlying a finite multi-graph G , then, suppressing vertices of degree two, the edges of X (free arcs) correspond to edges of G , and the ground space of X corresponds to the vertex set of G . Hence our conjecture naturally encompasses the second characterisation for finite Eulerian graphs. Also, Diestel and Kühn’s Eulerianity result [56, Theorem 7.2] for the Freudenthal compactification FG of a connected, locally finite graph G mentioned above falls under the scope of Conjecture 14.1.2: the ground space of FG consists of all vertices and ends of G , and edge cuts of FG correspond precisely

to the finite edge cuts of G .¹ The same holds for Georgakopoulos's [80] extension of this result to standard subspaces of Freudenthal compactifications of locally finite graphs.

For Peano continua, Harrold [92] showed in 1940 that every Peano continuum without free arcs is Eulerian,² and in 1994, Bula, Nikiel and Tymchatyn [41, Theorem 3, Example 2] showed that every Peano continuum obtained by adding a dense collection of free arcs to a Peano continuum is Eulerian.³ Both results are in line with Conjecture 14.1.2, as with connected ground spaces, these examples have no edge cuts whatsoever, and so the even-cut condition is trivially satisfied. In the same paper, Bula, Nikiel and Tymchatyn settled when so-called 'completely regular' continua are Eulerian. Call a continuum *graph-like*⁴ if its ground space is zero-dimensional, see [30, 49, 153]. In [70], Espinoza and the authors showed that a continuum is graph-like if and only if it is completely regular, and equivalently, if and only if it is a standard subspace of the Freudenthal compactification of a locally finite, connected graph. Hence, also these spaces fall under Conjecture 14.1.2.

14.1.1.3. *Towards the Eulerianity conjecture.* All previously known cases for Conjecture 14.1.2 fall under the dichotomy that there are either no free arcs at all, or the free arcs are dense. Our first result towards Conjecture 14.1.2, which we call the 'reduction theorem', clears the middle ground: the problem of establishing the Eulerianity Conjecture for a given space can always be reduced to a space with the same ground space in which the edges are dense. For brevity, such a Peano continuum in which the edges are dense will also be called a *Peano graph*. Note that Peano graphs are precisely the spaces that can be obtained as Peano compactifications of countable, locally finite graphs.

THEOREM 14.1.3 (Reduction Result). *If the Conjecture 14.1.2 holds for all [loopless] Peano graphs, then it holds in general.*

This result is proved in Theorems 14.2.12 and 14.2.14. The class of Peano graphs is still surprisingly complex: in Theorem 14.2.5 we observe that there is no restriction on the possible ground spaces of an (Eulerian) Peano graph. Our remaining results establish Conjecture 14.1.2 for three large classes of Peano continua, which together subsume and extend every result known about the Eulerianity of infinite graphs and of continua to date.

THEOREM 14.1.4. *Conjecture 14.1.2 holds for every Peano continuum whose ground space*

¹For every finite edge cut $E(A, V \setminus A)$ of the graph G , the properties of the Freudenthal compactification guarantee that A and $V \setminus A$ have disjoint closures in FG , and so $E_G(A, V \setminus A) = E_{FG}(\overline{A}, \overline{V \setminus A})$.

²To be precise, Harrold has shown in [92] that Peano continua in which *the non-local separating points are dense* are strongly irreducible images of I and S^1 . However, this condition is equivalent to not having free arcs, as remarked in Harrold's later paper [93].

³As stated, [41, Theorem 3] excludes edges which are loops, but this assumption is unnecessary.

⁴This notion of 'graph-like', by now firmly established in graph theory, is not to be confused with the notion of arc-like, tree-like and graph-like in continuum theory, which we shall not use in this paper.

- (A) consists of finitely many Peano continua, or
 (B) is homeomorphic to a product $V \times P$, where V is zero-dimensional and P a Peano continuum, or
 (C) is at most one-dimensional.⁵

Indeed, the main results of Harrold [92] and Bula-Nikiel-Tymchatyn [41, Theorem 3] follow either from (A) (where the ground space is a single Peano component, and the free arcs are either absent or dense) or (B) (by taking V to be a singleton). Diestel and Kühn’s results for Freudenthal compactifications of graphs, and the results about graph-like spaces from [70] are covered either by (B) (by taking P to be a singleton) or indeed (C).

However, (C) goes significantly beyond these results. Consider for example hyperbolic groups with one-dimensional boundaries, whose Gromov boundaries, provided the groups are one-ended, are either homeomorphic to S^1 , the Sierpinski carpet, or to the Menger curve [98, Theorem 4]. Interestingly, ‘generic’ finitely presentable groups are hyperbolic and have the Menger curve as boundary [46], thus falling once again under (C). A geometrically interesting class of spaces with S^1 boundary is given by the regular tessellations $T(n, k)$ of the hyperbolic plane where precisely k regular n -gons surround each vertex (for $1/k + 1/n < 1/2$). Since S^1 is connected, edge cuts in these spaces can only contain finitely many vertices on one side, so (C) implies that $T(n, k)$ is Eulerian if and only if k is even.




FIGURE 14.2. The spaces X and Y with ground-space in black and edges in red.

Our result (B) answers an open question in the literature, namely (a variant of) [41, Problem 3]. Its strength lies in supporting Conjecture 14.1.2 by providing non-trivial affirmative examples in all dimensions. To illustrate (B), consider the ‘fractal’ spaces X and Y with ground-space $\mathfrak{G}(X) = \mathfrak{G}(Y) = C \times [0, 1]$ in Figure 14.2. Both spaces X and Y clearly satisfy the even-cut condition and so are Eulerian by (B). Alternatively, due to the fractal nature of these specific examples, it is possible in both cases to give a geometric, recursive definition of an (edge-wise) Eulerian map in the spirit of Hilbert [96]. For a different example in which the free arcs are not necessarily dense, consider a Peano continuum X with ground-space a convergent sequence of unit squares, $\mathfrak{G}(X) = (\omega + 1) \times I^2$, satisfying the even-cut condition.

⁵Equivalently: the Eulerianity Conjecture holds for all one-dimensional Peano continua.




FIGURE 14.3. A Peano continuum satisfying the even-cut condition with ground space a convergent sequence of squares. Local connectedness implies that endpoints of edges are dense in the right limit square.

All three results in Theorem 14.1.4 rely on our earlier equivalences for Eulerianity given in Theorem 14.1.1. First, (A) follows from an appealing application of the equivalence (i) \Leftrightarrow (ii) in Theorem 14.1.4, and will be given, after introducing a modicum of notation, right at the end of the introduction in Section 14.1.3.3.

The other two results, (B) and (C), utilise the implication (iii) \Rightarrow (i) of Theorem 14.1.1, and, being rather more involved, occupy the final two chapters of this paper, Chapter 14.4 and 14.5. As indicated, for both cases the objective is, relying on nothing but the even-cut property, to construct an approximating sequence of Eulerian decompositions for these spaces, in other words, to show that the even-cut condition implies property (iii). Carrying out this program requires a combination of powerful techniques from both topology and graph theory. Topologically, we rely on Bing's [24, 25, 26] and Anderson's [7] theory of brick partitions, widely regarded as the single most effective structural tool in the theory of Peano continua. Combinatorially, we rely on the cycle space theory for locally finite graphs developed in the past 15 years by Diestel et al., see [53] for a survey, and its extension to graph-like spaces developed in [30, 70]. Roughly, these ingredients are then combined as follows: first, brick partitions are used to supply a preliminary decomposition of our spaces, whose parts are then carefully modified using combinatorial tools in order to gain control over the edge cuts of the individual parts.

(4) Open problems. The main open problem is to establish Conjecture 14.1.2 for all Peano continua. Motivated by the naturally occurring examples of hyperbolic boundaries, interesting partial results may be about Peano compactifications of locally finite graphs with remainder homeomorphic to S^2 , S^3 and generally S^n , and we hope that these examples can also be approached using our theory of *approximating sequences of Eulerian decompositions*. Slightly more general, a result saying that all 2-dimensional Peano graphs satisfy Conjecture 14.1.2 would be welcome, and might be in reach once the S^2 case has been settled.

14.1.2. Related Conjectures for the Eulerian Problem.

14.1.2.1. *Equivalent conjectures.* While calling the free arcs of a Peano continuum X ‘edges’, the points of $\mathfrak{G}(X) = X - E(X)$ should generally not be considered the ‘vertices’ of X . Instead the ‘vertices’ of X correspond to the connected components of $\mathfrak{G}(X)$. Let X_{\sim} denote the quotient of X where we collapse, one by one, each component of the ground space $\mathfrak{G}(X)$ to a point. Note that X_{\sim} is a continuum with $E(X_{\sim}) = E(X)$ and has zero-dimensional ground-space. In other words, the continuum X_{\sim} is a graph-like Peano continuum. Moreover, every edge cut of X corresponds to an edge cut of X_{\sim} and vice versa. Since we know from [70] that graph-like continua are Eulerian if and only if they satisfy the even-cut condition, the following is equivalent to the Conjecture 14.1.2:

CONJECTURE 14.1.5.

A Peano continuum X is Eulerian if and only if X_{\sim} is Eulerian.

Since points in a Peano continuum other than a finite graph may have infinite order, the definition of when a point has ‘even degree’ is problematic. Note that these difficulties for generalising Euler’s characterisation of Eulerian graphs occur already in the case of locally finite graphs, cf. [39, Fig. 2] and [20]. Nevertheless, from [70] we know that a graph-like continuum Y is Eulerian if and only if every point $y \in \mathfrak{G}(Y)$ has *even degree* in the sense that there exists a clopen neighbourhood A of y in $\mathfrak{G}(Y)$ such that for every clopen subset B of $\mathfrak{G}(Y)$ with $y \in B \subseteq A$, the edge cut $E(B, \mathfrak{G}(Y) \setminus B)$ is even. Thus another equivalent version of Conjecture 14.1.2 is that:

CONJECTURE 14.1.6.

A Peano continuum X is Eulerian if and only if every vertex of X_{\sim} has even degree.

14.1.2.2. *Circle decompositions.* Recall that another classical characterisation of finite Eulerian multi-graphs, due to Veblen, is that the edge set of the graph can be decomposed into edge-disjoint cycles, see [54, 1.9.1]. Accordingly, let us say that the edge set of a Peano continuum X can be *decomposed into edge-disjoint circles* if there is a collection of edge-disjoint copies of S^1 contained in X such that each edge of X is contained in precisely one of them. Generalising the corresponding equivalence for graphs due to Nash-Williams [123], we shall prove in Theorem 14.5.15 that a Peano continuum has the even-cut property if and only if its edge set can be decomposed into edge-disjoint circles. Consequently, another equivalent version of Conjecture 14.1.2 is that:

CONJECTURE 14.1.7. *A Peano continuum is Eulerian if and only if its edge set can be decomposed into edge-disjoint circles.*

14.1.2.3. *Open Eulerian spaces.* A finite multi-graph is *open Eulerian* if there is a walk starting and ending at distinct vertices, using every edge of the graph precisely once. The open Eulerian multi-graphs are precisely the connected graphs for which all but two

vertices have even degree. A Peano continuum X is *open Eulerian* if it is the strongly irreducible image of a map from the unit interval $I = [0, 1]$. Let $x \neq y \in X$, and let X_{xy} denote the Peano continuum where we add a new free arc from x to y . Then X is open Eulerian from x to y if and only if X_{xy} is Eulerian. Thus, Conjecture 14.1.2 may be used to characterise open Eulerian spaces. Moreover, applying the degree characterisation from [70] when a graph-like continuum is open Eulerian, the following is again equivalent, via the X_{xy} construction, to Conjecture 14.1.2:

CONJECTURE 14.1.8.

A Peano continuum X is open Eulerian if and only if all but two vertices of X_\sim have even degree.

To our knowledge, this conjecture is the first attempt to put forward a proposal for the characterisation of open Eulerian continua and, if correct, would provide a complete answer to [155, Problem 3]. Interestingly, if a Peano continuum X is open Eulerian from x to y for $x, y \in \mathfrak{G}(X)$, then Conjecture 14.1.2 predicts that X is also open Eulerian from x' to y' for all x' (respectively y') that lie in the same component of $\mathfrak{G}(X)$ as x (respectively y).

14.1.2.4. *The Bula-Nikiel-Tymchatyn conjecture.* Our conjecture is not the only contender to characterise Eulerian continua. Bula et al [41] have proposed an alternative, which is, however, difficult to verify in concrete cases, and implied by Conjecture 14.1.2.

A point x of a Peano continuum X is said to be *locally separating* if there is a connected open subset U of X such that $U \setminus \{x\}$ is disconnected. The set $N(X)$ denotes the set of all x in X such that x is *not* locally separating in X . By Y_X denote the quotient of X where we collapse every component of $\overline{N(X)}$ to a single point. By [41, Theorem 2], if Y_X is non-trivial then it is a (cyclically completely regular) Peano continuum, and if X is Eulerian then so is Y_X . The following is from [41, Problem 1]:

CONJECTURE 14.1.9 (Bula, Nikiel & Tymchatyn).

A Peano continuum X is Eulerian if and only if Y_X is Eulerian.

Since interior points of edges are locally separating, and $\mathfrak{G}(X)$ is closed, we have $\overline{N(X)} \subseteq \mathfrak{G}(X)$, and hence $(Y_X)_\sim = X_\sim$. In particular, edge cuts of Y_X are in bijective correspondence with edge cuts of X , and hence the truth of Conjecture 14.1.5 implies the truth of Conjecture 14.1.9. Furthermore, the difference between the two conjectures is not simply formal, as the two quotient spaces Y_X and X_\sim may differ: fix a finite tree T and add to it a dense, zero-sequence of loops. Denote the resulting Peano continuum by X , and note that $\mathfrak{G}(X) = T$. Since T is connected, X_\sim is a Hawaiian earring. However, as every point of T apart from the finitely many leaves remains locally separating in X , we have $X = Y_X$. For a more interesting example where Y_X and X_\sim differ, consider a topological sine curve Z . Form a Peano continuum X with $\mathfrak{G}(X) = Z$ by first adding a

dense collection of loops to Z (to guarantee $\mathfrak{G}(X) = Z$), and then also adding a nowhere dense collection of free arcs between points on the sine function-graph and points on the y -axis of Z (to make X locally connected). Again, X_\sim is the Hawaiian earring, but Y_X is an interval with a dense collection of free arcs, since $\overline{N(X)}$ corresponds precisely to the y -axis of Z .

14.1.2.5. *Further consequences.* Harrold has shown, generalising a result by Nöbling [132], that every Peano continuum X is the image of a map $g: S^1 \rightarrow X$ that sweeps through every free arc at most twice, [93, Theorem 1 ff.]. We observe here that this result is implied by Conjecture 14.1.2: for an arbitrary Peano continuum X , let \hat{X} denote the space where we add for each edge e of X one additional parallel edge \hat{e} . Then \hat{X} is again a Peano continuum (compare with Lemma 14.1.13 below) which now satisfies the even-cut condition. Hence, there is an Eulerian map $\hat{g}: S^1 \rightarrow \hat{X}$ that sweeps through every free arc of \hat{X} precisely once. But then it is clear that \hat{g} naturally induces a map $g: S^1 \rightarrow X$ that uses the original edge e a second time instead of \hat{e} for each $e \in E(X)$. By construction, g has the desired property that it sweeps through every free arc of X precisely twice.

14.1.3. Notation and Essentials. Throughout this paper, all topological spaces are metrisable, and all maps are continuous. A *continuum* is a compact connected metrisable space, a *Peano continuum* is a continuum which is locally connected, and a *Peano graph* is a Peano continuum in which the edges are dense. We write $\mathbb{N} = \{0, 1, 2, \dots\}$ and $[n] = \{1, 2, \dots, n\}$ for $n \in \mathbb{N}$. If A is a subset of the domain of a function g , then we denote by $g \upharpoonright A$ the restriction of g to A .

Let (X, d) be a metric space, and $A, B \subseteq X$ and \mathcal{A} a family of subsets of X . We use $A \sqcup B$ to denote disjoint union. A *clopen partition* of a space V is a partition of V into pairwise disjoint clopen subsets. If V is compact, then any clopen partition is finite, and we denote by $\Pi(V)$ the collection of clopen partitions of V . For $\varepsilon > 0$, let $B_\varepsilon(x)$ denote the open ε ball around x . Further, we write $\text{dist}(A, B) = \inf \{d(a, b) : a \in A, b \in B\}$, $\text{diam}(A) := \sup \{d(a, b) : a, b \in A\}$, and $\text{mesh}(\mathcal{A}) := \sup \{\text{diam}(A) : A \in \mathcal{A}\}$. Let X be a metrisable compactum. Then \mathcal{A} is said to be a *null-family*, if for any $\varepsilon > 0$, the collection $\{A \in \mathcal{A} : \text{diam}(A) > \varepsilon\}$ is finite. By compactness, this does not depend on the metric for X . Any null-family \mathcal{A} contains only countably many non-singleton sets. A countable null-family \mathcal{A} is said to be a *zero-sequence*. This is equivalent to saying that whenever an enumeration $\mathcal{A} = \{A_1, A_2, \dots\}$ is chosen, then $\text{diam}(A_n) \rightarrow 0$ as $n \rightarrow \infty$.

Let $A, B \subseteq X$ be disjoint closed subsets. An A – B -arc in X is an arc whose first endpoint lies in A , whose last end-point lies in B , and which is otherwise disjoint from $A \cup B$. Finally, a subset $A \subseteq X$ is *regular closed* if $A = \overline{\text{int}(A)}$.

14.1.3.1. *Edge cuts in Peano continua.* Free arcs in Peano continua behave much the same as edges in finite graphs, and statements to this effect can be found for example in [41] or [129]. To make this paper accessible for readers with more of a combinatorial

background, we offer brief indications how to prove these basic facts with a minimal topological background, relying only on the fact that Peano continua are (locally) arc-connected.

If e is an edge of X , then any point in $\partial e = \bar{e} \setminus e$ is called an *endpoint* of e . Moreover, with some fixed homeomorphism $e \cong (0, 1)$ in mind, we write $e(x) \in e$ for $x \in (0, 1)$ to mean the corresponding interior point on e , and also write $[a, b]_e$ for the set $\{e(x) : x \in [a, b]\}$ and similar for other subsets of the interval.

LEMMA 14.1.10. *Edges of a Peano continuum X are pairwise disjoint, unless $X = S^1$.*

PROOF. Suppose e and f are two distinct free arcs which intersect. Since each free arc is maximal with respect to set-inclusion, this amounts to the statement that all $e \setminus f$, $f \setminus e$ and $e \cap f$ are non-empty. Let A be a component of $e \cap f$. Then A is a proper subinterval of e , and so one endpoint a of A lies in $e \setminus f$. Now if there was a half-open interval $[a, a + \varepsilon)_e \subseteq e \setminus f$, then this contradicts maximality of f . But then connectedness of f implies that $e \setminus \{a\} \subseteq f$. However, it follows that $e \cup f = \bar{f} = f \cup \{a\}$ is homeomorphic to S^1 , and is clopen in X . So by connectedness, $X = S^1$. \square

For the remainder of this paper, when investigating Conjecture 14.1.2 for a space X we always implicitly assume that X is not a simple closed curve, implying that the edge set $E(X)$ consists of disjoint open sets and that $\mathfrak{G}(X)$ is non-empty.

LEMMA 14.1.11. *Let X be a Peano continuum.*

- (a) *Every edge (free arc) in X contains at most two endpoints.*
- (b) *Removing an edge from X creates at most two connected components which are again Peano continua. Thus, removing k edges from a Peano continuum results in at most $k + 1$ components, all of which are again Peano.*
- (c) *If $X \neq S^1$, the edges $E(X)$ form a zero-sequence of disjoint open subsets.*
- (d) *Every edge cut of X is finite.*

PROOF. (a) Consider a free arc $e \cong (0, 1)$ of a Peano continuum X . Write for the moment $e(0) = \overline{(0, \frac{1}{2}]} \setminus e$ and $e(1) = \overline{[\frac{1}{2}, 1)} \setminus e$. By symmetry, it suffices to show that $e(0)$ is a singleton. By compactness, it is certainly non-empty. Next, since X is locally arc-connected, there exists an $\{\frac{1}{2}\} - e(0)$ -arc α in X so that $(0, \frac{1}{2}] \subseteq \alpha$, and so $\alpha \setminus (0, \frac{1}{2}]$ is precisely the second end-point of α . However, compactness of α gives $\overline{(0, \frac{1}{2}]} \subseteq \alpha$, from which it is clear that $e(0)$ consists of at most one point.⁶

(b) Otherwise, for some edge e the space $X - e$ has a partition into three non-empty, pairwise disjoint compact subsets A, B, C . By (a), it follows that one of them, say A , does

⁶The assumption on local connectedness in (a) is necessary, as witnessed by the unique free arc of the topological sine curve, [121, 1.5].

not contain an endpoint of e . But then A against $B \cup C \cup \bar{e}$ forms a partition of X into two non-empty, pairwise disjoint compact subsets, contradicting connectedness of X .⁷

(c) As a collection of disjoint open subsets (Lemma 14.1.10) in a compact metrisable space, $E(X)$ must be countable, [63, 4.1.15]. Now if $E(X)$ does not form a zero-sequence, then there is $\varepsilon > 0$ and infinitely many distinct edges $\{e_1, e_2, e_3, \dots\} \subseteq E(X)$ each containing three successive points $x_n^1 < x_n^2 < x_n^3 \in e_n$ such that $d(x_n^i, x_n^j) \geq \varepsilon$ for all $i \neq j \in [3]$ and $n \in \mathbb{N}$. By moving to convergent subsequences and relabelling, we may assume that $x_n^i \rightarrow x^i$ for all $i \in [3]$ as $n \rightarrow \infty$, and so $d(x^i, x^j) \geq \varepsilon$ for all $i \neq j \in [3]$. However, by local arc-connectedness, for large enough n there exist arcs from x^2 to x^3 of diameter less than ε , a contradiction.⁸

(d) Trivial for $X = S^1$. Otherwise, the assertion follows from (c) since the sets of any topological disconnection $A \oplus B$ of $\mathfrak{G}(X)$ are disjoint compact, so have $\text{dist}(A, B) > 0$.⁹ \square

From now on, if e is an edge in a Peano continuum X , let $e(0), e(1) \in \mathfrak{G}(X)$ denote the two endpoints of that edge. If x is an end-point of an edge e , we also write $x \sim e$, or write $e = xy$ to mean that $e(i) = x$ and $e(1 - i) = y$ for $i = 0$ or $i = 1$. It is convenient to write $e(x)$ for $x \in (0, 1)$ to mean the corresponding interior point on e , where we choose our parametrisation so that $e(x)$ is continuous for $x \in [0, 1]$. Next, recall from the introduction that for a subset $F \subseteq E(X)$, we write for brevity $X - F := X \setminus \bigcup F$, and so $\mathfrak{G}(X) := X - E(X)$. If $F = \{f\}$ is a singleton, we write $X - f$ instead of $X - \{f\}$. Let $X[F] = \overline{\bigcup F} \subseteq X$ be the subspace of X induced by F . Similarly, for $U \subseteq \mathfrak{G}(X)$, write $E(U) = \{e = xy \in E(X) : \{x, y\} \subseteq U\}$ for the induced edge set of U , and set $X[U] = U \cup E(U)$. Finally, an edge set $F \subseteq E(X)$ is called *sparse (in X)* if $X[F]$ is a graph-like compactum. This notion will be of crucial importance in the final two chapters. Note that if F is sparse, then so is every $F' \subseteq F$.

A subspace Y of a Peano continuum X is a *standard subspace* if Y contains every edge from X it intersects. Finally, two standard subspaces Y_1, Y_2 of X are *edge-disjoint* if every edge of X is contained in at most one Y_i .

14.1.3.2. *Waiting times for maps from the circle.* A map $g: I \rightarrow X$ or $g: S^1 \rightarrow X$ which is nowhere constant is also called *light*. The first part of the next lemma is about ‘avoiding waiting times’: given a map $g: I \rightarrow X$, by contracting all non-trivial intervals in $g^{-1}(x)$ for each $x \in X$, one obtains an associated map that traces out the same path but is, by construction, nowhere constant. The second part describes, in a sense, the converse operation, and says that given a map $g: I \rightarrow X$, we may add a countable list of waiting intervals, so that the resulting map still traces out the same path.

⁷Alternatively, assertion (b) can be concluded from the *boundary bumping lemma* [121, 5.7].

⁸Alternatively, assertion (c) follows from *compactness of the hyperspace* [121, 4.14].

⁹Alternatively, for a proof that does not rely on (c), use normality to find disjoint open sets $U, V \subseteq X$ separating A from B , forming together with $E(A, B)$ an open cover of the compact X .

LEMMA 14.1.12. *Let X be a non-trivial Peano continuum.*

- (a) *For every continuous surjection $g: I \rightarrow X$, there is a continuous light surjection $\hat{g}: I \rightarrow X$ and a monotonically increasing $m: I \rightarrow I$ such that $g = \hat{g} \circ m$.*
- (b) *For every surjection $g: I \rightarrow X$ and any sequence (x_0, x_1, \dots) in X , there is a zero-sequence (J_0, J_1, \dots) of non-trivial disjoint closed intervals of I and monotonically increasing $m: I \rightarrow I$ such that $\tilde{g} = g \circ m: I \rightarrow X$ maps each J_n to x_n .*

Furthermore, the same assertions hold *mutatis mutandis* for maps $g: S^1 \rightarrow X$.

PROOF. Assertion (a) follows from the *monotone-light-factorisation* [121, 13.3], and relies on the fact that a quotient of I over closed intervals and points is again homeomorphic to I , cf. [121, 13.4 & 8.22]. For (b), pick points $y_n \in g^{-1}(x_n)$ and construct a uniformly converging sequence of monotone surjections $m_n: I \rightarrow I$ such that $m_n^{-1}(y_i)$ contains a non-trivial interval J_i for $i \in [n]$. The furthermore-part follows by viewing maps $g: S^1 \rightarrow X$ as maps $g: I \rightarrow X$ with $f(0) = f(1)$. \square

We first illustrate the use of Lemma 14.1.12(b) in following well-known fact.

LEMMA 14.1.13. *Suppose X is a compact metrisable space, and Y, Y_1, Y_2, \dots a zero-sequence of Peano subcontinua of X such that $Y \cap Y_n \neq \emptyset$ for all $n \in \mathbb{N}$. Then $Y' := Y \cup \bigcup_{n \in \mathbb{N}} Y_n \subseteq X$ is a Peano continuum.*

PROOF. Pick $y_n \in Y_n \cap Y$ for each $n \in \mathbb{N}$. By Lemma 14.1.12(b), there is a surjection $h: I \rightarrow Y$ and non-trivial disjoint closed intervals $J_n \subseteq I$ such that $h(J_n) = \{y_n\}$. Fix surjections $h_n: I \rightarrow Y_n$ such that $h_n(0) = h_n(1) = y_n$. Construct surjections $g_n: I \rightarrow Y \cup \bigcup_{i \in [n]} Y_i$ by replacing $h \upharpoonright J_i$ by h_i for $i \in [n]$. Then g_n converges uniformly to a continuous surjection $g: I \rightarrow Y'$ as desired. \square

Our second illustration of Lemma 14.1.12(b) lets us combine edge-wise Eulerian maps:

LEMMA 14.1.14. *Let X be a Peano continuum and suppose that Y, Y_1, Y_2, \dots is a zero-sequence of edge-disjoint standard Peano subcontinua of X with $X = Y \cup \bigcup_{n \in \mathbb{N}} Y_n$ such that $Y_n \cap Y \neq \emptyset$. If Y and all Y_n are edge-wise Eulerian, then so is X .*

PROOF. Follow the same proof as in Lemma 14.1.13, but start with edge-wise Eulerian surjections $h: S^1 \rightarrow Y$ and $h_n: I \rightarrow Y_n$. \square

14.1.3.3. *An application of the equivalence for edge-wise Eulerianity.* We conclude our introduction with a proof of Theorem 14.1.4(A). Indeed, given (ii) \Rightarrow (i) of Theorem 14.1.1, the proof of (A) reduces to the observation that for these types of spaces, there is a simple procedure for finding an edge-wise Eulerian surjection.

PROOF OF THEOREM 14.1.4(A) FROM THEOREM 14.1.1. Let X be a Peano continuum such that for its ground space we have $\mathfrak{G}(X) = Z_1 \oplus Z_2 \oplus \dots \oplus Z_\ell$ where each Z_i is a Peano continuum. Assume further that X has the even-cut property. By (i) \Leftrightarrow (ii)

of Theorem 14.1.1, to complete the proof it suffices to show the existence of an edge-wise Eulerian surjection onto X .

Partition the edge set $E(X) = E' \sqcup E''$ where $E' = \bigcup_{i \in [l]} E(Z_i, Z \setminus Z_i)$ consists of the finitely many cross edges between the components of $\mathfrak{G}(X)$, and $E'' = E \setminus E'$ consists of all the edges that have both endpoints attached to the same component of $\mathfrak{G}(X)$.

Since X satisfies the even-cut condition, $X_{\sim}[E']$ is a finite Eulerian multi-graph. Take any Eulerian walk W on $X_{\sim}[E']$ and extend to an edge-wise Eulerian surjection onto $Y = \mathfrak{G}(X) \cup \bigcup E'$ by inserting, between any two successive edges $eZ_i e'$ on W in $(X[E'])_{\sim}$ a surjection onto Z_i from the end vertex of e to the end vertex of e' in Z_i .

Now by Lemma 14.1.11, the set $E'' = \{e_n = x_n y_n : n \in K\}$ for $K \subseteq \mathbb{N}$ is either finite, or a zero-sequence of edges. Since Peano continua are uniformly locally arc-connected, [107, Ch. VI, §50, II Theorem 4], for each $n \in K$ there is an $x_n - y_n$ arc α_n in $\mathfrak{G}(X)$ such that $\text{diam}(\alpha_n) \rightarrow 0$. Then $Y_n = e_n \cup \alpha_n$ forms a zero-sequence of simple closed curves. Since Y and each Y_n are pairwise edge-disjoint standard subspaces which are all edge-wise Eulerian, it follows from Lemma 14.1.14 that $X = Y \cup \bigcup_{n \in K} Y_n$ is edge-wise Eulerian, too. \square

14.2. Eulerian maps and Peano graphs

14.2.1. Overview. Recall from the introduction that we had two, seemingly competing notions for generalised Euler tours in a Peano continuum X . First, the notion of an *Eulerian map*, a continuous surjection g from the circle that is strongly irreducible: no proper closed subset A of the circle satisfies $g(A) = g(S^1)$. And second the notion of an *edge-wise Eulerian map*, a continuous surjection from the circle that sweeps through every edge of X exactly once. In this chapter we show that both notions for an Eulerian space are in fact equivalent, and thus establish (i) \Leftrightarrow (ii) of Theorem 14.1.1: a Peano continuum is Eulerian if and only if it is edge-wise Eulerian. One implication, namely (i) \Rightarrow (ii), is straightforward.

LEMMA 14.2.1. *Every Eulerian map is edge-wise Eulerian.*

PROOF. Let us first note that by the intermediate value theorem, every strongly irreducible map $g: I \rightarrow I$ is injective. Otherwise, there are $a < b$ such that $g(a) = x = g(b)$. Since g being constant on $[a, b]$ results in an immediate contradiction, there exists $a < c < b$ such that say $g(c) > x$. By the intermediate value theorem, the interval $[x, g(c)]$ is covered by both $g \upharpoonright [a, c]$ and $g \upharpoonright [c, b]$. But then it is clear that for some non-trivial open interval $U \subseteq [a, c]$ with $g(U) \subseteq [x, g(c)]$ we have that $g(I \setminus U) = g(I)$, a contradiction.

To prove the lemma, suppose then there is a strongly irreducible map $g: S^1 \rightarrow X$ onto some Peano continuum X , an edge $e \in E(X)$ and an interior point $x \in e$ such that $g^{-1}(x)$ contains at least two distinct points a and b . By continuity, there are disjoint closed subintervals A and $B \subseteq S^1$ containing respectively a and b in their interior such that $g(A)$




FIGURE 14.4. Admissible trace of an edge-wise Eulerian map on the left, and an Eulerian map on the right.

and $g(B) \subseteq e$. By the first part, both $g \upharpoonright A$ and $g \upharpoonright B$ are injective embeddings, and so $g(A)$ and $g(B)$ are subintervals of e containing x in their interior. Thus, there is an open interval $V \subseteq e$ with $x \in V \subseteq g(A) \cap g(B)$. But then for some non-trivial open interval $U \subseteq A$ with $g(U) \subseteq V$ we have that $g(S^1 \setminus U) = X$, a contradiction. \square

The converse of Lemma 14.2.1, however, does not hold in general, and so the equivalence of Eulerian and edge-wise Eulerian spaces cannot hold function-wise: we already observed that edge-wise Eulerian maps are allowed to pause at points in the ground space. Much more significantly, however, consider for example the hyperbolic 4-regular tree Y from the introduction, where an edge-wise Eulerian map is allowed to trace out non-trivial paths on the boundary circle of Y , whereas an Eulerian map is not, as in the following Figure 14.4. Indeed, if say $g \upharpoonright [a, b]$ stays on the boundary for a non-trivial time interval $[a, b] \subseteq S^1$, then $g(S^1 \setminus (a, b))$, being closed and covering (the closure of) all edges of Y , must be the whole space (as $E(Y)$ is dense in Y), contradicting the defining property of an Eulerian map. Instead, to establish $(ii) \Rightarrow (i)$ in Theorem 14.1.1, we prove that if there exists an edge-wise Eulerian map g for X , then there also exists an Eulerian map h for X . First, in Section 14.2.2 we establish a number of equivalent definitions for ‘strongly irreducible’. Most importantly, in the context of Peano graphs (Peano continua whose edges are dense) we can add to the equivalent descriptions that a map g from S^1 onto a Peano graph X is Eulerian if and only if it is edge-wise Eulerian and never spends a positive time interval in the ground space of X (meaning that $g^{-1}(\mathfrak{G}(X))$ does not contain a non-empty open interval), Theorem 14.2.3. In other words, this behaviour of Eulerian maps that we have seen above is not only necessary, but also sufficient. This natural geometric formulation of ‘Eulerian map’ will be the key to our proof of $(ii) \Rightarrow (i)$.

In order to harness this geometric intuition, our next step in Section 14.2.3 is to establish our reduction result mentioned in the introduction so that we may restrict ourselves to Peano graphs. More explicitly, given a Peano continuum X define a Peano graph X' by attaching to X a zero-sequence of loops to a countable dense subset of the interior of the

ground space of X . It is immediate that X satisfies the even-cut condition if and only if X' does. Crucially we show that X has an Eulerian map if and only if X' has one. Going forward we may always restrict ourselves to Peano graphs, and thus rely on the geometric intuition of an Eulerian map as described above.

Now the strategy is clear: given an edge-wise Eulerian map g , we need to modify it so that it remains edge-wise Eulerian, but no longer spends non-trivial time intervals in the ground space. For the problem that edge-wise Eulerian maps may pause at points of the ground space, there is an easy remedy: given any surjection $g: S^1 \rightarrow X$ onto a non-trivial Peano continuum, by contracting all non-trivial intervals in $g^{-1}(x)$ for each $x \in X$, one obtains an induced edge-wise Eulerian map $\hat{g}: S^1 \rightarrow X$ which is, by construction, nowhere constant, see Lemma 14.1.12(a). This observation already establishes (ii) \Rightarrow (i) for the class of all graph-like continua, and hence in particular for Freudenthal compactifications of locally finite connected graphs, simply because of the fact that their ground spaces, being totally disconnected, do not contain non-trivial arcs. In fact, this argument shows that for every Peano continuum X whose ground space $\mathfrak{G}(X)$ contains no non-trivial arcs – if $\mathfrak{G}(X)$ is totally disconnected, but also if it is for example a pseudoarc or any other hereditarily indecomposable continuum [121, 1.23] – every nowhere constant edge-wise Eulerian map for X is Eulerian. Finally, the harder case, where the ground space does contain non-trivial arcs, will be dealt with in Section 14.2.4.

14.2.2. Equivalent Definitions for Eulerian Maps. We begin by recalling the following well-studied classes of continuous functions. Let $g: X \rightarrow Y$ be a continuous map between continua X and Y . Then:

- g is *almost injective* if the set $\{x : g^{-1}(g(x)) = \{x\}\}$ is dense in X ;¹⁰
- g is *irreducible* if for all proper subcontinua $K \subsetneq X$, we have $g(K) \subsetneq g(X)$;
- g is *hereditarily irreducible* if for every subcontinuum K of X we have that $g \upharpoonright K$ is irreducible (equivalently, for every pair of subcontinua $A \subsetneq B$ in X , we have $g(A) \subsetneq g(B)$);
- g is *strongly irreducible* if for all closed subsets $A \subsetneq X$, we have $g(A) \subsetneq g(X)$;
- g is *arcwise increasing* if for every pair of arcs $A \subsetneq B$ in X we have $g(A) \subsetneq g(B)$.

In this section we relate these different types of maps, particularly when X is I or S^1 . The arguments are elementary, and in most cases known or at least folklore. As the results are important for us, and for completeness, we provide brief proofs. For discussions on hereditarily irreducible and arc-wise increasing images of finite graphs see [1, 71].

LEMMA 14.2.2. *Let $g: I \rightarrow Y$ be a continuous surjection. Then the following are equivalent: (a) g is arcwise increasing; (b) g is hereditarily irreducible; (c) g is strongly irreducible; and (d) g is almost injective.*

¹⁰The set of points of injectivity for an almost injective function between compact spaces is not just dense but a dense G_δ , and so large (co-meager) in the sense of Baire category, [166, Theorem VIII.10.1].

PROOF. Clearly, $(b) \Leftrightarrow (a)$. For $(a) \Rightarrow (c)$, show the contrapositive. So suppose there is a proper closed subset A of I whose image is $g(A) = Y$. Without loss of generality, $A = I \setminus (s, t)$ where $0 < s < t < 1$. If $g([0, s]) = g([0, t])$ then certainly g is not arcwise increasing. Otherwise there is an r in (s, t) such that $g(r) \in U := Y \setminus g([0, s])$. By continuity of g at r there is a closed neighbourhood $[a, b]$ of r such that $g([a, b]) \subseteq U$. Since $Y = g(I) = g(A) = g([0, s]) \cup g([t, 1])$, we see that g maps $[a, b]$ into $g([t, 1])$. Now $g([b, 1]) = g([a, 1])$ and g is not arcwise increasing.

For $(c) \Rightarrow (d)$ show that if g is not almost injective then it is not strongly irreducible.¹¹ So assume that $\{x : g^{-1}(g(x)) = \{x\}\}$ misses an open interval $(s, t) \subseteq I$. This means for all $x \in (s, t)$ there exists $y_x \neq x$ such that $g(x) = g(y_x)$. By the Baire Category Theorem, there is $n \in \mathbb{N}$ and $(s', t') \subsetneq (s, t)$ such that $X := \{x \in (s, t) : |x - y_x| \geq 1/n\}$ is dense in (s', t') . Without loss of generality, $|t' - s'| < 1/n$. But now $g(I \setminus (s', t')) = Y$, since $g(I \setminus (s', t'))$ is closed in Y and contains the set $g(X)$, which was dense in $g(s', t')$.

For $(d) \Rightarrow (a)$ suppose f is almost injective, and pick subarcs $A \subsetneq B$ in I . Then $B \setminus A$ contains a non-empty open interval which must meet the dense set $\{x : g^{-1}(g(x)) = \{x\}\}$ say in x' . But then $g(x') \in g(B) \setminus g(A)$, as required for arcwise increasing. \square

Turning to the case of maps from the circle, we deduce that an Eulerian map satisfies all of the following equivalent conditions.

THEOREM 14.2.3. *For a continuous surjection $g: S^1 \rightarrow X$ onto a Peano continuum X , the following are equivalent: (a) g is arcwise increasing; (b) g is hereditarily irreducible; (c) g is strongly irreducible; (d) g is almost injective; and (e) g is irreducible.*

If, additionally, X is a Peano graph, then the preceding are also equivalent to: (f) g is edge-wise Eulerian and $g^{-1}(\mathfrak{E}(X))$ is zero-dimensional in S^1 .

PROOF. The equivalence of (a) through (e) follows from Lemma 14.2.2 and the fact that for S^1 , every proper closed subset is contained in a proper subcontinuum, giving $(c) \Leftrightarrow (e)$. Now additionally assume X is a Peano graph.

$(c) \Rightarrow (f)$. Suppose g is strongly irreducible. By Lemma 14.2.1, g is edge-wise Eulerian. Suppose for a contradiction that $g^{-1}(\mathfrak{E}(X))$ is not zero-dimensional. Then there is a non-trivial interval $[a, b] \subseteq S^1$ such that $g([a, b]) \subseteq \mathfrak{E}(X)$. However, then $g(S^1 \setminus (a, b)) \supseteq \overline{\bigcup E(X)} = X$, contradicting that g is strongly irreducible.

$(f) \Rightarrow (d)$. For any non-trivial open interval $J \subseteq S^1$, we have $J \setminus g^{-1}(\mathfrak{E}(X))$ is non-empty, so contains a point x which is mapped under g onto an interior point of some edge of X . Since g is edge-wise Eulerian, x is a point of injectivity of g . Since J was arbitrary, g is almost injective. \square

As mentioned above, the converse to Lemma 14.2.1 is false, and we may *not* add ‘ g is edge-wise Eulerian’ to our list of equivalences, even when restricting to Peano graphs.

¹¹See [166, Theorem VIII.10.2] for a generalisation of this implication.

Since edge-wise Eulerian maps have, by definition, the geometrically natural property of an ‘Eulerian path’ of sweeping through every edge exactly once, why do we take strongly irreducible as the primary definition of Eulerian?

The answer is twofold. First, consider, for example, the Gromov compactification of a locally finite hyperbolic graph G with Gromov boundary ∂G . By property (f), an Eulerian map on G is not allowed to spend any non-trivial time in the boundary ∂G . Hence, Eulerian maps therefore satisfy the natural property that if a subpath of the Eulerian map in G ‘disappears’ in some direction $x \in \partial G$ towards infinity along some ray, then it must also return from that very direction x into the graph G .

Our second, equally important reason is that for Peano graphs, Eulerian maps – unlike edge-wise Eulerian maps – can essentially be characterised *purely combinatorially* in terms of a cyclic order and orientation of the edge set, as follows.

First, fix a Peano graph X and an Eulerian map $g: S^1 \rightarrow X$. Note that the edges, E , of X inherit from g a natural cyclic order. Of course the circle, $S^1 = \{(\cos(2\pi t), \sin(2\pi t)) : t \in [0, 1)\}$, has a natural cyclic order and (anticlockwise) orientation. Then any family of open intervals in the circle have an induced cyclic order (pick one point in each interval and use the sub-order). We have just seen that g is edge-wise Eulerian and $g^{-1}(\mathfrak{E}(X))$ is closed, nowhere dense. But this means that the edges, E , are in bijective correspondence with the family $\mathcal{U} = \{g^{-1}(e) : e \in E\}$ of open intervals in S^1 , which, we note, has dense union. Then E inherits a cyclic order from \mathcal{U} .

Second, it is also intuitively clear that, through the natural orientation on S^1 , any (edge-wise) Eulerian map on a Peano graph crosses each edge once in a certain direction, and so induces an orientation of every edge. We make this precise as follows. For any spaces A and B let $\mathcal{H}(A, B)$ be the (possibly empty) set of all homeomorphisms from A to B , and define $\mathcal{H}(A) = \mathcal{H}(A, A)$ to be the set of all autohomeomorphisms of A . Every autohomeomorphism of $(0, 1)$ (respectively S^1) either preserves or reverses the (cyclic) order. For $e \in E(X)$ define an equivalence relation, \sim_o , on $\mathcal{H}((0, 1), e)$ by $h_1 \sim_o h_2$ if and only if there is an order-preserving σ in $\mathcal{H}((0, 1))$ such that $h_2 = h_1 \circ \sigma$. Then $\mathcal{H}((0, 1), e)$ has two equivalence classes under \sim_o , corresponding to the two different directions for crossing e . Fix a bijection, o_e , between $\mathcal{H}((0, 1), e)/\sim_o$ and $\{\pm 1\}$. (So o_e randomly assigns a ‘positive’ (+1) and ‘negative’ (−1) direction to the edge e .) Now suppose we also have an Eulerian map, $g: S^1 \rightarrow X$. Fix an edge e . Fix an order-preserving bijection, τ , between $(0, 1)$ and $g^{-1}(e)$, and define $o_g^*(e)$ to be $[g \upharpoonright g^{-1}(e) \circ \tau]_{\sim_o}$. (Note that $o_g^*(e)$ is independent of the choice of τ .) This gives a function $o_g: E \rightarrow \{\pm 1\}$ via $o_g(e) = o_e(o_g^*(e))$, the orientation of e induced by g .

In summary: for a fixed Peano graph X with edge set $E = E(X)$ choose (randomly) a direction +1 or −1 for each edge, then for any edge-wise Eulerian map g derive combinatorial data of a cyclic order \leq_g on E and a function $o_g: E \rightarrow \{\pm 1\}$ so that g crosses the edges in the order given by \leq_g and in the direction given by o_g .

Let us say that another map $g': S^1 \rightarrow X$ is *cyclically equivalent* to g if and only if there is an order-preserving autohomeomorphism, ϱ say, of S^1 such that $g' = g \circ \varrho$. Then it can be shown that g and g' give the same combinatorial data – \leq_g isomorphic to $\leq_{g'}$, and $o_g = o_{g'}$ – if and only if they are cyclically equivalent.

Now we see how to get from combinatorial data to a function. Fix a Peano graph X with fixed direction for each edge. Let \leq be a cyclic order on the edges, $E = E(X)$, and o any function from E into $\{\pm 1\}$. Define $g_{\leq, o}$ a function from S^1 to X as follows.

First select $\mathcal{U} = \mathcal{U}_{\leq, o}$, a dense family of open intervals in S^1 , which – in the induced cyclic order – is isomorphic to (E, \leq) (it is well-known that every countable cyclic order can be realised in this fashion), say via $\varphi: \mathcal{U} \rightarrow E$. For each U in \mathcal{U} , from the randomly assigned direction, ± 1 , to the edge $\varphi(U)$ compared to the value of $o(\varphi(U))$ we get a \sim_o equivalence class in $H((0, 1), \varphi(U))$ – let g_U^* be any element of this class. Now select an order preserving bijection, τ between U and $(0, 1)$, and define $g_U = g_U^* \circ \tau$. Define $g_{\leq, o}$ to be g_U on each U in \mathcal{U} , and extend, if possible, to a (unique, if it exists) continuous map from S^1 to X (and otherwise extend randomly).

THEOREM 14.2.4. *If X is a Peano graph, with edges $E = E(X)$ and fixed direction for each edge, then the following condition on a continuous surjection $g: S^1 \rightarrow X$ is also equivalent to it being an Eulerian map:*

- (g) *there is a cyclic order \leq on E and a function $o: E \rightarrow \{\pm 1\}$ such that g is cyclically equivalent to $g_{\leq, o}$.*

PROOF. For $(f) \Rightarrow (g)$, let g be as in (f) . Let $\leq = \leq_g$ and $o = o_g$. Let $\mathcal{U}_g = \{g^{-1}(e): e \in E\}$ be as above, with the induced cyclic order. Let $\mathcal{U} = \mathcal{U}_{\leq, o}$ be the dense family of open intervals used in the definition of $g_{\leq, o}$. It is well-known that since \mathcal{U} and \mathcal{U}_g are dense collections of open intervals which are order-isomorphic, there is an order-preserving autohomeomorphism $\varrho^* \in \mathcal{H}(S^1)$ inducing that order-isomorphism.

Now chasing the definitions, we see that the difference between g and $g_{\leq, o} \circ \varrho^*$ is caused by choosing the ‘wrong’ class representative on some (possibly, many) intervals U in \mathcal{U} . But we can modify ϱ^* to get ϱ which is still an order-preserving autohomeomorphism and which ‘corrects’ the mistakes, so $g = g_{\leq, o} \circ \varrho$, as required.

For $(g) \Rightarrow (f)$ note that a function cyclically equivalent to an Eulerian map is Eulerian. So suppose $g = g_{\leq, o}$, and $\mathcal{U} = \mathcal{U}_g = \mathcal{U}_{\leq, o}$. By construction, g is edge-wise Eulerian, and $g^{-1}(\mathfrak{G}(X)) = S^1 \setminus \bigcup \mathcal{U}$ is zero-dimensional, since \mathcal{U} is dense in S^1 . □

Finally, we note that Theorem 14.2.3(f) has the following interesting consequence: it says that if a Peano graph X is Eulerian via an Eulerian map g , then $X \cong S^1 / \approx$ is a quotient of the circle where \approx is the decomposition of S^1 into fibres $\{g^{-1}(x): x \in \mathfrak{G}(X)\}$ and points, [63, 3.2.11]. Turning this procedure around, we can engineer (open) Eulerian Peano graphs with prescribed ground spaces as follows:

THEOREM 14.2.5. *For any compact metrizable space Z there is a Peano graph X with $\mathfrak{G}(X) = Z$. Moreover, for all $x, y \in Z$, the space X can be chosen so that*

- (1) X is Eulerian, or
- (2) X is open Eulerian from x to y .

PROOF. Such a construction can be quickly achieved using the *adjunction space* construction, see [158, A.11.4] or [63, 2.4.12f]. Let Z be arbitrary. For (2), consider the Cantor middle third set $C \subseteq I$, and fix a surjection $h: C \rightarrow Z$ onto Z with $h(0) = x$ and $h(1) = y$ [121, 7.7]. Set $X = I \cup_h Z$, where $I \cup_h Z$ is the quotient of I given by the decomposition into fibres of h and points of $I \setminus C$. By [158, A.11.4], if $g: S^1 \rightarrow X$ denotes the quotient map, then $g \upharpoonright I \setminus C$ is a homeomorphism (onto the edge set of X) and $g(C)$ is homeomorphic to Z . Thus, $\mathfrak{G}(X) = Z$ and by Theorem 14.2.3(f), g is an open Eulerian map from x to y .¹²

For (1), add one further free arc $e = xy$ to the space X constructed so far. \square

14.2.3. Reduction to Peano Graphs. The main purpose of this section is to show that in order to prove the Eulerianity conjecture, it suffices to always restrict our attention to the case of Peano graphs, in other words, to Peano continua where the free arcs are dense. This will be done in Section 14.2.3.3. In preparation we introduce some background material on Peano continua, Bing's partition theory, and a technical result on almost injective maps from the circle in Section 14.2.3.2.

In Section 14.2.4 the reduction result is used to show the equivalence of Eulerianity and edge-wise Eulerianity, first in Peano graphs, and then in general Peano continua.

14.2.3.1. *Tools for Peano continua.* In the following we shall need Bing's notion of a *partition* of a Peano continuum – originally from [24, 25], but we use it in the form of [117].

DEFINITION 14.2.6 (ε -Peano covers and partitions). Let X be a Peano continuum. A *Peano cover* of X is a finite collection \mathcal{U} of Peano subcontinua of X such that \mathcal{U} covers X . A Peano cover consisting of regular closed Peano subcontinua additionally satisfying that $\text{int}(U)$ is connected and $\text{int}(U) \cap \text{int}(V) = \emptyset$ for all $U \neq V \in \mathcal{U}$ is called a *Peano partition*. If $\varepsilon > 0$, then a Peano cover (partition) \mathcal{U} is called an ε cover (partition) if $\text{mesh}(\mathcal{U}) \leq \varepsilon$.

THEOREM 14.2.7 (Bing's Partitioning Theorem, [24]). *Every Peano continuum admits a decreasing sequence, \mathcal{U}_n , of $1/n$ Peano partitions.*

14.2.3.2. *Controlling almost injective maps from the circle.* Harrold, in [92], showed that every Peano continuum without free arcs is the strongly irreducible (equivalently, almost injective) image of the circle, and so is Eulerian. We extend this result – and also one of Espinoza & Matsushashi, see [71] – so as to give more control of the map.

¹²For a more explicit construction, we refer the reader to the technique in [120, Lemma 2.2].

For this, we introduce the following notation. Let A and B be spaces. Denote by $\mathcal{C}(A, B)$ the set of all continuous maps from A to B . Let K and L be subsets of A and B , respectively. Write $\mathcal{S}(A, B; K, L)$ for all elements of $\mathcal{C}(A, B)$ taking K onto L , and abbreviate $\mathcal{S}(A, B; A, B)$ by $\mathcal{S}(A, B)$. If X is a Peano continuum, then both $\mathcal{C}(I, X)$ and $\mathcal{S}(I, X)$ endowed with the supremum metric d_∞ are (non-empty) complete metric spaces. If in addition K is closed, then $\mathcal{S}(I, X; K, L)$ is a closed subspace of $\mathcal{C}(I, X)$ and hence also a complete metric space under the sup-metric. For sets $T \subseteq I$ and $g \in \mathcal{S}(I, X)$, we put $\mathcal{S}(I, X, g, T) = \{h \in \mathcal{S}(I, X) : h \upharpoonright T = g \upharpoonright T\}$. Note that $\mathcal{S}(I, X, g, T)$ is a non-empty closed subspace of $\mathcal{S}(I, X)$, so it is itself a complete metric space under the sup-metric. Lastly, for $F \subseteq I$ and $\delta > 0$ we put

$$\mathcal{A}_{F,\delta}(I, X) = \{h \in \mathcal{S}(I, X) : h^{-1}(h(x)) \subseteq B_\delta(x) \text{ for each } x \in F\}$$

and

$$\mathcal{A}_F(I, X) = \bigcap_{n \in \mathbb{N}} \mathcal{A}_{F,1/n}(I, X) = \{h \in \mathcal{S}(I, X) : h^{-1}(h(x)) = \{x\} \text{ for each } x \in F\}.$$

LEMMA 14.2.8. *Let X be a non-trivial Peano continuum. For each $a \in I$ and $\delta > 0$, the set $\mathcal{A}_{\{a\},\delta}(I, X)$ is open in $\mathcal{S}(I, X)$.*

PROOF. This result is well-known, and was stated for example (though without proof) in [146, Lemma 2.3] and in [92]. We briefly sketch the argument.

We show that the complement of $\mathcal{A}_{\{a\},\delta}(I, X)$ is closed. Suppose that $\{g_n : n \in \mathbb{N}\}$ is a sequence of functions in the complement, so for each n there are $x_n, y_n \in I$ with $|x_n - y_n| \geq \delta$ and $g_n(x_n) = a = g_n(y_n)$, such that $g_n \rightarrow g$ uniformly. By moving to subsequences and relabeling, we may assume that $x_n \rightarrow x$ and $y_n \rightarrow y$. But then $|x - y| \geq \delta$ and $g(x) = a = g(y)$. Hence, $g \notin \mathcal{A}_{\{a\},\delta}(I, X)$, i.e. the complement is closed. \square

THEOREM 14.2.9. *Let X be a non-trivial Peano continuum. Let $T, T' \subseteq I$ and $g \in \mathcal{S}(I, X)$ such that*

- (1) $I = T \cup T'$,
- (2) T' is closed in I ,
- (3) $Q := g(T') \subseteq X$ is a Peano subcontinuum of X without free arcs, and
- (4) $Q \cap \text{int}(g(T)) = \emptyset$.

Then for each countable subset $F \subseteq I$ with

- (5) $F \cap \bar{T} = \emptyset$,

the set $\mathcal{S}(I, X, g, T) \cap \mathcal{S}(I, X; T', Q) \cap \mathcal{A}_F(I, X)$ is a dense G_δ -subset of $\mathcal{S}(I, X, g, T) \cap \mathcal{S}(I, X; T', Q) = \{h \in \mathcal{S}(I, X, g, T) : h(T') = g(T')\}$, and hence non-empty.

PROOF. As $\mathcal{S}(I, X, g, T) \cap \mathcal{S}(I, X; T', Q)$ is a closed, non-empty subspace of $\mathcal{S}(I, X)$ it is complete under the supremum metric. So the claim that $\mathcal{S}(I, X, g, T) \cap \mathcal{S}(I, X; T', Q) \cap$

$\mathcal{A}_F(I, X)$ is non-empty follows by the Baire Category Theorem once we show that it is a dense G_δ -subset of $\mathcal{S}(I, X, g, T) \cap \mathcal{S}(I, X; T', Q)$.

Since $\mathcal{A}_F(I, X) = \bigcap_{a \in F} \bigcap_{m \in \mathbb{N}} \mathcal{A}_{\{a\}, 1/m}(I, X)$, is a countable intersection of open (see Lemma 14.2.8) sets, it suffices to prove that for each $a \in F$ and each $m \in \mathbb{N}$, the set $\mathcal{A}_{\{a\}, 1/m}(I, X) \cap \mathcal{S}(I, X, g, T) \cap \mathcal{S}(I, X; T', Q)$ is a dense subset of $\mathcal{S}(I, X, g, T) \cap \mathcal{S}(I, X; T', Q)$.

So fix some $a \in F$ and $m \in \mathbb{N}$ and consider any map $k \in \mathcal{S}(I, X)$ such that k coincides with g on T , and $k(T') = Q$. Take any $\varepsilon > 0$. We have to find a map h in $\mathcal{A}_{\{a\}, 1/m}(I, X) \cap \mathcal{S}(I, X, g, T) \cap \mathcal{S}(I, X; T', Q)$ with $d_\infty(h, k) < \varepsilon$.

From $k(T) = g(T)$, $k(T') = g(T')$, and (3), (4) and (5), it is straightforward to find a $k' \in \mathcal{S}(I, X, g, T) \cap \mathcal{S}(I, X; T', Q)$ with $d_\infty(k', k) < \varepsilon/3$ and $k'(a) \notin k(T)$. Next, find a small Peano subcontinuum $P \subseteq X$ with $k'(a) \in \text{int}(P) \subseteq P \subseteq Q$ and $\text{diam}(P) < \varepsilon/3$ such that $k'^{-1}(P) \cap T = \emptyset$. After suitably reparameterising k' on $k'^{-1}(P)$ (so that it will be nowhere constant with value $k'(a)$) we obtain a $k'' \in \mathcal{S}(I, X, g, T) \cap \mathcal{S}(I, X; T', Q)$ such that: $d_\infty(k'', k') < \varepsilon/3$, $k''(a) = k'(a) \notin g(T) = k(T) = k'(T) = k''(T)$, and $k''^{-1}(k''(a))$ is nowhere dense in I .

Since X is Peano, there is a basis at $k''(a)$ consisting of Peano subcontinua, in other words, there is a nested sequence of connected, open subsets U_n , for $n \in \mathbb{N}$, such that: $P_n = \overline{U_n}$ is a Peano subcontinuum of X , $P_{n+1} \subseteq U_n$ for all $n \in \mathbb{N}$, $\bigcap_{n \in \mathbb{N}} U_n = \bigcap_{n \in \mathbb{N}} P_n = \{k''(a)\}$, $P_0 \subseteq P$, and $k''^{-1}(U_0) \cap T = \emptyset$.

We now claim that for some n , the compact set $k''^{-1}(P_{n+1})$ is covered by finitely many connected components $(a_1^n, b_1^n), \dots, (a_{N(n)}^n, b_{N(n)}^n)$ of the open set $k''^{-1}(U_n)$ such that $|b_i^n - a_i^n| < 1/m$ for all $1 \leq i \leq N(n)$. Indeed, if not, then by König's Infinity Lemma [54, Lemma 8.1.2], there is a choice of intervals $(a_{j(n)}^n, b_{j(n)}^n)$ such that: $|b_{j(n)}^n - a_{j(n)}^n| \geq 1/m$, and $(a_{j(n+1)}^{n+1}, b_{j(n+1)}^{n+1}) \subseteq (a_{j(n)}^n, b_{j(n)}^n)$ for all $n \in \mathbb{N}$. But then $(a, b) = \bigcap_{n \in \mathbb{N}} (a_{j(n)}^n, b_{j(n)}^n)$ is an interval of length at least $1/m$ with $(a, b) = \bigcap_{n \in \mathbb{N}} (a_{j(n)}^n, b_{j(n)}^n) \subseteq \bigcap_{n \in \mathbb{N}} k''^{-1}(U_n) = k''^{-1}(k''(a))$ contradicting the fact that $k''^{-1}(k''(a))$ is nowhere dense in I .

So let us fix an $n \in \mathbb{N}$ as in the claim and consider $P_{n+1} \subseteq U_n \subseteq P_n$. Without loss of generality, assume $a \in (a_{N(n)}^n, b_{N(n)}^n)$. Pick arcs $\alpha_i: [a_i^n, b_i^n] \rightarrow P_n$ for $1 \leq i < N(n)$ from $k''(a_i^n)$ to $k''(b_i^n)$ inside P_n , and note that since U_{n+1} contains no free arcs by (3), the space $\bigcup \alpha_i$ is nowhere dense in U_{n+1} . In particular, there is a point $x \in U_{n+1}$ which is not yet covered by any of the α_i . Using the Hahn-Mazurkiewicz Theorem, pick a space filling curve $\alpha_{N(n)}: [a_{N(n)}^n, b_{N(n)}^n] \rightarrow P_n$ from $k''(a_{N(n)}^n)$ to $k''(b_{N(n)}^n)$, which we may parameterise such that $\alpha_{N(n)}(a) = x$.

Finally, the map h obtained from k'' by replacing each $k'' \upharpoonright [a_i^n, b_i^n]$ with α_i for $i \in [N(n)]$ is as desired. Clearly, h is onto by construction, and $h^{-1}(h(a)) = h^{-1}(x) \subseteq [a_{N(n)}^n, b_{N(n)}^n]$, so has diameter $< 1/m$ as desired. Further, k'' and h differ only within P_n , and so $d_\infty(h, k'') \leq \text{diam}(P_n) \leq \text{diam}(P_0) < \varepsilon/3$. Next, since $k''^{-1}(U_0) \cap T = \emptyset$, we have $h \upharpoonright T =$

$k'' \upharpoonright T = k \upharpoonright T$ and $h(T') = k''(T') = k(T')$. Finally, we have

$$d_\infty(h, k) \leq d_\infty(h, k'') + d_\infty(k'', k') + d_\infty(k', k) < \varepsilon/3 + \varepsilon/3 + \varepsilon/3 = \varepsilon$$

and so we have found our surjection $h \in \mathcal{A}_{\{a\}, 1/m}(I, X) \cap \mathcal{S}(I, X, g, T) \cap \mathcal{S}(I, X; T', Q)$ with $d_\infty(h, k) < \varepsilon$, completing the proof. \square

COROLLARY 14.2.10. *Let X be a non-trivial Peano continuum without free arcs. Let $T \subseteq I$ be nowhere dense, and let $g \in \mathcal{S}(I, X)$ such that $g(T)$ is nowhere dense in X . Then there is an almost injective map $h: I \rightarrow X$ with $h \upharpoonright T = g \upharpoonright T$.*

PROOF. As T is nowhere dense, we can find a dense countable subset $F \subseteq I$ with $F \cap \overline{T} = \emptyset$. Since $g(T)$ is nowhere dense by hypothesis, applying Theorem 14.2.9 with $T' = S^1$, we obtain an almost injective map h with $h \upharpoonright T = g \upharpoonright T$. \square

REMARK 14.2.11. *All the results above on almost-injective maps from the closed unit interval, I , extend naturally (with the obvious notational changes) to maps from the circle, S^1 . To see this, note that maps $\hat{g}: S^1 \rightarrow X$ naturally correspond to maps $g: I \rightarrow X$ such that $g(0) = g(1)$ and in applying the results, always add 0 and 1 to T .*

14.2.3.3. The reduction result. We now show we can reduce the general case the Eulerianity conjecture (for Peano continua, possibly with some free arcs) to the special case where the free arcs are dense, in other words, to the case of Peano graphs.

Indeed, let X be a Peano continuum with free arcs indexed by E . Define $X' = X \cup L$ to be the space obtained by attaching a zero-sequence of loops, L , to points in a countable dense subset of the part $X \setminus \overline{E}$ of the ground space where the free arcs are not dense. Then X' is a Peano graph by Lemma 14.1.13. It is immediate that X' satisfies the even-cut condition if and only if X does. And the next theorem says that X' is Eulerian if and only if X is Eulerian, and so, if the Eulerianity Conjecture holds for X' , then it holds for X .

THEOREM 14.2.12 (Reduction Result). *Let X be a Peano continuum, and D a countable dense subset of $X \setminus \overline{E}$. Define a Peano graph X' by attaching a zero-sequence of loops $L = \{\ell_d: d \in D\}$ to points in D .*

Then X' is Eulerian if and only if X is Eulerian.

PROOF. Enumerate $D = \{d_n: n \in \mathbb{N}\}$. First, if X is a Peano continuum, then so is $X' = X \cup \bigcup_{n \in \mathbb{N}} \overline{\ell_{d_n}}$ by Lemma 14.1.13. Moreover, if X is Eulerian, then so is X' , as any almost injective map $g: S^1 \rightarrow X$ lifts to an almost injective map $g': S^1 \rightarrow X'$ by incorporating the loops ℓ_{d_n} into g using the results from Section 14.1.3.2.

Conversely, assuming that X' is Eulerian, we show X is also Eulerian. To this end, fix an almost injective map $g: S^1 \rightarrow X'$. Pick a sequence of decreasing $1/n$ -Peano partitions \mathcal{P}_n for X (see Definition 14.2.6 and Theorem 14.2.7). Let $\mathcal{P}'_{n+1} \subseteq \mathcal{P}_{n+1}$ be the collection of all $P \in \mathcal{P}_{n+1}$ such that P is disjoint from \overline{E} , but the unique Q in \mathcal{P}_n containing P

meets \overline{E} . Let $\{P_j : j \in \mathbb{N}\}$ be an enumeration of $\bigcup_{n \in \mathbb{N}} \mathcal{P}'_n$ such that $\varepsilon_j = \text{diam}(P_j)$ is monotonically decreasing to 0 as $j \rightarrow \infty$. Note that $\text{int}(P_i) \cap P_j = \emptyset$ whenever $i \neq j$ and that $D \subseteq \bigcup_{j \in \mathbb{N}} P_j$. Indeed, for the last statement note that every $d \in D$ by construction has positive distance from \overline{E} , so when the mesh of \mathcal{P}_n is smaller than that distance, there is $P \in \mathcal{P}_n$ such that $d \in P$ and $P \cap \overline{E} = \emptyset$. Finally, observe that each P_j is a Peano subcontinuum of X without free arcs, and so may play the rôle of the set Q in item (3) of the previous theorem.

We now define a countable dense set $F \subseteq S^1$ and a sequence of continuous surjections $g_i : S^1 \rightarrow X_i$ where $X_i = X' \setminus \{\ell_d : d \in \bigcup_{j < i} P_j\}$ such that for all $i \in \mathbb{N}$

- the set F witnesses that g_i is almost injective,
- $g_i(F) \cap \partial P_j = \emptyset$ for all $j \in \mathbb{N}$,
- g_{i+1} agrees with g_i on $S^1 \setminus \text{int}(g_i^{-1}(P_i[X_i]))$, and
- $g_{i+1}(g_i^{-1}(P_i[X_i])) = P_i$.

[Where for a subcontinuum $P \subseteq X$ we denote by $P[X_i] = P \cup \{\ell_d \in E(X_i) : d \in P\}$, in other words, P with all loops from L that are still present in the space X_i .]

Once the construction is complete, we claim that $h = \lim g_i$ is the desired, almost injective surjection from S^1 onto $X = \bigcap_{i \in \mathbb{N}} X_i$. Indeed, as we change our function value for each point of S^1 at most once, and do so inside the target sets $P_i[X_i]$ which are decreasing in size, the sequence is Cauchy and converges to a surjection onto X . Moreover, since the sequence $(g_i)_{i \in \mathbb{N}}$ is pointwise eventually constant, it is immediate from the first bullet point that F witnesses that also h is almost injective.

It remains to complete the construction. Define $g_1 = g$ and let $F \subseteq g_1^{-1}(E(X'))$ be a countable dense subset of S^1 witnessing that g is almost injective (possible by Theorem 14.2.3(f)). Next, suppose recursively that g_i has already been defined. Consider $T'_i := g_i^{-1}(P_i[X_i]) \subseteq S^1$, a closed, compact subspace with non-empty interior (as a positive amount of time is needed to cover the loops ℓ_d with $d \in \text{int}(P_i)$). Let $\{[a_m, b_m] : m \in \mathbb{N}\}$ be an enumeration of the maximal non-trivial intervals contained in $g_i^{-1}(P_i[X_i])$. Then clearly, $g_i(a_m), g_i(b_m) \in \partial P_i = \partial P_i[X_i]$. Consider the natural quotient map $q_i : X_i \rightarrow X_{i+1}$ which collapses every loop ℓ_d in $P_i[X_i]$ onto its base point d . Let $g'_i = q_i \circ g_i : S^1 \rightarrow X_{i+1}$. We then may apply Theorem 14.2.9 for maps on S^1 (see Remark 14.2.11) to the map $g'_i \in \mathcal{S}(S^1, X_{i+1})$ in order to find a surjection $g_{i+1} \in \mathcal{S}(S^1, X_{i+1}, g'_i, T_i) \cap \mathcal{S}(S^1, X_{i+1}, T'_i, Q_i) \cap \mathcal{A}_{F_i}(S^1, X_{i+1})$ where $T_i = S^1 \setminus \bigcup_{m \in \mathbb{N}} (a_m, b_m)$, $T'_i = g_i^{-1}(P_i[X_i])$, $Q_i = g'_i(T'_i) = P_i$ and $F_i = \bigcup_{m \in \mathbb{N}} (a_m, b_m) \cap F$.

We claim that g_{i+1} is as desired. That it satisfies the properties of the third and fourth bullet points follows from the fact that it is an element of $\mathcal{S}(S^1, X_{i+1}, g'_i, T_i)$ and of $\mathcal{S}(S^1, X_{i+1}, T'_i, Q_i)$ respectively. For the first bullet point, we verify that all points of F are points of injectivity of g_{i+1} . Since $g_{i+1} \in \mathcal{A}_{F_i}(S^1, X_{i+1})$, this is clear for points of $F_i \subseteq F$. Suppose for a contradiction that some $x \in F \setminus F_i$ is no longer a point of injectivity for

g_{i+1} . Since $g_{i+1} \upharpoonright T_i = g'_i \upharpoonright T_i = g_i \upharpoonright T_i$ and x was a point of injectivity for g_i , it must be the case that there is $x' \in (a_m, b_m)$ for some $m \in \mathbb{N}$ such that $g_{i+1}(x) = g_{i+1}(x')$. This, however, implies that $g_{i+1}(x) \in \partial P_i$, but since $g_{i+1}(x) = g_i(x)$, this contradicts the property of the second bullet for g_i . Lastly, it remains to verify that $g_{i+1}(F) \cap \partial P_j = \emptyset$ for all $j \in \mathbb{N}$. This is clear for points in $F \setminus F_i$ as their values are unchanged, and follows for points in F_i from the fact that $g_{i+1} \in \mathcal{A}_{F_i}(S^1, X_{i+1}) \cap \mathcal{S}(S^1, X_{i+1}, g_{i+1}, T_i)$ readily implies that $g_{i+1}(F_i) \subseteq \text{int}(P_i)$. \square

14.2.4. Equivalence of Eulerianity and Edge-Wise Eulerianity. Recall we have defined a Peano continuum X to be *edge-wise Eulerian* if there is a surjection $g: S^1 \rightarrow X$ such that g sweeps through every free arc of X precisely once, and we have seen that every Eulerian continuum is edge-wise Eulerian. We now establish the converse, the proof of which establishes the assertion for Peano graphs first, and then, utilizing the reduction result, for general Peano continua.

THEOREM 14.2.13. *A space is Eulerian if and only if it is edge-wise Eulerian.*

PROOF. By Lemma 14.2.1, only the backwards implication requires proof. We first prove this implication for Peano graphs, in other words, when the edges are dense.

The circle has a natural cyclic order where $x \leq y \leq z$ if we visit y as we travel anticlockwise around the circle starting at x and ending at z . Then we say a surjection $g: S^1 \rightarrow X$ is *edge-wise monotone* if for every edge e of X its inverse image, $g^{-1}(e)$ is a single open interval in S^1 (so g crosses e exactly once) and, after orienting e appropriately, g is monotone (if $x \leq y \leq z$ in $g^{-1}(e)$ then $g(x) \leq g(y) \leq g(z)$ in e) from $g^{-1}(e)$ and e (so g may pause when crossing e , but does not backtrack). Clearly edge-wise Eulerian maps are edge-wise monotone, but observe, also, that if g is edge-wise monotone then, as explained in Lemma 14.1.12(a), we can eliminate the waiting times to get an edge-wise Eulerian map with nowhere dense fibres. In any case, it suffices to show that if X has an edge-wise Eulerian map with nowhere dense fibres then it has an Eulerian map. We do this in two steps.

First of all, let us write $\mathcal{M}(S^1, X) \subseteq \mathcal{S}(S^1, X)$ for the space of edge-wise monotone maps with the sup-metric. We will show that this is a closed subspace, and hence a G_δ set. Let us write $\mathcal{W}(S^1, X) \subseteq \mathcal{S}(S^1, X)$ for the space of edge-wise Eulerian maps which have all fibres nowhere dense, with the sup-metric. Fix a countable subset D of S^1 . Noting that a map g from S^1 onto X has nowhere dense fibres if and only if for every distinct d and d' from D and every x strictly between them ($d < x < d'$) either $g(x) \neq g(d)$ or $g(x) \neq g(d')$, we see that $\mathcal{W}(S^1, X) = \mathcal{M}(S^1, X) \cap \bigcap_{d \neq d' \in D} U_{d,d'}$ where $U_{d,d'} = \bigcup_{d < x < d'} \{g \in \mathcal{S}(S^1, X) : g(d) \neq g(x) \text{ or } g(d') \neq g(x)\}$ is an open set. Thus $\mathcal{W}(S^1, X)$ is a non-empty G_δ subset of $\mathcal{S}(S^1, X)$, which is complete, and so itself is complete, [63, 4.3.23]. Hence – by the Baire Category Theorem – dense G_δ subsets of $\mathcal{W}(S^1, X)$ are non-empty.

Now to show that $\mathcal{M}(S^1, X)$ is indeed closed, suppose we have a sequence $(g_n : n \in \mathbb{N})$ in $\mathcal{M}(S^1, X)$ and $g \in \mathcal{S}(S^1, X)$ with $d_\infty(g_n, g) \rightarrow 0$. We need to show that $g \in \mathcal{M}(S^1, X)$, which in turn means we need to show that for every edge $e \in E(X)$, we have g is monotone on the interval $g^{-1}(e)$. Fix an edge e . It can be oriented in one of two ways. Since the g_n 's converge uniformly to g , and every g_n is monotone on the interval $g_n^{-1}(e)$ for some orientation of e , eventually the orientations must all be the same. So without loss of generality, let us assume e is oriented the same way for all n in \mathbb{N} . Take any x, z in $g^{-1}(e)$ and any y between them, $x \leq y \leq z$. Then again by uniform convergence of the g_n 's to g and the intermediate value theorem, if g does not respect the order, so we do not have $g(x) \leq g(y) \leq g(z)$, then for some large enough n , g_n will also not respect the order - contradicting g_n being edge-wise monotone. Now it follows both that y is in $g^{-1}(e)$, which is therefore an interval, and that g is monotone on that interval. Hence, $g \in \mathcal{M}(S^1, X)$ and we have established that $\mathcal{M}(S^1, X)$ is closed.

The second step (for X a Peano graph) is to show that for every a in S^1 and $\delta > 0$, the set $\mathcal{A}_{\{a\}, \delta}(S^1, X) \cap \mathcal{W}(S^1, X) = \{g \in \mathcal{W}(S^1, X) : g^{-1}(g(a)) \subseteq B_\delta(a)\}$ (where $\mathcal{A}_{\{a\}, \delta}(S^1, X)$ is as defined in Section 14.2.3.2) is dense in $\mathcal{W}(S^1, X)$. Since it is open, see Lemma 14.2.8, taking any countable dense subset $F \subseteq S^1$, by Baire Category, there is a function in $\bigcap_{n \in \mathbb{N}} \bigcap_{a \in F} \mathcal{A}_{\{a\}, 1/n}(S^1, X) \cap \mathcal{W}(S^1, X)$. This function is then almost injective, so Eulerian by Theorem 14.2.3, as desired.

So it remains to check for density. For this, let $g \in \mathcal{W}(S^1, X)$, a in S^1 and $\varepsilon > 0$ be arbitrary. Our task is to find $h \in \mathcal{A}_{\{a\}, \delta}(S^1, X) \cap \mathcal{W}(S^1, X)$ with $d_\infty(g, h) < \varepsilon$. Since X is Peano, there is a basis at $g(a)$ consisting of Peano subcontinua, so in particular there are connected, open subsets U_0 and U_1 such that: $\text{diam}(U_0) < \varepsilon/2$, $P_1 = \overline{U_1}$ is a Peano subcontinuum of X , and $a \in U_1 \subseteq P_1 \subseteq U_0$. Clearly, the compact set $g^{-1}(P_1)$ is covered by finitely many connected components $(a_1, b_1), \dots, (a_k, b_k)$ of the open set $g^{-1}(U_0)$. Relabelling if necessary, assume $a \in (a_1, b_1)$. Let us write g_i for $g \upharpoonright [a_i, b_i]$ where $1 \leq i \leq k$. We deal with two cases depending on whether or not g_1 crosses an edge of X .
Case 1. Suppose g_1 crosses an edge of X . Then we can reparameterise g_1 to get g'_1 so that $g'_1(a)$ is in e . Now define the map h on the circle to be g'_1 on $[a_1, b_1]$ and g elsewhere. Then h is as desired, indeed $d_\infty(g, h) < \varepsilon/2$, $h^{-1}(h(a)) = \{a\}$ and as g is never constant on a non-trivial interval, by construction of h , it too has nowhere dense fibres.

Case 2. Otherwise, by the boundary bumping lemma we know that the image, $\text{ran } g_1$, of g_1 is a non-trivial subcontinuum of $\mathfrak{G}(X) \cap U_0$. In particular, let us fix distinct points $x_1, \dots, x_{2k-1} \in \text{ran } g_1$, and – this is where we assume X is a Peano graph, and the edges are dense – for each of them a sequence of edges $e_n^i \in U_1$ such that $e_n^i \rightarrow x_i$ as $n \rightarrow \infty$. Now, as g is edge-wise Eulerian, each edge e_n^i must be crossed by precisely one function g_j for $2 \leq j \leq k$. By the pigeon hole principle we see that for each i , at least one function $g_{j(i)}$ crosses infinitely many of $\{e_n^i : n \in \mathbb{N}\}$. Moreover, since we have $2k - 1 = 2(k - 1) + 1$ many points x_i , but only $k - 1$ functions, by the pigeon hole principle again, there is

one function, say (relabelling if necessary) g_2 , that is used at least three times, say (after relabelling) for x_1, x_2, x_3 .

Now by construction, there are points $y_1, y_2, y_3 \in (a_2, b_2)$ and $(z_m^i : m \in \mathbb{N})$ for $i \in [3]$ such that such: $g_2(y_i) = x_i$, $g_2(z_m^i) \in e_{n_m}^i$ and $z_m^i \rightarrow y_i$ as $m \rightarrow \infty$.

Relabelling if necessary, let us assume that $y_1 < y_2 < y_3$, and further, for all $m \in \mathbb{N}$ we have $y_1 < z_m^2 < y_2$. This means, in particular, that $g_2 \upharpoonright [y_1, y_2]$ starts and ends in $\text{ran}(g_1)$ and crosses an edge. Pick $x \leq y \in [a_1, b_1]$ such that $g_1(x) = x_1$ and $g_1(y) = x_2$. Then define g' on S^1 to be g except swap $g_1 \upharpoonright [x, y]$ with $g_2 \upharpoonright [y_1, y_2]$. Clearly g' is edge-wise Eulerian, has nowhere dense fibres (by construction, given that g has the same property) and has distance $< \epsilon/2$ from g . Now apply the argument of *Case 1* to g' to get the map h . This h is as required: $d_\infty(g, h) \leq d_\infty(g, g') + d_\infty(g', h) < \epsilon/2 + \epsilon/2 = \epsilon$, and h is in $\mathcal{A}_{\{a\}, \delta}(S^1, X) \cap \mathcal{W}(S^1, X)$.

To complete the proof, consider now an arbitrary Peano continuum X which is edge-wise Eulerian. Let $g: S^1 \rightarrow X$ be a surjection that sweeps through every free arc of X precisely once. Let X' be the Peano continuum where we attached a dense zero-sequence of loops of the ground space of X , as in Theorem 14.2.12. Then X' is a Peano graph, and g clearly lifts to a surjection $g': S^1 \rightarrow X'$ that sweeps through every free arc of X' precisely once by Lemma 14.1.14. Hence X' is edge-wise Eulerian, and so Eulerian by the first part of this proof. By Theorem 14.2.12, it follows that X is Eulerian, as well. \square

Finally, we conclude this chapter with a further reduction result reducing to the case where we do not have loops.

THEOREM 14.2.14 (Loopless reduction result). *It suffices to prove the Eulerianity conjecture for Peano graphs without loops. More precisely, Conjecture 14.1.2 holds for a Peano continuum X provided it holds for all loopless Peano graphs Z with $\mathfrak{G}(Z) = \mathfrak{G}(X)$.*

PROOF. By the first reduction result, it suffices to consider Peano graphs X only. Since the Eulerianity conjecture holds for spaces X where $\mathfrak{G}(X)$ is a singleton (in which case X is either a circle, a wedge of finitely many circles, or a Hawaiian earring), we may assume that $|\mathfrak{G}(X)| > 1$. So consider such a Peano graph X with $|\mathfrak{G}(X)| > 1$ satisfying the even-cut condition, and let $L = \{e \in E(X) : e(0) = e(1)\} \subseteq E(X)$ be the collection of loops in X . Then $Y = X - L$ is a Peano continuum, but may no longer be a Peano graph. Let $U = \text{int}\left(\overline{\bigcup L}\right) \cap \mathfrak{G}(X)$. If $U = \emptyset$, set $F := \emptyset$. Otherwise, let $D = \{d_1, d_2, \dots\}$ be a countable dense subset of U . Since $X \neq S^1$, no d_n is isolated in $\mathfrak{G}(X)$. For each d_n consider a small Peano continuum neighbourhood $P_n \subseteq X$ with $d_n \in \text{int}(P_n) \subseteq P_n \subseteq \text{int}\left(\overline{\bigcup L}\right)$. Then $P_n - L \subseteq \mathfrak{G}(X)$ is a non-trivial Peano continuum. Hence, there exists a small non-trivial arc $\alpha_n \subseteq \mathfrak{G}(X)$ from d_n to say x_n of diameter $\leq 2^{-n}$. Add a new edge / free arc f_n from d_n to x_n of length $\text{dist}(d_n, x_n) \leq 2^{-n}$, and set $F = \{f_n : n \in \mathbb{N}\}$. Then $Z = Y + F$ is a Peano graph with $\mathfrak{G}(Z) = \mathfrak{G}(X)$. Moreover, Z inherits the even-cut condition from X ,

since loops in L and edges in F each have both their end points in the same component of $\mathfrak{G}(X) = \mathfrak{G}(Z)$, and hence to not appear in any finite edge cut. By assumption, there exists an edge-wise Eulerian map g_Z for Z . This turns naturally into an edge-wise Eulerian map g_Y for Y , by replacing every newly added edge f_n by α_n . But using Lemma 14.1.14, we may incorporate the zero-sequence of loops in L into g_Y in order to obtain an edge-wise Eulerian map g_X for X . By Theorem 14.2.13, it follows that X is Eulerian. \square

14.3. Approximating by Eulerian decompositions

From the introduction we know that the key task facing us is the construction of Eulerian maps for Peano continua with the even-cut condition. From the last chapter, we know that we may restrict our attention to constructing edge-wise Eulerian maps. The goal for this chapter is then to provide one such construction. In order to do so, we introduce a versatile framework which we call ‘approximating sequences of Eulerian decompositions’, and then show that these can indeed be used to give an edge-wise Eulerian map, thus completing the proof $(ii) \Leftrightarrow (iii)$ announced in Theorem 14.1.1. The implication $(ii) \Rightarrow (iii)$ is proved in Theorem 14.3.6 and $(iii) \Rightarrow (ii)$ is proved in Theorem 14.3.12.

The idea behind this framework of Eulerian decompositions lies in the observation that any edge-wise Eulerian map induces a countable cyclic order on the edge set $E(X)$ of our Peano continuum X . As in the case of graph-like spaces [70], we want to approximate such a cyclic order on a finitary version of X , and then choose a sequence of compatible approximations that ‘converge’ to the desired cyclic order on X . In this chapter, we formalise this idea. We describe what we understand about finite approximations and lay down a set of rules that these have to satisfy in order to make the ideas of ‘compatible’ and ‘converging’ mathematically sound, and then state and prove our main mapping result, Theorem 14.3.12, for constructing edge-wise Eulerian maps.

14.3.1. Eulerian Decompositions. An important tool in structural graph theory is the notion of a *tree-decomposition*, due to Halin [88], and rediscovered and made widely known by Robertson and Seymour in their graph-minors project [135]. Roughly, a tree decomposition (T, τ) of a graph G consists of a tree T and a map τ such that $\tau(t)$ is a subgraph of G for every $t \in V(T)$, such that the various subgraphs (‘parts’) $\{\tau(t) : t \in V(T)\}$ form a cover of the graph G whose elements are roughly arranged like T , see also [54, §12.3].

In analogy, we will now consider Eulerian decompositions: covers of a Peano continuum X by finitely many parts which are arranged roughly like an Eulerian graph.

14.3.1.1. Setup and definitions.

DEFINITION 14.3.1. Let X be a Peano continuum. A subspace $Y \subseteq X$ is called *standard* if Y contains all edges of X it intersects.

Recall that for an edge e of a finite multi-graph or a Peano continuum, we write $e(0)$ and $e(1)$ for the two end vertices of e (if e is a loop, then $e(0) = e(1)$), see Lemma 14.1.11.

DEFINITION 14.3.2 (Eulerian decomposition). Let X be a Peano continuum, G be a finite multi-graph with bipartitioned edge set $E(G) = F \sqcup D$, and η be a map with domain $V(G) \cup E(G)$ such that

- (E1) $\eta(v)$ is a non-empty standard Peano subcontinuum of X for all $v \in V(G)$,
- (E2) $\eta(f) \in E(X)$ for all $f \in F$, and
- (E3) $\eta(d) \subseteq \mathfrak{G}(X)$ is a (possibly trivial) arc for all $d \in D$.

The pair (G, η) is called a *decomposition*¹³ of X if it satisfies the following four conditions:

- (E4) the family $\{\eta(x) : x \in V \cup F\}$ forms a cover of X ,
- (E5) the elements of $\{\eta(x) : x \in V \cup F\}$ are pairwise $E(X)$ -edge-disjoint,¹⁴
- (E6) $(\eta(f))(j) \in \eta(f(j))$ for all $f \in F$ and $j \in \{0, 1\}$, and
- (E7) $(\eta(d))(j) \in \eta(d(j))$ for all $d \in D$ and $j \in \{0, 1\}$.

The *width* of a decomposition is $w(G, \eta) := \max \{\text{diam}(\eta(v)) : v \in V\}$. The edges in F are also called *real* or *displayed* edges, and the edges in D are the *dummy* edges of G . The elements $\{\eta(v) : v \in V\}$ are called *tiles* of the decomposition. A decomposition (G, η) where G is Eulerian, is called an *Eulerian decomposition* of X .

Dummy edges d between vertices v, w of G represent the possibility of moving from tile $\eta(v)$ to $\eta(w)$ through a common point in their overlap (if $\eta(d)$ is a singleton) or through an arc contained in the ground space of X (if $\eta(d)$ is a non-trivial arc). As an illustration, consider two Eulerian decompositions of the hyperbolic 4-regular tree X .




FIGURE 14.5. Two Eulerian decompositions (G, η) and (G', η') for X with tiles in pink and black (single vertices), displayed edges in blue, dummy edges $\eta(d_i) = \{\delta_i\} = \eta'(d_i)$ in red, and $\eta(v) = \{x\} = \eta'(v_i)$.

¹³Note that due to (E2) and (E3), the information $E(G) = F \sqcup D$ is encoded in η .

¹⁴This implies that $\eta \upharpoonright F$ is injective; however, for distinct vertices v and w of G , $\eta(v) = \eta(w)$ could be the same tile, which must then be contained in the ground space. Note also that $\eta(v)$ could contain free arcs which are not free in X . These don't play a role for the requirement of edge-disjoint.

Recall that an *edge-contraction* is the combinatorial analogue of collapsing the closure of an edge in a topological graph to a single point. Formally, given an edge $e = xy$ in a multi-graph $G = (V, E)$ (with parallel edges and loops allowed), the contraction G/e is the graph with vertex set $V \setminus \{x, y\} \sqcup \{v_e\}$ and edge set $E \setminus \{e\}$, and every edge formally incident with x or y of G is now incident with v_e . Note that all edges parallel to e are now loops in G/e . If e was a loop in G , then $G/e = G - e$. The contraction of more than one edge is denoted by $G/\langle e_1, \dots, e_k \rangle$. The order in which we contract edges does not matter. Any such graph G' which can be obtained by a sequence of contractions from G is called a *contraction minor* of G , denoted by $G' \preceq G$.

LEMMA 14.3.3 (Contractions on Eulerian decompositions.). *Suppose $\mathcal{D} = (G, \eta)$ is an [Eulerian] decomposition of X with edge partition $E = E(G) = F \sqcup D$. Then for an arbitrary edge $e = xy \in E$, there is an [Eulerian] decomposition $\mathcal{D}/e := (G', \eta')$ where $G' = G/e$, $E' = E - e$ with induced partition $F' \sqcup D'$, and the function η' given by*

- (C1) $\eta'(v_e) = \eta(x) \cup \eta(e) \cup \eta(y)$,
- (C2) $\eta'(v) = \eta(v)$ for all $v \neq v_e$, and
- (C3) $\eta'(e') = \eta(e')$ for all $e' \in E'$.

PROOF. By property (E6) and (E7) for \mathcal{D} (depending on whether $e \in F$ or $e \in D$ respectively), we have that $\eta'(v_e)$ is a standard subcontinuum of X . The remaining properties are easily verified.

Finally, it is clear that if G is Eulerian, then so is G' . □

DEFINITION 14.3.4. For two decompositions $\mathcal{D}_1 = (G_1, \eta_1)$ and $\mathcal{D}_2 = (G_2, \eta_2)$ of X , we say that \mathcal{D}_2 *extends* \mathcal{D}_1 , in symbols $\mathcal{D}_1 \preceq \mathcal{D}_2$, if there is a sequence of edges $e_1, \dots, e_k \in E(G_2)$ such that $\mathcal{D}_1 = \mathcal{D}_2/\langle e_1, \dots, e_k \rangle$.

In particular, $\mathcal{D}_1 \preceq \mathcal{D}_2$ implies that $G_1 \preceq G_2$, and conversely, every contraction minor $G_2/\langle e_1, \dots, e_k \rangle$ gives rise to a corresponding Eulerian decomposition which is extended by G_2 . For illustration, consider the following decompositions of the hyperbolic tree X .

DEFINITION 14.3.5. A sequence of [Eulerian] decompositions $(\mathcal{D}_n : n \in \mathbb{N})$ for a Peano continuum X is called an *approximating sequence of [Eulerian] decompositions* for X , if

- (A1) $\mathcal{D}_n \preceq \mathcal{D}_{n+1}$ for all $n \in \mathbb{N}$, and
- (A2) $w(\mathcal{D}_n) \rightarrow 0$ as $n \rightarrow \infty$.

14.3.1.2. *From Eulerian maps to Eulerian decompositions.* One motivation behind the definition of an Eulerian decomposition is they can be generated from every (edge-wise) Eulerian map $g: S^1 \rightarrow X$. In fact, any such map yields a surprising simple approximating sequence as follows:

THEOREM 14.3.6. *Every edge-wise Eulerian space admits an approximating sequence $((G_n, \eta_n) : n \in \mathbb{N})$ of Eulerian decompositions, where each G_n is a cycle of length n .*




FIGURE 14.6. Eulerian decompositions $(G_1, \eta_1) \preccurlyeq (G_2, \eta_2)$ with dummy edges satisfying $\eta_1(d_i) = \delta_i$ for $i \in [2]$ and $\eta_2(d_i) = \delta_i$ for $i \in [6]$. Note that $G_1 \preccurlyeq G_2$ by contracting all edges inside the dotted subgraphs of G_2 .

PROOF. Suppose that $g: S^1 \rightarrow X$ is an edge-wise Eulerian map. Then the preimages $I_e := g^{-1}(e) \subseteq S^1$ for edges $e \in E(X)$ form a collection of disjoint open intervals on S^1 . Let $E(X) = \{e_j: j \in J\}$ for some (possibly finite) $J \subseteq \mathbb{N}$ be an enumeration of the edge set of X , and let $\Delta = \{\delta_1, \delta_2, \dots\}$ be a countable dense subset of $S^1 \setminus \overline{\bigcup \{I_e: e \in E(X)\}}$. Set $E_n = \{e_i: i \in [n]\}$ and $\Delta_n = \{\delta_i: i \in [n]\}$ (if Δ is empty, Δ_n is empty, too).

For $n \in \mathbb{N}$, let $C_n = \{J_1^n, \dots, J_{k_n}^n\}$ denote the set of connected components of $S^1 \setminus (\Delta_n \cup \bigcup \{I_e: e \in E_n\})$. Let $V_n = \{v_J: J \in C_n\}$, $F_n = \{f_e: e \in E_n\}$ and $D_n = \{d_\delta: \delta \in \Delta_n\}$ be duplicate sets of C_n , E_n and Δ_n respectively. In our Eulerian decomposition (G_n, η_n) , the graph G_n will be a cycle with vertex set V_n and edge set $E(G_n) = F_n \sqcup D_n$.

Define $\eta_n(v_J) := g(\bar{J})$ for each $v \in V_n$. By construction, $\eta_n(v)$ is a standard Peano subcontinuum of X , giving (E1). Set $\eta_n(f_e) := e$ and $\eta_n(d_\delta) := \delta$ for (E2)–(E5). Since every interval in $\{I_e: e \in E_n\}$ and every point in Δ_n is incident with the closure of precisely two components of C_n , transferring this assignment to G_n satisfies (E6) and (E7) (formally, if $\bar{I}_e \cap J \neq \emptyset$ we put $f_e \sim v_J$, and similarly, if $\delta \in \bar{J}$, put $d_\delta \sim v_J$). Hence, all properties of Definition 14.3.2 are satisfied, and so (G_n, η_n) is an Eulerian decomposition of X .

To see that $((G_n, \eta_n): n \in \mathbb{N})$ is an approximating sequence, note that for (A1), it is easily verified that $(G_{n+1}, \eta_{n+1}) / \langle e_{n+1}, d_{n+1} \rangle = (G_n, \eta_n)$. For (A2), note that by our density assumption on Δ , it follows that $\text{mesh}(V_n) \rightarrow 0$. By elementary topological arguments, this implies that also $\text{mesh}(\{\eta_n(v): v \in V_n\}) \rightarrow 0$, i.e. $w(G_n, \eta_n) \rightarrow 0$. \square

14.3.1.3. *A link between even-cut property and Eulerian decompositions.* Our second motivation for Eulerian decompositions is that by permitting the model graph G to be Eulerian, and not necessarily only a cycle, such decompositions can be built assuming just the even-cut condition, as demonstrated by the following observation which forms the blueprint for the more intricate constructions in the later chapters.

For the construction, we recall the following notion:

DEFINITION 14.3.7 (Intersection graph). For \mathcal{U} a family of subsets of X , the associated *intersection graph* $G_{\mathcal{U}}$ is the graph with vertex set \mathcal{U} , and an edge UV for $U \neq V \in \mathcal{U}$ whenever $U \cap V \neq \emptyset$.

If \mathcal{U} is a finite cover of a Peano continuum X , it follows from the connectedness of X that $G_{\mathcal{U}}$ is a finite connected graph.¹⁵

BLUEPRINT 14.3.8. *Suppose X satisfies the even-cut condition. Then any Peano partition of X into standard subspaces gives rise to an Eulerian decomposition for some suitable choice of dummy edges.*

PROOF. Let \mathcal{U} be a (finite) Peano partition of X into standard subspaces. Let $F \subseteq \mathcal{U}$ denote the collection of standard subspaces consisting of a single edge, and put $V = (\mathcal{U} \setminus F) \cup S$ where S is the finite collection of isolated points of $X - F$.

Now let G' be any graph with vertex set V and edge set F satisfying (E4), (E5) and (E6) of Definition 14.3.2. Our task is to add some new dummy edges D to G' to form a supergraph G that will be the desired Eulerian decomposition satisfying (E7).

Towards this, consider the auxiliary graph $H = (V, E_H)$ given by the intersection graph G_V on V associated with the cover V of $X - F$. We shall prove that we can find a multi-subset $D \subseteq E_H$ as desired.

As a first step, we claim that for each component C of H , the number of odd-degree vertices of G' in C is even. To see the claim, note first that $X - F$ has finitely many connected components, Lemma 14.1.11, and for every component C of H , the underlying subset $\bigcup C$ is a connected component of $X - F$ by (E1). Thus, the bipartition (C, D) with $D = V - C$ of $V = V(H) = V(G')$ induces a bipartition of $\mathfrak{G}(X)$, and hence an edge cut $B := E(\bigcup C, \bigcup D) \subseteq F$ of X , which must be even by assumption. However, property (E6) of G' implies that $E(C, D) = B$ is also an edge cut of G' containing the same edges. In particular, the quotient graph G'_C of G' where we collapse D to a single vertex v_D has the property that v_D has even degree, as v_D is adjacent precisely to the evenly many edges in B , plus possibly some loops (which do not affect the parity of the vertex degree). By the *Handshaking Lemma*, the number of odd-degree vertices in G'_C is even. Since v_D has even degree, it follows that the number of odd-degree vertices of G'_C in C (and hence also of G' in C) is even, and thus the claim follows.

Hence, we may pair up the odd-degree vertices of G' such that pairs lie in the same component of H . For each such pair $\{u, v\}$, consider a $u - v$ path in H . By taking the mod-2 sum over the edge sets of all these paths, we obtain an edge set $D_1 \subseteq E_H$ such that by adding D_1 to G' , one obtains an even graph G'' .

Since the intersection graph H is connected, we may find an edge set $D_2 \subseteq E_H$ such that adding D_2 to G'' results in a connected graph. Then define $G := G'' \cup 2 \cdot D_2$, i.e. for

¹⁵For a cover \mathcal{U} , the intersection graph $G_{\mathcal{U}}$ is sometimes also called the *nerve* of the cover.

every edge in D_2 we add two parallel dummy edges to G , in order to ensure connectedness without affecting the degree parity conditions.

Finally, to make sure that property (E7) of Definition 14.3.2 is satisfied, note that by definition of the intersection graph H , for every $d = xy \in E_H$, the sets $x, y \in V$ intersect, and hence we may choose a point (i.e. a trivial arc) $\eta(d)$ contained in $x \cap y \subseteq X$, satisfying property (E7) as required. \square

14.3.2. Obtaining an Edge-Wise Eulerian Map.

14.3.2.1. *Translating combinatorial information to topology.* For the benefit of clarity, and because we will need to jump between combinatorial and topological graphs, we denote for a combinatorial multi-graph G by $|G|$ the underlying topological space. Recall that for an edge e of a finite multi-graph or a Peano continuum, we write $e(0)$ and $e(1)$ for the two end vertices of e , and $e(x)$ for $x \in (0, 1)$ for the corresponding interior point on e .

DEFINITION 14.3.9 (Usc function, covering function). For a topological space X let $2^X = \{A \subseteq X : A \text{ nonempty, closed}\}$. A function $g: Y \rightarrow 2^X$ is *upper semi-continuous* (usc) if for all $y \in Y$ and all open sets $U \supset g(y)$ there is an open neighbourhood V of y such that $\bigcup_{y' \in V} g(y') \subseteq U$. The function g is said to *cover* X if $X = \bigcup \{g(y) : y \in Y\}$.

LEMMA 14.3.10. *Suppose (G, η) is an Eulerian decomposition of some Peano continuum X . Then the map $\hat{\eta}: |G| \rightarrow 2^X$ given by*

- $\hat{\eta}(v) := \eta(v)$ for all $v \in V$, and
- $\hat{\eta}(e(y)) := \{(\eta(e))(y)\}$ for all $e \in E(G)$ and $y \in (0, 1)$

defined on the 1-complex $|G|$ of G is upper semi-continuous, covers X , and is injective and acts as identity for points on real edges.¹⁶ Moreover, $\text{diam}(\hat{\eta}(y)) \leq w(G, \eta)$ for all $y \in |G|$.

PROOF. First, it is immediate from property (E4) that $\hat{\eta}$ covers X . Next, the usc-condition for $\hat{\eta}$ is evidently satisfied for interior points on edges of G . So consider a vertex $v \in G$ and an open set $U \subseteq X$ with $P = \eta(v) \subseteq U$. To simplify notation, let us write $f_X := \eta(f)$ for every edge $f \in F$, and similarly $d_X := \eta(d)$ for every edge $d \in D$.

By (E6), every edge $f \in F$ incident with v in G , say $f(j) = v$, satisfies that $f_X(j) \in \eta(v)$, and hence $\overline{f_X} \cap U$ is an open neighbourhood of $f_X(j) \in \overline{f_X} \subseteq X$. Since $\hat{\eta}$ acts as the identity between f and f_X , there is an open neighbourhood V_f of v in \overline{f} such that $\bigcup_{y' \in V_f} \hat{\eta}(y') = \overline{f_X} \cap U$. By (E7), we similarly obtain an open set V_d for every $d \in D$. Together, this yields that

$$V = \{v\} \cup \bigcup \{V_f : f \in F, f \sim v\} \cup \bigcup \{V_d : d \in D, d \sim v\}$$

is an open neighbourhood in $|G|$ of the vertex v satisfying that $\bigcup_{x' \in V} \hat{\eta}(x') \subseteq U$, which establishes that $\hat{\eta}$ is upper semi-continuous.

¹⁶Interior points of a dummy edge d for which $\eta(d)$ is trivial are mapped constantly to that singleton.

That $\hat{\eta}$ is injective and acts as identity for points on real edges follows from (E5). Finally, that $\text{diam}(\hat{\eta}(y)) \leq w(G, \eta)$ for all $y \in |G|$ is clear from construction. \square

Lastly, we record how the usc-maps corresponding to two comparable Eulerian decompositions relate to each other:

LEMMA 14.3.11. *Let X be a Peano continuum. For two Eulerian decompositions $\mathcal{D}_1 = (G_1, \eta_1)$ and $\mathcal{D}_2 = (G_2, \eta_2)$ of X with $\mathcal{D}_1 \preceq \mathcal{D}_2$, let $\varrho: |G_2| \rightarrow |G_1|$ denote the edge-contraction map corresponding to $G_1 \preceq G_2$. Then the associated usc-maps $\hat{\eta}_1$ and $\hat{\eta}_2$ satisfy $\hat{\eta}_2(y) \subseteq \hat{\eta}_1(\varrho(y))$ for all $y \in |G_2|$.*

PROOF. It suffices to prove the lemma in the case where we contract a single edge, say $\mathcal{D}_1 = \mathcal{D}_2/e$ with $e = ab$. In this case,

$$\varrho: |G_2| \rightarrow |G_1|, z \mapsto \begin{cases} z & \text{for all } z \in |G_2| \setminus \bar{e}, \text{ and} \\ v_e & \text{for all } z \in \bar{e} = \{a\} \cup e \cup \{b\}. \end{cases}$$

Also, according to Lemma 14.3.3, we have $G_1 = G_2/e$ and η_1 is given by

- $\eta_1(v_e) = \eta_2(a) \cup \eta_2(e) \cup \eta_2(b)$,
- $\eta_1(v) = \eta_2(v)$ for all $v \neq v_e$, and
- $\eta_1(f) = \eta_2(f)$ for all $f \in E(G_2) \setminus \{e\}$.

To verify the assertion of the lemma, consider some $z \in |G_2|$. If z is an interior point of some edge $f \neq e$, then it follows from the statement in the third bullet point that $\hat{\eta}_1(\varrho(z)) = \hat{\eta}_1(z) = \hat{\eta}_2(z)$. Similarly, if z is a vertex other than a or b , then it follows from the second bullet point that $\hat{\eta}_1(\varrho(z)) = \hat{\eta}_1(z) = \hat{\eta}_2(z)$. Finally, if z is an end vertex or interior point of e , then it follows from the first bullet point that $\hat{\eta}_1(\varrho(z)) = \hat{\eta}_1(v_e) = \eta_2(a) \cup \eta_2(e) \cup \eta_2(b) \supseteq \hat{\eta}_2(z)$. \square

14.3.2.2. *Construction of edge-wise Eulerian maps.* We now prove our main theorem of this chapter that every approximating sequence of Eulerian decompositions gives rise to an edge-wise Eulerian map, completing the proof of (iii) \Rightarrow (ii).

THEOREM 14.3.12 (Mapping Theorem). *Any Peano continuum X admitting an approximating sequence of Eulerian decompositions is edge-wise Eulerian.*

PROOF. Let $(\mathcal{D}_n: n \in \mathbb{N})$ with $\mathcal{D}_n = (G_n, \eta_n)$ be an approximating sequence of Eulerian decompositions for X , each G_n with edge bipartition $E_n = F_n \sqcup D_n$ into real and dummy edges. Note that by property (A1) and Definition 14.3.4, we have G_n is a contraction minor of G_{n+1} for all $n \in \mathbb{N}$, and hence the sequence $(G_n: n \in \mathbb{N})$ forms an inverse system of finite Eulerian multi-graphs under contraction bonding maps. Hence, the inverse limit $\Gamma = \varprojlim G_n$ is an Eulerian graph-like continuum, see [70, Thm. 13, Prop. 17]. Write $F = \bigcup F_n$ and $D = \bigcup D_n$. Then $E(\Gamma) = F \sqcup D$. Note that there is a natural bijection between F and $E(X)$ via $\eta(f) := \eta_n(f)$ if $f \in F_n$, which is well defined by property (C3).

Further, it is readily checked that (A2) and (E4) imply that η is onto, while (E5) implies that η is injective.

We now construct a continuous surjection $\hat{\eta}: |\Gamma| \rightarrow X$ such that $\hat{\eta}$ is injective for interior points on $f \in F$ and $\hat{\eta} \upharpoonright f: f \rightarrow \eta(f)$ is a homeomorphism for interior points on $f \in F \subseteq E(\Gamma)$ to its associated edge $\eta(f) \in E(X)$ for all $f \in F$. For the construction of $\hat{\eta}$, consider first for each $n \in \mathbb{N}$ the function

$$q_n: |\Gamma| \rightarrow 2^X, \quad z = (z_i: i \in \mathbb{N}) \mapsto \hat{\eta}_n(z_n),$$

which, by Lemma 14.3.10, is upper semi-continuous, covering, and is injective and acts as identity for points on edges $f \in F$. Moreover, Lemma 14.3.11 shows that

$$(\ddagger) \quad q_{n+1}(z) \subseteq q_n(z)$$

for all $n \in \mathbb{N}$ and $x \in |\Gamma|$. Thus, $\bigcap_{n \in \mathbb{N}} q_n(z) \subseteq X$ is a nested intersection of non-empty closed subsets of X , and so it follows from compactness of X that this intersection is non-empty. At the same time, however, we have $\text{diam}(q_n(z)) \leq w(G_n, \eta_n) \rightarrow 0$ by Lemma 14.3.10 and (A2), and so this intersection must be a singleton for each $z \in |\Gamma|$. Hence, there is a function

$$\hat{\eta}: |\Gamma| \rightarrow X \quad \text{defined by} \quad \{\hat{\eta}(z)\} = \bigcap_{n \in \mathbb{N}} q_n(z) \quad \text{for all } z \in |\Gamma|.$$

As the image of each q_n is an upper semi-continuous function that covers X and satisfies (\ddagger) , it follows from [121, General Mapping Theorem 7.4] that the map $\hat{\eta}: |\Gamma| \rightarrow X$ is a continuous surjection as desired. Further, it is clear by the definition of $\hat{\eta}$ that for every real edge $f \in F$ we have $\hat{\eta}^{-1}(\eta(f)) = f$ and $\hat{\eta} \upharpoonright f$ acts a identity from $f \in F$ onto $\eta(f) \in E(X)$.

In order to complete the proof, note that since Γ is an Eulerian graph-like continuum, there is an Eulerian map $h: S^1 \rightarrow |\Gamma|$. In particular, h is a continuous surjection with the property that for every open edge $f \in E(\Gamma)$ (dummy and real edges alike) we have $I_f := h^{-1}(f)$ is an interval on S^1 and $h \upharpoonright I_f$ is a homeomorphism from I_f onto f .

We now claim that $g = \hat{\eta} \circ h: S^1 \rightarrow X$ is the desired edge-wise Eulerian map. Clearly, as the composition of surjective functions, g is itself a surjection from the circle onto X . To see that g is edge-wise Eulerian, we need to check that g sweeps through each edge of X precisely once. So let $e \in E(X)$ be arbitrary. By our considerations above, there is a unique $f \in F$ with $\eta(f) = e$. But $g^{-1}(e) = h^{-1} \circ \hat{\eta}^{-1}(e) = I_f$. Since $h_f = h \upharpoonright I_f$ is a homeomorphism from I_f onto f , and $\hat{\eta}_f = \hat{\eta} \upharpoonright f$ acts as identity between interior points of f and e , it follows that $g \upharpoonright I_f$ is as the composition of the homeomorphisms $\hat{\eta}_f \circ h_f$ itself a homeomorphism from I_f onto $\eta(f) = e$. Thus, we have verified that g is an edge-wise Eulerian map, and hence that X is edge-wise Eulerian. \square

14.3.3. Simplicial Maps. In this last section on Eulerian decompositions, we describe an equivalent condition to Definition 14.3.4 about compatible Eulerian decompositions, which lends itself better to the constructions in the next two chapters.

DEFINITION 14.3.13 (Contraction map, edge-contraction map). We call a surjective map $\varrho: G_2 \rightarrow G_1$ between two graphs $G_i = (V_i, E_i)$ a *contraction map* if

- (Q1) $\varrho(V_2) = V_1$,
- (Q2) ϱ restricts to a bijection between $E_2 \setminus \varrho^{-1}(V_1)$ and E_1 ,
- (Q3) $\varrho(e(j)) = (\varrho(e))(j)$ for all $e \in E_2 \setminus \varrho^{-1}(V_1)$ and $j \in \{0, 1\}$, and
- (Q4) $\varrho(e(j)) = \varrho(e)$ for all $e \in E_2 \cap \varrho^{-1}(V_1)$ and $j \in \{0, 1\}$.

If additionally,

- (Q5) $\varrho^{-1}(v)$ is a connected subgraph of G_2 for all $v \in V(G_1)$,

then the map ϱ is called an *edge-contraction map*.

Thus, an edge-contraction map $\varrho: G_2 \rightarrow G_1$ is precisely a map witnessing that $G_1 \preceq G_2$, whereas a contraction map may identify vertices that are not necessarily connected by an edge.

DEFINITION 14.3.14. Let $\mathcal{D}_1 = (G_1, \eta_1)$ and $\mathcal{D}_2 = (G_2, \eta_2)$ be decompositions of a Peano continuum X . A contraction map $\varrho: G_2 \rightarrow G_1$ is called *η -compatible* if

$$\eta_1(x) = \bigcup \{ \eta_2(y) : y \in \varrho^{-1}(x) \}$$

for all $x \in V(G_1) \cup E(G_1)$.

LEMMA 14.3.15. Suppose $\mathcal{D}_1 = (G_1, \eta_1)$ and $\mathcal{D}_2 = (G_2, \eta_2)$ are both decompositions of a Peano continuum X . Then $\mathcal{D}_1 \preceq \mathcal{D}_2$ if and only if there is an η -compatible edge-contraction map $\varrho: G_2 \rightarrow G_1$.

PROOF. This follows from the observation that $G_1 \cong G_2 / \langle e_1, \dots, e_k \rangle$ if and only if there is an edge contraction map $\varrho: G_2 \rightarrow G_1$ such that $\varrho^{-1}(V_1) = \{e_1, \dots, e_k\}$. \square

14.4. Product-structured ground spaces

14.4.1. Introduction. In this chapter we establish that the Eulerianity conjecture holds for Peano continua X whose ground space has a product structure, in other words, where $\mathfrak{G}(X) = V \times P$ is the product of a (compact) zero-dimensional space V with a Peano continuum P , thereby proving the second case (B) of our main result Theorem 14.1.4 stated in the introduction.

THEOREM 14.4.1. Let X be a Peano continuum with ground space $\mathfrak{G}(X) = V \times P$ where V is a compact zero-dimensional space and P a Peano continuum. Then X is Eulerian if and only if it satisfies the even-cut condition.

Bula, Nikiel and Tymchatyn have asked whether the Eulerianity Conjecture holds for spaces with ground set $C \times K$, where C is the Cantor set and K is any continuum (not necessarily Peano), [41, Problem 3]. For this question, our Theorem 14.4.1 gives a strong answer in the case where $P = K$ is a Peano continuum. For our result, the assumption that P is Peano is crucial. To demonstrate this, recall that Bula, Nikiel and Tymchatyn have also asked whether a Peano continuum X with ground space a continuum (not necessarily Peano) satisfies the Eulerian conjecture [41, Problem 2]. We believe that this question is, maybe unexpectedly so, at least as hard as the situation discussed in Theorem 14.4.1: indeed, with the techniques from this chapter one can establish the Eulerianity conjecture for spaces X with ground space a Cantor fan, or even a generalised fan of the form $\mathfrak{G}(X) = (V \times P)/\{(v, p) : v \in V\}$ for some $p \in P$.

14.4.1.1. *Blanket assumptions.* Given our work in Chapter 14.2, for our proof of Theorem 14.4.1, we may assume throughout this chapter, without any loss of generality, that our Peano continuum X satisfies the following additional assumptions:

- X is a Peano graph without loops by the second reduction result, Theorem 14.2.14.
- X has diameter bounded by 1.
- P is not a singleton (as otherwise, X is a graph-like continuum, a class for which the Eulerianity conjecture is already known to hold [70]).

14.4.1.2. *Proof strategy.* After having established Theorem 14.1.1, by (iii) \Rightarrow (i) we need to construct an approximating sequence of Eulerian decompositions for X . The first ingredient to construct this approximation is the observation that every Peano graph X with ground space $\mathfrak{G}(X) = V \times P$ exhibits a fractal-like behaviour as follows: for every point $(v, p) \in V \times P$ and every $\varepsilon > 0$ there exists $V' \times P' \subseteq V \times P$ such that $v \in V' \subseteq V$ is clopen, $p \in \text{int}(P') \subseteq P' \subseteq P$ and P' is a regular subcontinuum of P , and $X' := X[V' \times P']$ is again a Peano graph of the same form as in the theorem, see Lemma 14.4.16. Let us call such a space X' a *tile* of X . Utilising this fractal-like behaviour, our main technical result in this chapter is the so-called *decomposition theorem*, Theorem 14.4.21, which says roughly that any Peano-continuum with product-structured ground space can be decomposed into edge-disjoint tiles all of arbitrarily small diameter plus some finitely many cross edges that go between tiles, such that most of the tiles now satisfy the even-cut condition.

Crucially, to control all edge cuts simultaneously, we borrow and extend in Section 14.4.2 the techniques of topological spanning trees, fundamental circuits and infinite thin sums from recently developed infinite graph and infinite matroid theory, see [54, §8.7] and [30, 38].

In the final section of this chapter, Section 14.4.5, we then demonstrate how this decomposition theorem can be used, now using the assumption that the original space X satisfied the even-cut condition for the first time, to construct an approximating sequence of Eulerian decompositions for X .

14.4.2. Spanning Trees and the Even-Cut Condition. Before we embark on our proof, we need some preliminary results about *spanning trees* in graph-like continua. These notions are by now standard in the theory of infinite graphs (see e.g. [54, §8] and [53]) and they do generalise nicely to graph-like continua. Indeed, this is not by accident and could be seen as a corollary to the general theory of infinite matroids and matroids induced by graph-like spaces, see [30, 38]. However, as there are direct proofs for the results we need, and so as to make it easier for the reader, we simply state and prove what we need.

LEMMA 14.4.2. *The following are equivalent for a standard subspace T of a graph-like continuum Z :*

- (1) T is edge-minimally connected,
- (2) T is uniquely arc-connected,
- (3) T is connected and does not contain a non-trivial cycle, and
- (4) T is a dendrite.

PROOF. Recall that a graph-like continuum is hereditarily locally connected, so every subcontinuum of Z is automatically Peano [70, Corollary 8]. The equivalence of (3) and (4) holds by the definition of *dendrite* (see [121, 10.1]). The equivalence of (2) and (3) is easy. To see that (1) and (3) are equivalent, note that if T contains a cycle, then deleting an edge on that cycle does not disconnect T , and conversely, if deleting an edge $e = xy$ does not disconnect T , then for any $x - y$ arc P in $T - e$, we have $P \cup e$ is a cycle. \square

DEFINITION 14.4.3 (Spanning tree). A subspace Y of a graph-like continuum (X, V, E) is called *spanning* if $V \subseteq Y$. A spanning standard subspace T of a graph-like continuum Z is called a *spanning tree* of Z provided it satisfies one (and therefore every) condition in Lemma 14.4.2.

Spanning trees of graph-like continua are easy to construct, because connectivity is preserved under nested intersections—so in order to obtain a standard subspace with property (1), one only needs to enumerate all edges from a graph-like continuum, and then delete the next edge in line as long as it is not a bridge at that current stage.

DEFINITION 14.4.4 (Fundamental cuts; fundamental cycles). Let T be a spanning tree of a graph-like continuum Z .

- If $f \in E(T)$, then by Lemma 14.1.11 and property (1) in Lemma 14.4.2, the space $T - f$ has two connected components with vertex sets say A and B which form a clopen partition of $V(T) = V(Z)$. The corresponding edge cut $E(A, B)$ of Z is also called the *fundamental cut* of f , denoted by D_f .
- If $e \notin E(T)$, then T contains a unique standard arc A between the endpoints of e . The *fundamental cycle* C_e is given by the edge set $E(A) \cup \{e\}$. Note that $Z[C_e]$ is indeed homeomorphic to S^1 .

Observe that for $f \in E(T)$ and $e \notin E(T)$ one has $e \in D_f$ if and only if $f \in C_e$.

DEFINITION 14.4.5 (Thin family). Let E be a set. A multi-set $(C_j: j \in J)$ of subsets of E is called *thin* if for all $e \in E$, we have $|\{j \in J: e \in C_j\}| < \infty$.

DEFINITION 14.4.6 (Thin sum). For a thin family $(C_j: j \in J)$, the sum

$$C = \sum_{j \in J} C_j := \{e \in E: |\{j \in J: e \in C_j\}| \text{ is odd}\}$$

is well-defined. We say that C is the *thin sum* over the $(C_j: j \in J)$.

The following theorem is in some sense a natural generalisation of the corresponding theorem for finite and infinite graphs [54, Theorems 1.9.5 and 8.7.1] respectively.

THEOREM 14.4.7. *Let $X = (V, E)$ be a graph-like continuum, and $D \subseteq E$. Then all topological cuts of $X[D]$ are even if and only if D is a thin sum of fundamental cycles of any spanning tree of X .*

PROOF. Compare to [54, 8.7.1], where this statement is proved for Freudenthal compactifications of locally finite graphs (which form a proper subclass of the class of graph-like continua). For additional background, see [56].

To see that a thin sum of cycles satisfies the even-cut condition, recall that by [70, Lemma 6], any single cycle C intersects any topological cut of X in an even number of edges. This extends immediately to finite symmetric differences, as is easily verified. But then this also extends to thin sums of cycles: since cuts are finite, only finitely many cycles in our thin sum can meet the cut, and so the result follows.

For the converse implication, suppose $X[D]$ satisfies the even-cut condition and fix any spanning tree T of X . We show that $D = \sum_{e \in D \setminus E(T)} C_e$. To see that this sum is well-defined, observe that $f \in C_e$ if and only if $e \in D_f$. Since fundamental cuts are finite, the above is the sum over a thin family. To prove the equality, we claim that $D' := D + \sum_{e \in D \setminus E(T)} C_e = \emptyset$. First, it is clear that $D' \subseteq E(T)$, since every edge $e \in D \setminus E(T)$ has been eliminated by the corresponding C_e (and all other edges in C_e lie in $E(T)$ by construction).

Second, the existence of an edge $f \in D'$ leads to a contradiction as follows: since $f \in D' \subseteq E(T)$, it follows that $f \in D_f \cap D' \subseteq D_f \cap E(T) = \{f\}$.

Thus, D_f is a topological cut meeting D' in an odd number of edges. This contradicts the fact that both D (by assumption) and the thin sum $\sum_{e \in D \setminus E(T)} C_e$ (by virtue of the first proven implication) meet every cut in an even number of edges. \square

14.4.3. Sparse Edge Sets.

14.4.3.1. *Properties of sparse edge sets.* Given a Peano graph X with ground set $\mathfrak{G}(X) = V \times P$, we will now investigate under which conditions certain (infinite) edge sets can be removed without harming local connectedness or density. Recall from Section 14.1.3.1 that a subset $F \subseteq E(X)$ of edges is called *sparse (in X)* if $X[F]$ is a graph-like compactum (i.e. if $\overline{\bigcup F} \setminus \bigcup F$ is zero-dimensional). Note that the property of an edge set F being sparse is inherited by subsets of F .

LEMMA 14.4.8. *Let X be a Peano continuum [Peano graph] X and $F \subseteq E(X)$ a sparse edge set. Then the following assertions hold.*

- (i) *The non-trivial components of $X - F$ form a zero-sequence of standard Peano continua [Peano graphs].*
- (ii) *If $\mathfrak{G}(X)$ contains no 1-point components, then $\mathfrak{G}(X - F) = \mathfrak{G}(X)$.*
- (iii) *If for some $\delta > 0$ all components of $\mathfrak{G}(X)$ have diameter at least δ , then $X - F$ consists of finitely many Peano continua [Peano graphs], so is locally connected.*

PROOF. Let \mathcal{D} denote the collection of components of $X - F$. It is clear that each element of \mathcal{D} is a standard subcontinuum. We first show that \mathcal{D} forms a null-family. Otherwise, for some $\varepsilon > 0$ there are infinitely $D_n \in \mathcal{D}$ with $\text{diam}(D_n) \geq \varepsilon$ for all $n \in \mathbb{N}$. By sequential compactness of the hyperspace [121, 4.18], we may assume that $D_n \rightarrow D$, i.e. D_n converges to a continuum D in the Hausdorff metric [121, 4.2]. And since $\text{diam}(D_n) \geq \varepsilon$ for all $n \in \mathbb{N}$, we have – by the properties of the Hausdorff metric – that $\text{diam}(D) \geq \varepsilon$, too. Moreover, since edges are open, we necessarily have $D \subseteq \mathfrak{G}(X)$. But now, since D is a non-trivial continuum and $\overline{\bigcup F} \setminus \bigcup F$ is zero-dimensional, there is $x \in D$ and a connected neighbourhood U of x in X with $U \cap X[F] = \emptyset$. However, since $D_n \rightarrow D$ there exists $N \in \mathbb{N}$ such that $D_n \cap U \neq \emptyset$ for all $n \geq N$. Therefore, $D \cup U \cup D_N$ is a connected subset of $X - F$, contradicting that D_N was a component. This contradiction establishes that \mathcal{D} forms a null-family, and hence that the subfamily $\mathcal{D}' \subseteq \mathcal{D}$ of non-trivial elements of \mathcal{D} forms a zero-sequence.

To see that each $D \in \mathcal{D}'$ is a Peano continuum, note that by construction, $D \setminus \overline{F}$ is open, so hence locally connected, and moreover dense in D . It follows that the interior of D is locally connected with zero-dimensional boundary (as the boundary is a subset of the zero-dimensional $X[F] \cap \mathfrak{G}(X)$), and so D must be a Peano continuum, since if a continuum fails to be locally connected at some point, then it fails to be locally connected at all points of a non-trivial subcontinuum, [121, 5.13].

Finally, if X is a Peano graph, then each $D \in \mathcal{D}'$ is a Peano graph too, i.e. has dense edge set. Suppose to the contrary that for some non-trivial component D , its edge set $E(D) = \{e \in E(X) : e \subseteq D\}$ is not dense in D . Since $\overline{F} \setminus F$ is zero-dimensional, there is $x \in D$ and a connected open neighbourhood U of x in X with $U \cap \overline{\bigcup (E(D) \cup F)} = \emptyset$. Since by assumption $E(X)$ is dense in X and forms a zero-sequence by Lemma 14.1.11,

there is an edge $e \in E(X)$ completely contained in U . But since $U \subseteq D$, this implies $e \in E(D)$, a contradiction.

For (ii), note that the inclusion $\mathfrak{G}(X - F) \subseteq \mathfrak{G}(X)$ holds for all edge sets $F \subseteq E(X)$ and all X , as free edges in $E(X) \setminus F$ remain free in $X - F$. For the converse inclusion to hold, however, the additional assumptions of the statement are necessary. So suppose there was $x \in \mathfrak{G}(X) \setminus \mathfrak{G}(X - F)$. Then there is a free arc α in $X - F$ with $x \in \alpha$. But then $\bar{\alpha} \cup X[F]$ is a compact graph-like space in X forming a neighbourhood of x in X , from which it follows that x forms a singleton component in X .

For (iii), it now follows from the previous step that every component $X - F$ has diameter at least δ , and so by (i), $X - F$ must consist of finitely many Peano continua. \square

14.4.3.2. *Sparse spanning trees.* The purpose of this section is to give a fairly general procedure how to find non-trivial sparse edge sets.

LEMMA 14.4.9. *Let X be a Peano continuum. For every zero-dimensional compact set $Y \subseteq \mathfrak{G}(X)$, there exists a standard graph-like continuum $Z \subseteq X$ with $Y \subseteq Z$.*

PROOF. The proof modifies an idea by Ward of *approximating a Peano continuum by finite trees*, see [161] and [162].

Let $(\mathcal{U}_n : n \in \mathbb{N})$ be a refining sequence of finite 2^{-n} Peano covers of X where $U_0 = \{X\}$ is the trivial cover. For a subset $A \subseteq X$, define $\mathcal{U}_n \upharpoonright A := \{U \in \mathcal{U}_n : U \cap A \neq \emptyset\}$. Recursively, we will define finite, i.e. compact trees $T_n \subseteq X$ and finite vertex sets $V_n \subseteq T_n$ such that for all $n \in \mathbb{N}$,

- (1) $T_n \subseteq T_{n+1}$ as topological subspaces,
- (2) $V_n \subseteq V_{n+1}$,
- (3) V_n is the set of branch- and end-vertices of T_n ,
- (4) $\mathcal{U}_n \upharpoonright Y \subseteq \mathcal{U}_n \upharpoonright T_n$, and
- (5) $\mathcal{U}_n \upharpoonright Y$ covers $T_{n+1} \setminus T_n$, and
- (6) $\mathcal{U}_n \upharpoonright Y$ covers $V_{n+1} \setminus V_n$.

Let $T_0 = V(T_0) = \{t_0\}$ be an arbitrary singleton tree. Since $U_0 = \{X\}$, this satisfies (4). All other conditions are trivial or vacuous at this point. This completes the base case. For the recursion step, suppose that T_0, \dots, T_n are already defined according to (1) – (6), and pick finitely many points $A = \{a_1, \dots, a_k\}$ such that $\mathcal{U}_{n+1} \upharpoonright Y = \mathcal{U}_{n+1} \upharpoonright A$. Let $S_0 := T_n$, $V(S_0) := V_n$ and suppose we already have constructed a sequence of finite tree $S_0 \subseteq S_1 \subseteq \dots \subseteq S_i$ for $i < k$ such that S_i contains $\{a_1, \dots, a_i\}$ and such that $S_i \setminus T_n$ is covered by $\mathcal{U}_n \upharpoonright Y$. Consider a_{i+1} . Again, if $a_{i+1} \in S_i$, set $S_{i+1} := S_i$. Otherwise, pick $U \in \mathcal{U}_n$ such that $a_{i+1} \in U$, and also pick $t \in T_n \cap U$ (possible by (4)). Pick an arc $\alpha : I \rightarrow U$ from t to a_{i+1} . Since S_i is compact, there is a maximal $x_{i+1} < 1$ such that $\alpha(x_{i+1}) \in S_i$. Define $S_{i+1} = S_i \cup \alpha([x_{i+1}, 1])$, and $V(S_{i+1}) = V(S_i) \cup \{\alpha(x_{i+1}), a_{i+1}\}$. Since α was an arc completely contained in U , we have $S_{i+1} \setminus T_n$ is covered by $\mathcal{U}_n \upharpoonright Y$.

In the end, put $T_{n+1} := S_k$ and $V_{n+1} = V(S_k)$. Clearly, T_{n+1} is a finite tree with vertex set V_{n+1} . Moreover, by choice of A , it satisfies (4). Finally, (5) and (6) follow since all S_i satisfied that $S_i \setminus T_n$ is covered by $\mathcal{U}_n \upharpoonright Y$, and so then does $S_k = T_{n+1}$. This completes the recursive construction.

Define $T = \bigcup_{n \in \mathbb{N}} T_n$, and $V = \bigcup V_n$. Our aim is to show that $Z = \overline{T}$ is a graph-like continuum containing Y . Clearly, T is connected, and hence Z is compact connected. To see that Z covers Y , note that for any $y \in Y$, since $W_n := \bigcup (\mathcal{U}_n \upharpoonright \{y\})$ has vanishing diameter for $n \rightarrow \infty$, the family $\{W_n : n \in \mathbb{N}\}$ forms a neighbourhood base of y in X . By property (4), every W_n intersects T , and so $y \in \overline{T}$. Since $y \in Y$ was arbitrary, this shows $Y \subseteq \overline{T} = Z$. Finally, the proof that Z is graph-like essentially relies on the following observation:

Claim: *For every $p \notin Y$ there is a open set $U \subseteq X$ with $p \in U$ such that for some $n \in \mathbb{N}$ we have $U \cap \overline{T} \subseteq T_n$ and $U \cap \overline{V} \subseteq V_n$.*

To see the claim, note that if $p \notin Y$, then $\varepsilon = \text{dist}(p, Y) > 0$, and so there is n large enough such that $2^{-n} < \varepsilon$. Let $W := \bigcup (\mathcal{U}_n \upharpoonright Y)$ and $U = X \setminus W$. Then U is open and $p \in U$. Moreover, $\overline{T} \cap U = \overline{T} \setminus W = \left(\overline{T_n} \cup \overline{T \setminus T_n} \right) \setminus W \subseteq \overline{T_n} = T_n$ by property (5), and the fact that T_n is compact. Similarly, $\overline{V} \cap U = \overline{V} \setminus W \subseteq \overline{V_n} = V_n$ by property (6), and the fact that V_n is finite. This establishes the claim.

Finally, we argue that the set $V(Z) := Y \cup V$ is a vertex set for Z witnessing that Z is graph-like. First, by the claim, $V(Z)$ is closed in X and hence compact. Moreover, since each V_n is finite and Y is zero-dimensional, also $V(Z)$ is zero-dimensional by the countable sum theorem for dimension, [62, Thm. 1.5.2].

Finally, we need to show that each $p \in Z \setminus V(Z)$ has a neighbourhood homeomorphic to an open interval. So let $p \in Z \setminus V(Z)$. Let U be as in the claim, i.e. U is a neighbourhood of p such that $U \cap Z = U \cap \overline{T} \subseteq T_n$. Then $U \setminus V_n$ is open, and $(U \setminus V_n) \cap Z \subseteq T_n \setminus V_n$ consists of finitely many connected components, each homeomorphic to an open interval.

Finally, to make Z standard, define $Z' = Z \setminus \bigcup \{e : e \cap Z \neq \emptyset \neq Z \setminus e\}$. Since $Y \subseteq \mathfrak{G}(X)$, we still have $Y \subseteq Z'$, and further, Z' is still connected, as no half edge is needed for connectivity in Z . \square

DEFINITION 14.4.10 (Sparse spanning tree). Let X be a Peano continuum. A spanning tree T of X_\sim is *sparse* if its edge set $E(T)$ is sparse in X .

LEMMA 14.4.11 (Existence of sparse spanning trees). *Every Peano continuum X with $\mathfrak{G}(X) = V \times P$ admits a sparse spanning tree.*

PROOF. Pick $p \in P$, and put $Y := V \times \{p\}$, a compact zero-dimensional subset of $\mathfrak{G}(X)$. By Lemma 14.4.9, there exists a standard graph-like continuum $Z \subseteq X$ with $Y \subseteq Z$. Let $\pi : X \rightarrow X_\sim$ be the quotient map. Since Y intersects every component of

$\mathfrak{G}(X)$, it follows that $\pi(Z)$ is a spanning graph-like subcontinuum of X_\sim . Let $T \subseteq \pi(Z)$ be a spanning tree of X_\sim . Then $E(T) \subseteq E(X_\sim) = E(X)$, and since Z was graph-like, it is evident that $\overline{E(T)} \subseteq Z$ is a graph-like compactum, i.e. $E(T)$ is sparse in X . \square

14.4.4. Tiles in Peano Graphs with Product-Structured Ground Spaces. We discuss fractal properties of Peano continua X with ground space $\mathfrak{G}(X) = V \times P$.

14.4.4.1. *Tiles via horizontal restriction.* First, we discuss tiles that result by restricting to well-behaved subsets of V .

LEMMA 14.4.12. *Every locally connected compactum X with ground set $\mathfrak{G}(X) = V \times P$ [and dense edge set] is of the form $X = \bigoplus_{A \in \mathcal{A}} X_A$, where \mathcal{A} is a (finite) clopen partition of V and $X_A \subseteq X$ is a standard Peano continuum [Peano graph] with ground space $\mathfrak{G}(X_A) = A \times P$.*

PROOF. As a locally connected compactum, X has finitely many components, [107, VI §49, II Theorem 7]. Moreover, since P is connected, each component C is of the form $C = X[A_C \times P]$ with $A_C \subseteq V$. Since C is closed, it follows from compactness and the continuity of projection maps that $A_C \subseteq V$ is closed. Moreover, for distinct components $C \neq C'$ we clearly have $A_C \cap A_{C'} = \emptyset$. Therefore, every A_C is a clopen subset of V . Hence, the collection \mathcal{A} of such clopen $A_C \subseteq V$ is the desired (finite) clopen partition of V . \square

COROLLARY 14.4.13. *If X is a Peano graph with $\mathfrak{G}(X) = V \times P$, and $F \subseteq E$ is sparse, then there is a (finite) clopen partition \mathcal{A} of V such that $X - F = \bigoplus_{A \in \mathcal{A}} X_A$ where each $X_A \subseteq X$ is a standard Peano graph with ground space $\mathfrak{G}(X_A) = A \times P$.*

PROOF. By Lemma 14.4.8(iii), the space $X - F$ is locally connected with ground space $\mathfrak{G}(X) = V \times P$, so the assertion follows from Lemma 14.4.12. \square

COROLLARY 14.4.14. *If X is a Peano graph with $\mathfrak{G}(X) = V \times P$ and $B \subseteq V$ is clopen, then there is a (finite) clopen partition \mathcal{B} of B such that $X[B \times P] = \bigoplus_{B \in \mathcal{B}} X_B$ where each $X_B \subseteq X$ is a standard Peano graph with ground space $\mathfrak{G}(X_B) = B \times P$.*

PROOF. Since $F = E(B \times P, (V \setminus B) \times P)$ is a (finite) edge cut of X , the edge set F is sparse, and so the result follows from the previous Corollary 14.4.13, by taking \mathcal{B} to be the subcollection of \mathcal{A} of elements that intersect B . \square

14.4.4.2. *Tiles via vertical restriction.* Next, we discuss tiles that result by restricting to well-behaved subsets of P .

LEMMA 14.4.15. *Let X be a Peano graph, $x \in \mathfrak{G}(X)$, and $U \subseteq X$ a connected set such that $U \cap \mathfrak{G}(X)$ is a neighbourhood of x in $\mathfrak{G}(X)$. Then for every $\varepsilon > 0$ there is a connected neighbourhood V of x in X such that $V \subseteq B_\varepsilon(U)$.*

PROOF. If y is an endpoint of some edge e , write $B_\delta^e(y)$ (where $0 < \delta \leq 1$) for the half-open interval with end-point y of diameter δ on e . Then put

$$V := U \cup \{B_\epsilon^e(y) : e \in E \text{ and } y \in \bar{e} \cap U\} \subseteq X.$$

Then V is connected, and it is a neighbourhood of x in X (as almost all edges in E have diameter $< \epsilon$), and by construction, we have $V \subseteq B_\epsilon(U)$. \square

LEMMA 14.4.16. *For every Peano graph X with ground set $\mathfrak{G}(X) = V \times P$, every $W \subseteq P$ a regular closed Peano subcontinuum and for every $\epsilon > 0$, there is a (finite) clopen partition \mathcal{A} of V with $\text{mesh}(\mathcal{A}) \leq \epsilon$ such that $X[A \times W]$ is a Peano graph for all $A \in \mathcal{A}$.*

PROOF. By Lemma 14.4.12 it suffices to show that the induced subspace $X_W = X[V \times W]$ inherits local connectedness from X . This is trivial for points in the interior of X_W , i.e. interior points of edges, and points in $V \times \text{int}(W)$. So consider an arbitrary point $x = (v, w)$ for $v \in V$ and $w \in \partial W$, and fix $\delta > 0$. Our task is to find a connected open neighbourhood V of x in X_W of diameter at most δ . First, pick a connected open neighbourhood U of w in W with $\text{diam}(U) < \delta/3$. Then $V \times (U \cap \text{int}(W))$ is a non-empty open subset of X , and so it follows from local connectedness of X that there are $A \subseteq V$ clopen with $v \in A$, $B \subseteq U \cap \text{int}(W)$ open, and a connected open set $Y \subseteq X$ with $\text{diam}(Y) < \delta/3$, $Y \subseteq U$ and $X[A \times B] \subseteq Y$.

But then $Y' = Y \cup X[A \times U]$ is connected, and restricts to a neighbourhood of (v, w) in $\mathfrak{G}(X_W)$ of diameter $\text{diam}(Y') \leq \delta/3$. So applying Lemma 14.4.15 to Y' with $\epsilon = \delta/3$ provides a connected neighbourhood as desired. \square

14.4.4.3. Ground-space covering tiles.

LEMMA 14.4.17. *Suppose for a Peano continuum P with edges $E = E(P)$ and ground space $Z = Z(P)$, we have a set of edges F such that $Z \cup \bigcup F$ is locally connected. Then $Z \cup \bigcup F'$ is locally connected for all $F \subseteq F' \subseteq E$.*

PROOF. Let $Y = Z \cup \bigcup F$. By local connectedness, all components of Y are open, and so it follows from compactness that Y has finitely many components. Moreover, since the edges in $F' \setminus F$ form a zero-sequence of Peano subcontinua, the result now follows from (a natural adaption of) Lemma 14.1.13. \square

Relying on the results established above about sparse spanning trees, our aim for this short section is to prove the following theorem.

THEOREM 14.4.18. *The edge set $E(X)$ of every Peano graph X with ground space $\mathfrak{G}(X) = V \times P$ (with P non-degenerate) admits a bipartition $E(X) = E_1 \sqcup E_2$ into two edge sets both dense for $\mathfrak{G}(X)$ such that both $X_i = X[E_i]$ are locally connected.*

PROOF. Let $(\mathcal{U}_n : n \in \mathbb{N})$ be a decreasing sequence of 2^{-n} partitions for P with $\mathcal{U}_0 = \{P\}$. Let $\mathcal{R} = (R, \leq)$ be the corresponding *refinement tree*, that is $\mathcal{R}(n)$, the n th level of \mathcal{R} , indexes the elements of \mathcal{U}_n , so $\mathcal{U}_n = \{U_r : r \in \mathcal{R}(n)\}$, and $r \leq r'$ if and only if $U_r \supseteq U_{r'}$. Recall that each \mathcal{U}_n is finite, and so \mathcal{R} is a locally finite tree. Write $\mathcal{R}(\leq n) := \bigcup_{i \leq n} \mathcal{R}(i)$ and similarly $\mathcal{R}(< n) := \bigcup_{i < n} \mathcal{R}(i)$.

We now recursively construct

- a family of finite multicuts $\{\mathcal{A}_r : r \in \mathcal{R}\}$ of V , and
- subtrees $T_{r,A} \subseteq X_\sim$ for $r \in \mathcal{R}$ and $A \in \mathcal{A}_r$

such that

- (1) $r \leq r' \in \mathcal{R}$ implies $\mathcal{A}_r \succcurlyeq \mathcal{A}_{r'}$,
- (2) $\text{mesh}(\mathcal{A}_r) \leq 2^{-n}$ for $r \in \mathcal{R}(n)$,
- (3) for each $r \in \mathcal{R}(n)$ and $A \in \mathcal{A}_r$, the space

$$X_{r,A} = X[A \times U_r] \setminus \bigcup \{E(T_{A',s}) : s \in \mathcal{R}(< n), A' \in \mathcal{A}_s\}$$

is a Peano graph,

- (4) $T_{r,A}$ is a sparse spanning tree for $X_{r,A}$ for all $r \in \mathcal{R}$ and $A \in \mathcal{A}_r$ (unless $(X_{r,A})_\sim$ has a single vertex, in which case $T_{r,A}$ consists of an arbitrary edge from $X_{r,A}$).

For $n = 0$, and $r \in \mathcal{R}(0)$ the unique root of \mathcal{R} , the trivial (finite) clopen partition $\mathcal{A}_r = \{V\}$ is clearly sufficient. Now let $n \in \mathbb{N}$ and suppose we have already defined finite multicuts $\{\mathcal{A}_r : r \in \mathcal{R}(\leq n)\}$ of V , and subtrees $T_{r,A} \subseteq X_\sim$ for $r \in \mathcal{R}(\leq n)$ and $A \in \mathcal{A}_r$ according to (1)–(4). Consider $r \in \mathcal{R}(n)$. Since $X_{r,A}$ is a Peano graph by (3), we may use Lemma 14.4.11 to find sparse spanning trees $T_{r,A}$ for $X_{r,A}$ for each $A \in \mathcal{A}_r$, unless A is a singleton, in which case we let $T_{r,A}$ consist of an arbitrary edge from $X_{r,A}$. Then property (4) is satisfied. By Corollary 14.4.13, each

$$X'_{r,A} := X_{r,A} \setminus \bigcup \{E(T_{A',s}) : s \in \mathcal{R}(n), A' \in \mathcal{A}_s\}$$

remains locally connected. Consider an arbitrary successor s of r , i.e. some $s \in \mathcal{R}(n+1)$ with $r < s$. By Corollary 14.4.14 and Lemma 14.4.16, there is a (finite) clopen partition $\mathcal{B}_{A,s}$ of A with $\text{mesh}(\mathcal{B}_{s,A}) \leq 2^{-(n+1)}$ such that $X'_{r,A}[B \times U_s]$ is a Peano continuum for each $B \in \mathcal{B}_{s,A}$. Then $\mathcal{A}_s := \bigcup \{\mathcal{B}_{s,A} : A \in \mathcal{A}_r\}$ satisfies (1), (2) and (3).

Once the recursion is complete, let us write $L_n := \bigcup \{E(T_{r,A}) : r \in \mathcal{R}(n), A \in \mathcal{A}_r\}$ for the edge set of all trees on level $n \in \mathbb{N}$, and note that it follows from properties (3) and (4) that $L_n \cap L_m = \emptyset$ for all $n \neq m \in \mathbb{N}$. Thus, by defining

$$E'_1 = \bigcup_{n \in \mathbb{N}} L_{2n} \quad \text{and} \quad E'_2 = \bigcup_{n \in \mathbb{N}} L_{2n+1}$$

we obtain two disjoint edge sets of E . So it remains to check that E'_1 and E'_2 each are dense in $V \times P$ and induce a locally connected subspace of X . This will complete the proof, as then by Lemma 14.4.17, any partition $E = E_1 \sqcup E_2$ with $E_1 \supseteq E'_1$ and $E_2 \supseteq E'_2$ satisfies the assertion of the lemma.

Indeed, to see that $X[E'_1]$ is locally connected and dense, pick $(v, p) \in V \times P$ and $\delta > 0$ arbitrarily, and let $k = 2n$ large enough so that $\text{mesh}(\mathcal{A}_k) < \delta/2$ and $\text{mesh}(\mathcal{U}_k) < \delta/4$ by (1). Pick $A \in \mathcal{A}_k$ with $v \in A$ and let $U = \bigcup \{U' \in \mathcal{U}_k : p \in U'\}$. Then $\text{diam}(U) < \delta/2$ and $p \in \text{int}(U)$. By choice of $T_{r,A}$ in (4) (where $r \in \mathcal{R}(k)$ is the index of an element $U_r \subseteq U$) we have $(A \times U) \cup T_{r,A} \subseteq X[E'_1]$ is connected, of diameter at most δ , and contains at least one edge. Using Lemma 14.4.15, and the fact that δ was arbitrary, this establishes local connectedness and density for E'_1 . The case E'_2 is similar after choosing k to be odd. \square

14.4.4.4. *A decomposition theorem.* The following result combines the combinatorial techniques from Section 14.4.2 with the topological techniques from the previous Sections 14.4.3 and 14.4.4. It will be used to prove our main decomposition theorem below. Recall that ∂A denotes the boundary operator.

LEMMA 14.4.19. *Let Q_1 and Q_2 be Peano subcontinua of some non-degenerate Peano continuum P such that (a) $Q_1 \cup Q_2 = P$, (b) $Q_1 \setminus Q_2$ and $Q_2 \setminus Q_1$ are non-empty regular closed subcontinua with connected interior, and (c) $Q_1 \cap Q_2 = W = W_1 \oplus \cdots \oplus W_k$ is a finite disjoint union of regular closed Peano continua W_i each with connected interior such that $\text{int}(W)$ separates Q_1 from Q_2 .¹⁷ Then for any locally connected compactum X with dense edge set and $\mathfrak{G}(X) = V \times P$, there is a partition $E(X) = E_1 \sqcup E_2 \sqcup F$ such that*

- (1) $X[E_i]$ is locally connected, and $\partial E_i = V \times Q_i$ for $i = 1, 2$,
- (2) $|F| < \infty$,
- (3) $X[E_2]$ satisfies the even-cut condition.

PROOF. We may assume that $X[V \times W_1]$ is connected – as otherwise, by (c) and Lemma 14.4.16, there is a clopen partition \mathcal{B} of V such that $X[B \times W_1]$ is a Peano continuum for all $B \in \mathcal{B}$. Assign the finitely many cross-edges of the clopen partition associated with \mathcal{B} to F and apply the following argument to each $X[B \times P]$ individually. Hence we may find, by Lemma 14.4.11, a sparse spanning tree $T \subseteq X_{\sim}$ such that for any edge $e \in E(T)$, both its endpoints lie in $V \times W_1$. By Lemma (iii), the remaining space $X' := X[V \times P] - E(T)$ is a locally connected, metrisable compactum with a dense collection of edges.

Hence, by Lemma 14.4.16 and Theorem 14.4.18, we can partition each edge set of $X'[V \times W_i]$ into E_1^i and E_2^i such that both $(V \times W_i) \cup E_j^i$ are locally connected with E_j^i being a dense collection of edges for all $i \in [k]$ and $j \in [2]$. Let

$$E'_1 = \bigcup \{E_1^i : i \in [k]\} \cup \{e = xy \in E(X) : x \in V \times (Q_1 \setminus Q_2), y \in V \times Q_1\}$$

and

$$E'_2 = \bigcup \{E_2^i : i \in [k]\} \cup \{e = xy \in E(X) : x \in V \times (Q_2 \setminus Q_1), y \in V \times Q_2\}.$$

¹⁷For a typical example let $P = S^1$, and Q_1 a clockwise arc on P from 8 to 4 o'clock, and Q_2 a clockwise arc on P from 2 to 10 o'clock.

We claim that $\partial E'_j = V \times Q_j$ and $(V \times Q_j) \cup E'_j$ is locally connected for $j = 1, 2$. Consider the case $j = 1$ (the other case is similar). By (b) and Lemma 14.4.16, it follows that $X[V \times (Q_1 \setminus \text{int}(Q_2))]$ is locally connected. And by construction, we also have $(V \times W) \cup \bigcup \{E'_i : i \in [k]\}$ is locally connected. Hence, it follows that their union is a locally connected space with ground set $V \times Q_1$ whose edge set is a subset of E'_1 . But then it follows from Lemma 14.4.17 that we may add all remaining edges from E'_1 without harming local connectedness or density. The claim is established.

By this point, we have accounted for all edges in $E(X)$ apart from edges of T , and edges of $F := E(V \times (Q_1 \setminus Q_2), V \times (Q_2 \setminus Q_1))$. Note that F is finite: since $\text{int}(W)$ separates Q_1 from Q_2 , the sets $(Q_1 \setminus Q_2)$ and $(Q_2 \setminus Q_1)$ have positive distance from another, and so since $E(X)$ forms a zero-sequence, only edges of sufficiently large diameter can be in F .

Thus, it remains to distribute the edges of T between E'_1 and E'_2 . We will do this as to make sure that $X[E_2]$ satisfies the even-cut condition, and let $E_2 = \sum \{C_e : e \in E'_2\}$, i.e. consider the thin sum of fundamental cycles of edges in E'_2 with respect to T , Definitions 14.4.4 and 14.4.6. Note that $E'_2 \subseteq E_2 \subseteq E'_2 \cup E(T)$, so $\partial E_2 = V \times Q_2$. Moreover, since E_2 is the thin sum of circuits, it follows from Theorem 14.4.7 that $X[E_2]$ satisfies the even-cut condition. Finally, let $E_1 := E(X) \setminus (E_2 \cup D)$. Then also $E'_1 \subseteq E_1 \subseteq E'_1 \cup E(T)$, so $\partial E_1 = V \times Q_1$. Moreover, as E_1 and E_2 are supersets of E'_1 and E'_2 respectively, so both $(V \times Q_i) \cup E_i$ are locally connected by Lemma 14.4.17. \square

Recall the definition of a Peano partition from Definition 14.2.6. We can visualize the way the different elements of a partition \mathcal{U} interact by its intersection graph $G_{\mathcal{U}}$, see Definition 14.3.7. Note that if \mathcal{U} is a finite cover of a Peano continuum X , it follows from the connectedness of X that $G_{\mathcal{U}}$ is a finite connected graph.

LEMMA 14.4.20. *Let \mathcal{U} be a finite Peano partition of a connected set X , $G_{\mathcal{U}}$ its associated intersection graph, and $U \in \mathcal{U}$. If we denote by $N(U)$ all neighbours of U in $G_{\mathcal{U}}$, then U and $\bigcup V(G_{\mathcal{U}}) \setminus (U \cup N(U))$ are disjoint closed sets in X , and therefore have some positive distance.*

PROOF. They are disjoint by the definition of intersection graph and neighborhood, and they are closed as a finite union of closed sets. \square

THEOREM 14.4.21 (Decomposition Theorem). *For every $\varepsilon > 0$ and every Peano continuum P , there exists a finite cover $\mathcal{P} = \{P_1, \dots, P_k\}$ of P consisting of Peano subcontinua with $\text{mesh}(\mathcal{P}) < \varepsilon$ such that every locally connected compactum $X = (V \times P) \cup E$ admits a finite partition $E = E_1 \sqcup \dots \sqcup E_k \sqcup F$ such that*

- (1) $|F| < \infty$,
- (2) $\partial E_i = V \times P_i$,
- (3) $X_i := X[E_i]$ is locally connected for all $i \in [k]$,
- (4) X_i satisfies the even-cut condition for all $i \neq 1$.

Note that while $\{P_1, \dots, P_k\}$ is not a Peano partition of P , but only a cover (i.e. $P_i \cap P_j$ may have non-empty interior), the resulting tiles $\{X_1, \dots, X_k\}$ of the decomposition theorem together with the finitely many edges from F do form a Peano partition of X : for all these tiles and edges are edge-disjoint, and as the edges of X are dense, this means they all have pairwise disjoint interiors.

PROOF. Suppose for a contradiction that the statement is false for some $\varepsilon > 0$, and consider the class \mathcal{C} of all Peano continua that witness the failure of ε . For each $P \in \mathcal{C}$ let $k_P \in \mathbb{N}$ denote the minimum size over all $\varepsilon/3$ Peano partitions of P , and fix $P \in \mathcal{C}$ such that $k = k_P$ is minimal. Let \mathcal{U} be a $\varepsilon/3$ Peano partition of P with $|\mathcal{U}| = k$, which exists by Theorem 14.2.7.

Clearly, we must have $k \geq 3$, as otherwise, $\text{diam}(P) < \varepsilon$ and there is nothing to do. Now pick a spanning tree T for its associated intersection graph $G = G_{\mathcal{U}}$ (see Definition 14.3.7), and let U be a leaf of this tree, and denote by $N_G(U)$ the neighbourhood of U in $G_{\mathcal{U}}$. Set $P' := U \cup \bigcup N(U)$ and $P'' = \bigcup V(T) \setminus \{U\}$. Since U was a leaf of T , the induced subgraph $G_{\mathcal{U}} - \{U\}$ is connected, P' and P'' are both Peano subcontinua of P together covering P such that $\text{int}(P' \cap P'') = \text{int}(\bigcup N(v))$ consists of finitely many Peano subcontinua of P separating P' from P'' , see Lemma 14.4.20. Also note that $\text{diam}(P') \leq \varepsilon$.

Further, note that $\mathcal{U}' := \mathcal{U} \setminus \{U\}$ is an $\varepsilon/3$ Peano partition for the Peano continuum P'' . By minimality of k_P , it follows that $P'' \notin \mathcal{C}$ and so there is a finite cover $\mathcal{Q} = \{P_1, \dots, P_\ell\}$ of P'' satisfying the conclusion of the theorem. To obtain the final contradiction, we show that the finite cover $\mathcal{P} = \{P_1, \dots, P_\ell, P'\}$ of P witnesses that P could not have been a counterexample. Clearly, $\text{mesh}(P) < \varepsilon$.

To see the other assertions, consider an arbitrary locally connected compactum $X = (V \times P) \cup E$ with V compact zero-dimensional and the collection of free arcs E being dense. By construction of P' and P'' we may apply Lemma 14.4.19 to find a partition $E = E' \sqcup E'' \sqcup F'$ of the edge set E of X such that

- $\partial E' = V \times P'$ and $X' = (V \times P') \cup E'$ is locally connected and satisfies the even-cut condition,
- $\partial E'' = V \times P''$ and $(V \times P'') \cup E''$ is locally connected, and
- $|F'| < \infty$.

Next, by the assumptions on the cover \mathcal{Q} of P'' , we may find a further partition $E'' = E_1 \sqcup \dots \sqcup E_\ell \sqcup F''$ such that

- $|F''| < \infty$,
- E_i is dense in $V \times P_i$,
- $X_i := V \times P_i \cup E_i$ is locally connected for all $i \in [\ell]$,
- X_i satisfies the even-cut condition for all $i \neq 1$.

But then we see that the edge partition $E = E_1 \sqcup \dots \sqcup E_\ell \sqcup E' \sqcup F$ for $F := F' \cup F''$ witnesses that \mathcal{P} does satisfy the assertion of the theorem after all. \square

14.4.5. Approximating Sequences of Eulerian Decompositions.

14.4.5.1. *Covering the ground-set by tiles.* The plan is now to apply the decomposition Theorem 14.4.21 recursively, in order to construct an approximating sequence of Eulerian decompositions for X as in Theorem 14.3.12. So let us fix a Peano graph X with ground space $\mathfrak{G}(X) = V \times P$ and edge set $E = E(X)$ throughout this section, satisfying the blanket assumptions of Section 14.4.1.1 explained at the beginning of this chapter.

First, we recursively construct a sequence $(\mathcal{P}_n: n \in \mathbb{N})$ of finite covers of P and a locally finite tree \mathcal{R} with levels $\mathcal{R}(n)$ such that for all $n \in \mathbb{N}$ we have

- (COVER) (a) $\mathcal{P}_0 = \{P\} = \{P_r\}$ for $\{r\} = \mathcal{R}(0)$ the root of \mathcal{R} ,
 (b) $\text{mesh}(\mathcal{P}_n) \leq 2^{-n}$, and
 (c) $\mathcal{P}_{n+1} \preceq \mathcal{P}_n$ witnessed by the *refinement tree* \mathcal{R} , i.e. for all $r < r'$ with $r \in \mathcal{R}(n)$ and $r' \in \mathcal{R}(n+1)$ we have $P_r \in \mathcal{P}_n$, $P_{r'} \in \mathcal{P}_{n+1}$ and $P_r \subseteq P_{r'}$,
 (d) For $r \in \mathcal{R}(n)$ writing $r^+ := \{s \in \mathcal{R}(n+1): r < s\}$, we have that $\{P_s: s \in r^+\}$ is a finite cover of P_r satisfying the assertions of Theorem 14.4.21 for P_r .

The base case is given in (a). Now whenever \mathcal{P}_n is already constructed, pick for each $Q \in \mathcal{P}_n$ a cover \mathcal{P}_Q of $\text{mesh}(\mathcal{P}_Q) \leq 2^{-(n+1)}$ according to the Decomposition Theorem 14.4.21 for Q , and let $\mathcal{P}_{n+1} := \bigcup \{\mathcal{P}_Q: Q \in \mathcal{P}_n\}$. Moreover let $\mathcal{R} = (R, \leq)$ be the corresponding *refinement tree*, that is $\mathcal{R}(n)$, the n th level of \mathcal{R} , indexes the elements of \mathcal{P}_n , so $\mathcal{P}_n = \{P_r: r \in \mathcal{R}(n)\}$, and $r < r'$ for $r \in \mathcal{R}(n)$ and $r' \in \mathcal{R}(n+1)$ if and only if $P_r \in \mathcal{P}_{r'}$.

To formulate our next properties, we use the following piece of notation: if $r \in \mathcal{R}(n)$, then r^- denotes the unique node in $\mathcal{R}(n-1)$ with $r^- < r$. In fact, note that \mathcal{R} embeds into the tree $\mathbb{N}^{<\mathbb{N}}$ of finite natural sequence ordered by extension. Thus, without loss of generality, we assume from now on that $\mathcal{R} \subseteq \mathbb{N}^{<\mathbb{N}}$. In particular, the root of \mathcal{R} will be denoted by \emptyset , each level $\mathcal{R}(n) = \mathcal{R} \cap \mathbb{N}^n$ consists of the n -element sequences in \mathcal{R} , and for every $r \in \mathcal{R}$ we may assume that $r^+ = \{r \frown 0, r \frown 1, \dots, r \frown k_r\}$ for some suitable $k_r \in \mathbb{N}$, with $r \frown i$ denoting the extension of the finite sequence r by a new last element i .

We now construct by recursion on $n \in \mathbb{N}$

- a family $\{\mathcal{A}_r: r \in \mathcal{R}(n)\}$ of (finite) clopen partitions of V ,
- a family $\{E_{r,A}: r \in \mathcal{R}(n), A \in \mathcal{A}_r\}$ of pairwise disjoint subsets of E , and
- a family $\{F_{r,A}: r \in \mathcal{R}(n), A \in \mathcal{A}_r\}$ of pairwise disjoint, finite subsets of E ,

such that for all $r \in \mathcal{R}$ the following holds:

- (CUT) (a) $\mathcal{A}_r = \{V\}$ for the unique node $r \in \mathcal{R}(0)$,
 (b) $\text{mesh}(\mathcal{A}_r) \leq 2^{-n}$ for all $r \in \mathcal{R}(n)$,
 (c) $r \leq r' \in \mathcal{R}$ implies $\mathcal{A}_r \succcurlyeq \mathcal{A}_{r'}$,
- (EDGE) (a) $E_{r,V} = E$ for the unique node $r \in \mathcal{R}(0)$,
 (b) $E_{r,A} = F_{r,A} \sqcup \bigsqcup \{E_{s,A'}: s \in r^+, A' \in \mathcal{A}_s\}$ for all $A \in \mathcal{A}_r$,
- (TILE) (a) $X_{r,A} = X[E_{r,A}]$ is a Peano graph with $\mathfrak{G}(X_{r,A}) = A \times P_r$ for all $A \in \mathcal{A}_r$,
 (b) all tiles $X_{A,s}$ for $s \in r^+ \setminus \{r \frown 0\}$ and $A \in \mathcal{A}_s$ satisfy the even-cut condition.

CONSTRUCTION. By recursion on $n \in \mathbb{N}$. The base case is clear as for the unique node $r \in \mathcal{R}(0)$ we have $X = (V \times P) \cup E = (A \times P_r) \cup E_{r,A} = X_{r,A}$ for $A \in \mathcal{A}_r = \{V\}$. Now suppose the construction has progressed up to some tile $X_{r,A}$ with $r \in \mathcal{R}(n)$ and $A \in \mathcal{A}_r$, which is a Peano graph with ground space $A \times P_r$ by TILE(a). By Corollary 14.4.14 there is a (finite) clopen partition \mathcal{B}_A of A with $\text{mesh}(\mathcal{B}_A) \leq 2^{-(n+1)}$ such that $X_{r,B} = X_{r,A}[B \times P_r]$ is a Peano graph with ground space $B \times P_r$ for each $B \in \mathcal{B}_A$. Let $F(\mathcal{B}_A)$ denote the finite set of cross-edges the clopen partition \mathcal{B}_A induces in $X_{r,A}$.

By property COVER(d) for P_r , the Decomposition Theorem 14.4.21 applied to $X_{r,B}$ returns a finite partition

$$E_{r,B} = E_{r \frown 0, B} \sqcup \cdots \sqcup E_{r \frown k_r, B} \sqcup F_{r,B}$$

so that the corresponding tiles $Y_{i,B} := (B \times P_{r \frown i}) \cup E_{r \frown i, B}$ are locally connected with a dense collection of edges for all $i \leq k_r$, and so that $Y_{i,B}$ satisfies the even-cut condition for all $i \neq 0$. By Lemma 14.4.12, for each $Y_{i,B}$ there is a (finite) clopen partition $\mathcal{A}_{r \frown i, B}$ of B so that $Y_{i,B} = \bigoplus_{A' \in \mathcal{A}_{r \frown i, B}} X_{r \frown i, A'}$ where $X_{r \frown i, A'} \subseteq Y_{i,B}$ is a standard Peano graph with ground space $\mathfrak{G}(X_{r \frown i, A'}) = A' \times P_{r \frown i}$ and edge set say $E_{r \frown i, A'}$, giving TILE(a), and $F_{r,A} = F(\mathcal{B}_A) \cup \bigcup_{B \in \mathcal{B}_A} F_{r,B}$ is finite, satisfying EDGE(b). Further, by the moreover-part of Lemma 14.4.12, each $X_{r \frown i, A'}$ for $A \in \mathcal{A}_{r \frown i, B}$ with $i \neq 0$ satisfies the even cut condition, giving TILE(b). Now for each $i \leq k_r$ define $\mathcal{A}_{r \frown i} = \bigcup_{A \in \mathcal{A}_r} \bigcup_{B \in \mathcal{B}_A} \mathcal{A}_{r \frown i, B}$, which is a (finite) clopen partition of V satisfying CUT(b) and (c). Then by construction, for all $A' \in \mathcal{A}_{r \frown i}$ we have $X_{r \frown i, A'} = X[E_{r \frown i, A'}]$ is a Peano graph. The construction is complete. \square

We need the following elementary results, the proofs of which are evident.

LEMMA 14.4.22. $X = \bigcup_{i \in [n]} X_i$. Then X satisfies the even-cut condition if and only if each X_i satisfies the even-cut condition. \square

LEMMA 14.4.23. Let Z be a compact graph-like space satisfying the even-cut condition. Suppose that $E(Z) = E_0 \sqcup \cdots \sqcup E_k$ such that $Z[E_i]$ satisfies the even-cut condition for all $1 \leq i \leq k$. Then also $Z[E_0]$ satisfies the even-cut condition. \square

For $k \in \mathbb{N}$ and $s \in \mathbb{N}^{\mathbb{N}}$, write $s \frown 0^k := s \frown \underbrace{0 \frown 0 \frown \cdots \frown 0}_{k \text{ times}}$. When using this notation, we usually require that s does not end on 0.

For $r \in \mathcal{R}$, let $E_r := \bigcup_{A \in \mathcal{A}_r} E_{r,A}$, $F_r := \bigcup_{A \in \mathcal{A}_r} F_{r,A}$ and $X_r = X[E_r]$. Then $X_r = \bigoplus_{A \in \mathcal{A}_r} X_{r,A}$, and hence it follows by property TILE(b) and Lemma 14.4.22 that whenever r does not end on 0, then X_r satisfies the even cut condition.

The following simple observation is the key for constructing an Eulerian decomposition.

LEMMA 14.4.24. For every $t \in \mathbb{N}^{\mathbb{N}}$, and s not ending on 0 with $t = s \frown 0^n$, the graph-like space $Z_t := X_{\sim}[E_t \sqcup \bigsqcup_{k=0}^{n-1} F_{s \frown 0^k}]$ has the even-cut property.

PROOF. First, if $n = 0$, then $Z_t = X_{\sim}[E_t]$ has the even-cut property by assumption if $t = \emptyset$, and otherwise by TILE(b) and Lemma 14.4.22. Now consider $t = s \frown 0^{n+1}$, let $r = s \frown 0^n$ and assume inductively that Z_r has the even-cut property. Recall that by EDGE(b), we have $E_r = F_r \sqcup \bigsqcup \{E_s : s \in r^+\}$. Since each $s \neq r \frown 0$ has the even-cut property, it follows from Lemma 14.4.23 that also the complement of these sets in Z_r has the even-cut property. But clearly, the edge-complement of $\{E_s : s \in r^+\}$ is precisely Z_t . \square

14.4.5.2. *Three auxiliary graphs.* To build an approximating sequence of Eulerian decompositions, we will now construct suitable Eulerian multi-graphs (G_n, η_n) approximating the decomposition constructed above in TILE(a). We will do this in three stages reminiscent of the steps in the blueprint from Observation 14.3.8.

- First, construct a sequence of auxiliary multi-graphs $(G'_n : n \in \mathbb{N})$ each living on the tiles at stage n and has as edge set F_n of all remaining edges of X at stage n .
- Second, we form a sequence of even¹⁸ multi-graphs $(G''_n : n \in \mathbb{N})$, where each G''_n is a supergraph of G'_n formed by adding some type-E dummy edges. This step is the critical part of the argument, relying on the even-cut properties in TILE(b).
- Finally, form a sequence of even, connected multi-graphs $(G_n : n \in \mathbb{N})$, where each G_n is a super-graph of G''_n formed by adding some type-C dummy edges to G''_n ,¹⁹

making sure in all steps that we always have compatible inverse limits $\varprojlim G'_n \hookrightarrow \varprojlim G''_n \hookrightarrow \varprojlim G_n$, each with contraction maps (Definition 14.3.13) as bonding maps. The reader may picture this process as in the following two figures, Figures 14.7 and 14.8.




FIGURE 14.7. A sketch of $E_0 = E_0 \sqcup F_0 \sqcup E_1$ and the corresponding tiles on the left. On the right, the first auxiliary graph G'_1 with edge set F_0 .

¹⁸A finite graph is called *even* if all its vertices have even degree.

¹⁹The purpose of *type-E edges* will be to make all degrees of G_{n+1} even, and the purpose of *type-C edges* is to make G_{n+1} connected.

²⁰We remark that for ease of formalisation, our algorithm will add additional type-C edges not drawn in this picture.




FIGURE 14.8. Type-E dummy edges in blue turn G'_1 into an even graph, with their η_1 images drawn as dotted arcs. Type-C dummy edges in green make G_1 connected, with their common η_1 image being a trivial arc.²⁰

Building the first auxiliary graph. For every $n \in \mathbb{N}$ we recursively construct decompositions (G'_n, η'_n) with G'_n a finite multi-graph encoding the edge patterns between the tiles at step n . So formally, the graph G'_n has vertex set V_n and edge set F_n where

- $V_n = \{v_{r,A} : r \in \mathcal{R}(n), A \in \mathcal{A}_r\}$ and
- $F_n := \bigcup \{F_{r,A} : r \in \mathcal{R}(<n), A \in \mathcal{A}_r\}$.²¹

and η'_n is defined by $\eta'_n \upharpoonright F_n = \text{id}$ and $\eta_n(v_{r,A}) := X_{r,A}$ for all vertices in V_n . Note that on our way to build a decomposition, (G'_n, η_n) satisfies (E1), (E2), (E4) and (E5) of a decomposition according to Definition 14.3.2. Edge-vertex incidence in G'_n is defined recursively in n ²² so as to satisfy (E6) and Definition 14.3.4 for F_n . For this, observe that for every $n \in \mathbb{N}$ there is a natural (surjective) contraction map

$$\varrho'_n : G'_{n+1} \rightarrow G'_n, v_{r,A} \mapsto v_{r-,A'} \text{ and } f \mapsto \begin{cases} f & \text{if } f \in F_n, \\ v_{r,A} & \text{if } f \in F_{n+1} \setminus F_n, f \in F_{r,A}. \end{cases}$$

which clearly corresponds to the relation $X_{r,A} \subseteq X_{r-,A'}$ where A' is the unique element of \mathcal{A}_{r-} satisfying $A' \supseteq A$. Indeed, it is straightforward to check that properties (Q1) – (Q4) in Definition 14.3.13 for a contraction map are satisfied.

Since G'_0 is the unique edge-less graph on a single vertex, there is nothing to do. Suppose that G'_n has already been defined so that (E6) and Definition 14.3.4 are satisfied for the finite sequence $(G'_i : i \leq n)$. Consider $f \in E(G'_{n+1}) = F_{n+1}$. If $f \in F_n$, and say $f_{G'_n}(0) = v_{r,A}$ for some $r \in \mathcal{R}(n)$ and $A \in \mathcal{A}_r$, then by our recursive assumptions we have $f(0) := (x, y) \in A \times P_r$. Choose any $s \in r^+$ such that $y \in P_s \subseteq P_r$ and let A' be the unique

²¹ F_n should not be confused with $F_{(n)}$ where (n) is a one-element sequence on the first level of \mathcal{R} .

²²If one such displayed free arc $f \in F_n$ has an endpoint $(x, y) \in V \times P$ in X , then all vertices $v_{r,A} \in V_n$ with $y \in P_r$ and $x \in A$ are potential candidates for the corresponding endvertex of f in G'_n . This is where we make a recursive choice.

element of \mathcal{A}_s satisfying $A' \subseteq A$, and define $f_{G'_{n+1}}(0) = v_{s,A'}$. Similarly, if $f \in F_{n+1} \setminus F_n$, i.e. $f \in F_{r,A}$ for some $r \in \mathcal{R}(n)$ and $A \in \mathcal{A}_r$, then if say $f(0) := (x, y) \in V \times P$ choose any $s \in r^+$ such that $y \in P_s$ and let A' be the unique element of \mathcal{A}_s satisfying $A' \subseteq A$, and define $f_{G'_{n+1}}(0) = v_{s,A'}$, and similarly for $f(1) := (x', y') \in V \times P$.

Summary: Each $\mathcal{D}'_n = (G'_n, \eta'_n)$ forms a decomposition of X (cf. Definition 14.3.2), and $\varrho'_n: G'_{n+1} \rightarrow G'_n$ is an η -compatible contraction map (cf. Definition 14.3.13 and 14.3.14).

Building the second auxiliary graph. For our second auxiliary graph G''_n , for each edge e of G'_n , we will add two corresponding *type-E dummy edges* $d^{e(0)}$ and $d^{e(1)}$ to G'_n , making sure that (E3) and (E7) are satisfied for each (G''_n, η''_n) . We also make sure that ϱ'_n extends to a contraction map $\varrho''_n: G''_{n+1} \rightarrow G''_n$.

DEFINITION 14.4.25. For $e \in E(X)$, write $e(i) = (x_{e(i)}, y_{e(i)}) \in V \times P$ for its endpoints $e(0)$ and $e(1)$ in X . For every $e \in E(X)$, there is a unique index $m = m(e)$ such that $e \in F_{m+1} \setminus F_m$, and so there is a unique $s = s^e \in \mathcal{R}(m)$ such that $e \in E_{s,A}$ for some $A \in \mathcal{A}_s$. For every $k \geq m$, let $s^e(k) = s \frown 0^{k-m} \in \mathcal{R}(k)$. Note that for every edge e , the set $\{P_{s^e(k)} : k \geq m(e)\}$ is a nested zero-sequence of subcontinua of P , and hence there is a unique point contained in the intersection $\bigcap_{k \geq m(e)} P_{s^e(k)}$ which we denote by $\sigma(e)$. Further, for $k \geq m$ and $i \in \{0, 1\}$, let $A^{e(i)}(k) \in \mathcal{A}_{s^e(k)}$ be the unique element with $x_{e(i)} \in A^{e(i)}(k)$. For $e \in E$, and $k > m(e)$ we write $v^{e(i)}(k) := v_{s^e(k), A^{e(i)}(k)} \in V_k$, and call this vertex the *root vertex associated with the endpoint $e(i)$ at stage k* . Finally, fix arcs $\alpha^{e(i)} \subseteq \{x_{e(i)}\} \times P_{s^e}$ from $e(i) = (x_{e(i)}, y_{e(i)})$ to $(x_{e(i)}, \sigma(e))$ for each $e \in F_n$ and $i \in \{0, 1\}$.

Define (G''_n, η''_n) by adding to G'_n a set of dummy edges $D''_n = \{d^{e(0)}, d^{e(1)} : e \in F_n\}$, and extend η'_n to a map η''_n by defining $\eta''_n(d^{e(i)}) = \alpha^{e(i)}$ on the newly added dummy edges. By construction of the arcs α , this assignment satisfies (E7) for η''_n . Further, edge-vertex incidence for type-E dummy edges in G''_n is given by $d^{e(i)}_{G''_n}(0) := e_{G''_n}(i)$ and $d^{e(i)}_{G''_n}(1) := v^{e(i)}(n)$, that is to say, the edge $d^{e(i)}$ connects an endpoint of e in G_n to the root vertex associated with the endpoint at stage n .

Moreover, we extend the map ϱ'_n to a contraction map $\varrho''_n: G''_{n+1} \rightarrow G''_n$ by defining

$$\varrho''_n(d^{e(i)}) = \begin{cases} \varrho'(e) & \text{if } d^{e(i)} \in D_{n+1} \setminus D_n \\ d^{e(i)} & \text{if } d^{e(i)} \in D_n. \end{cases}$$

THEOREM 14.4.26. *Each G''_{n+1} is an even multi-graph, $\mathcal{D}''_n = (G''_n, \eta''_n)$ forms a decomposition of X , and $\varrho''_n: G''_{n+1} \rightarrow G''_n$ is an η -compatible contraction map.*

PROOF. It is routine to check that $\mathcal{D}''_n = (G''_n, \eta''_n)$ forms a decomposition of X . Moreover, the map $\varrho''_n: G''_{n+1} \rightarrow G''_n$ is a contraction map, because we added new dummy edges only between vertices in the same fibre of ϱ' . Hence, (Q4) of a contraction map is still satisfied, and the other properties are inherited from ϱ'_n . To see that G''_{n+1} is even, we make use of the following observation, which is immediate from the construction.

Observation: For every $n \in \mathbb{N}$, the edge set of G''_n can be partitioned into a family of edge-disjoint trails²³ $\{T_n(e) : e \in F_n\}$ whose vertex-edge sequence is given by

$$T_n(e) = v^{e(0)}(n), d^{e(0)}, e_{G_n}(0), e, e_{G_n}(1), d^{e(1)}, v^{e(1)}(n).$$

We are now ready to calculate the parity of vertex degrees in G''_n , relying on the elementary fact that every inner vertex of a trail T has even degree in the subgraph induced by T , and every end-vertex of an open trail T (i.e. a trail with distinct start and end-vertices) has odd degree in the subgraph induced by T . So consider some vertex $v = v_{t,A} \in V(G_n)$. Write $t = s \frown 0^j$ where s does not end on zero and $j \in \mathbb{N}$. By Lemma 14.4.24, A induces an even edge cut C in $Z_t := X_{\sim}[E_t \sqcup \bigsqcup_{k=0}^{j-1} F_{s \frown 0^k}]$. Furthermore, since $X_{t,A}$ with ground set $A \times P_t$ is a connected component of $X[E_t]$, it follows that $C \subseteq \bigsqcup_{k=0}^{j-1} F_{s \frown 0^k}$.

Claim: The vertex v has odd degree in $T_n(e)$ if and only if $e \in C$.

The claim implies the theorem, since the number of trails in which v has odd degree is even. To prove the claim, note that $e \in C$ if and only if $x_{e(0)} \in A$ and $x_{e(1)} \notin A$ (or vice versa), which happens – since $C \subseteq \bigsqcup_{k=0}^{j-1} F_{s \frown 0^k}$ – if and only if $v^{e(0)}(n) = v$ and $v^{e(1)}(n) \neq v$ (or vice versa), i.e. if and only if v has odd degree in $T_n(e)$. □

Building the Eulerian decompositions. To build Eulerian (i.e. even and connected) graphs G_n from G''_n so that the maps ϱ_n become edge-contractions, it now suffices to recursively add further dummy edges to G''_{n+1} only between vertices of the same fibre $\varrho''_{n-1}(v)$ such that every such fibre becomes connected. By induction, this will imply that each G_n is connected.

The Eulerian decompositions (G_n, η_n) are built recursively. Since $2^0 = 1$, both $G_0 = G''_0$ are the unique graph on a single vertex without loops. Now suppose G_n has already been defined. Assume inductively that

- (‡1) every dummy edge $d = v_{t,A}v_{t',A'} \in E(G_n) \setminus E(G''_n)$ has an associated point $\eta(d) = (q_V(d), q_P(d)) \in V \times P$ which is contained in the intersection of the corresponding tiles $X_{t,A} \cap X_{t',A'}$.
- (‡2) Moreover, assume there is an equivalence relation \sim on the dummy edges in $E(G_n) \setminus E(G''_n)$ such that every equivalence class consists of precisely two dummy edges which are parallel in G_n .

To build G_{n+1} , first obtain a graph G_{n+1}^{**} by displaying all dummy edges of G_n such that (‡1) and (‡2) are satisfied, and so that $\varrho_n : G_{n+1}^{**} \rightarrow G''_n$ is a contraction map (when ambiguous, make an arbitrary choice). Note in particular that (‡2) and the fact that G''_{n+1} was even imply that G_{n+1}^{**} is an even graph.

To obtain a connected even graph G_{n+1} from G_{n+1}^{**} , first of all, for each $r \in \mathcal{R}(n)$, let us pick a spanning tree S_r for the intersection graph formed by the cover $\{P_{r'} : r' \in r^+\}$

²³Recall that a *trail* is a walk without repeated edges.

on P_r . Next, for each edge $P_s P_{s'}$ of S_r fix an arbitrary point $y_{ss'} \in P_s \cap P_{s'}$. We now add type-C *dummy edges* to G_{n+1}^{**} according to the following rule:

- (C) Fix a vertex $v_{r,A} \in V_n$ with $A \in \mathcal{A}_r$. Let \mathcal{B} denote the finite partition of V which is the least common refinement of the family $\{\mathcal{A}_{r'} : r' \in r^+\}$. Pick a vertex $x_B \in B$ for each $B \in \mathcal{B}$. Now for every x_B and every edge $P_s P_{s'}$ of S_r , we add two parallel type-C dummy edges $d_1 \sim d_2$ with the same associated point $\eta_{n+1}(d_1) = \eta_{n+1}(d_2) := (x_B, y_{ss'}) \in V \times P$ between the two vertices v_{s,A_s} and $v_{s',A_{s'}}$ where A_s and $A_{s'}$ are the unique elements of \mathcal{A}_s and $\mathcal{A}_{s'}$ respectively with $B \subseteq A_s$ and $B \subseteq A_{s'}$. Finally, we extend the map ϱ_n to these newly inserted edges by defining $\varrho_n(d_1) = \varrho_n(d_2) := v_{r,A}$. This arrangement for d_1 and d_2 satisfies $(\dagger 1)$ and $(\dagger 2)$.

THEOREM 14.4.27. *Each G_{n+1} is a finite Eulerian multi-graph, $\mathcal{D}_n = (G_n, \eta_n)$ is an Euler decomposition of X , and $\varrho_n : G_{n+1} \rightarrow G_n$ is an η -compatible edge-contraction map. Thus, $(\mathcal{D}_n : n \in \mathbb{N})$ is an approximating sequence of Eulerian decompositions for X .*

PROOF. We first show that $\varrho_n : G_{n+1} \rightarrow G_n$ is an edge-contraction map, i.e. that it has connected fibres, see $(Q5)$ of Definition 14.3.13. Interpreted as a continuous map, this translates to the fact that ϱ_n is monotone. In particular, this will imply inductively that each G_n is connected: Indeed, G_0 is trivially connected, and if G_n is connected, then it follows from the fact that since $\varrho_n : G_{n+1} \rightarrow G_n$ is a continuous, monotone surjective map from a compact spaces onto a connected space, then also the domain G_{n+1} must be connected, see e.g. [63, Theorem 6.1.29].

To see that ϱ_n has connected fibres, fix some $v_{r,A} \in V_n$, and consider $H := \varrho_n^{-1}(v_{r,A})$, a subgraph of G_{n+1} . By definition, the vertex set of H is precisely the set

$$V_H = \{v_{s,A'} : s \in r^+, A' \in \mathcal{A}_s\}.$$

Let $C \subseteq V_H$ be the vertex set of a component of the graph H . We have to show $C = V_H$. For this, note that if $v_{s,A'} \in C$ and $v_{t,A''} \in V_H$ with $A' \cap A'' \neq \emptyset$, then $v_{t,A''} \in C$. Indeed, let $P \subseteq S_r$ denote the unique $P_s P_t$ path in the tree S_r . Fix $x_B \in B \subseteq A' \cap A''$. Then the dummy edges in $\eta_n^{-1}(\{(x_B, y_{uu'}) : uu' \in E(P)\}) \subseteq H$ which have been added according to rule (C) witness connectivity between $v_{s,A'}$ and $v_{t,A''}$.

Therefore,

$$A_C := \bigcup \{A' : v_{s,A'} \in C\} \quad \text{and} \quad A_{-C} := \bigcup \{A' : v_{s,A'} \in V_H \setminus C\}$$

gives rise to a clopen bipartition (A_C, A_{-C}) of A . We claim that $A_{-C} = \emptyset$. This would imply that $C = V_H$, proving that H is connected. Otherwise, (A_C, A_{-C}) is a non-trivial clopen bipartition of A , and so since $X_{r,A} = (A \times P_r) \cup E_{r,A}$ was a Peano continuum by (TILE)(a), it follows that $E_{r,A}(A_C, A_{-C})$ is a non-empty edge cut of $X_{r,A}$. Pick f in $E_{r,A}(A_C, A_{-C})$ arbitrarily. Then $f \in F_{r,A} \subseteq F_{n+1}$ by (EDGE)(b), and hence $f \in E(H)$.

However, it now follows from (E6) that $f \in E_H(C, V_H \setminus C)$, witnessing that C was not maximally connected, a contradiction.

That the ϱ_n are η -compatible is easily verified, and so it follows from Lemma 14.3.15 that $(\mathcal{D}_n: n \in \mathbb{N})$ is indeed an approximating sequence of Eulerian decompositions for X . Note that $w((D_n, \eta_n)) \rightarrow 0$ follows from COVER(b), CUT(b), and the fact that we assumed that X contained no loops, implying that $\text{diam}(X_{r,A}) \rightarrow 0$ as $|r| \rightarrow \infty$. \square

The proof of our main result is now complete:

PROOF OF THEOREM 14.4.1. Let X be a Peano continuum with $\mathfrak{G}(X) = V \times P$. We may assume that X is a Peano graph without loops with the even-cut property, such that P is non-trivial. Then by Theorem 14.4.27, the space X has an approximating sequence of Eulerian decompositions, and hence X is Eulerian by Theorem 14.1.1. \square

14.5. One-dimensional spaces

14.5.1. Overview. The purpose of this final chapter is to prove the following theorem.

THEOREM 14.5.1. *A one-dimensional Peano continuum is Eulerian if and only if it satisfies the even-cut condition.*

More precisely, using (iii) \Rightarrow (i) of Theorem 14.1.1, what we will show here is that every one-dimensional Peano continuum satisfying the even-cut condition admits an approximating sequence of Eulerian decompositions.

Let us briefly remark that for $n \geq 1$, the dimension of a Peano continuum X is n if and only if the ground space $\mathfrak{G}(X)$ has dimension n . This is a consequence of the well-known *sum theorem for dimension*, [62, Thm. 1.5.2], by applying it to X considered as a countable union of $\mathfrak{G}(X)$ and one-cells \bar{e} for $e \in E(X)$. In particular, Theorem 14.1.4(C) is indeed equivalent to Theorem 14.5.1.

14.5.1.1. *Proof strategy.* Consider a one-dimensional Peano continuum X for which we aim to construct an approximating sequence of Eulerian decompositions. As described in the Blueprint 14.3.8, any Peano partition \mathcal{U} for X into standard subspaces gives rise to a corresponding Eulerian decomposition for X , provided that X satisfies the even-cut condition. Note that the even-cut assumption on X is a necessary one, for if \mathcal{U} displays an odd edge cut of X , then no such corresponding Eulerian decomposition can exist. Now if we could find a Peano partition \mathcal{U} such that each partition element $U \in \mathcal{U}$ individually still has the even-cut property, we could continue this procedure recursively to construct an extending sequence of Eulerian decompositions (cf. Definition 14.3.4).

Recall, however, that there is a second objective for constructing an approximating sequence of Eulerian decompositions: Not only should the Eulerian decompositions extend each other (property (A1) of Definition 14.3.5), but their widths should also decrease to

zero (property (A2) of Definition 14.3.5). This second requirement, however, is at odds with our earlier idea that partition elements of \mathcal{U} individually always continue to have the even-cut property, as the even-cut property generally prohibits single edges to be displayed (cf. Blueprint 14.3.8), and so the width of our recursively constructed decompositions will be bounded from below by the diameter of the largest edge.

We resolve these issues by the following approach: given X , we construct in Theorem 14.5.20 a Peano partition \mathcal{U} into standard subspaces of X such that each partition element $U \in \mathcal{U}$ individually still has the even-cut property, and so that each U contains a finite set of edges F_U such that each component of $U - F_U$ has somewhat smaller diameter than X . Then the partition \mathcal{U}' consisting of the components of $U - F_U$ for $U \in \mathcal{U}$ and individual edges in $\bigcup_{U \in \mathcal{U}} F_U$ gives rise to an Eulerian decomposition of smaller width as desired. And the fact that each U satisfied the even-cut condition leaves enough traces in $U - F_U$ (almost all vertices of $U - F_U$ have even degree) so that we may continue the recursive construction, see Theorem 14.5.4.

Before we come to these results, we gather in Section 14.5.2 a number of auxiliary results whose purpose is first to set up the language for arranging the even-cut property in terms of inverse limits, and second to deal with the fact that edges of some partition element $U \in \mathcal{U}$ are not a priori edges of X , which requires us to generalise our concept of ground space and edges.

14.5.2. Admissible Vertex Sets and Combinatorial Alignment.

14.5.2.1. *Admissible vertex sets.* In the introduction, we stated in Sections 14.1.1.2 and 14.1.3.1 the even-cut condition for the class of Peano continua X in terms of their ground spaces $\mathfrak{G}(X)$. For this chapter we generalise these notions in two directions: first, we generalise the notion of ground space to that of *admissible vertex sets*, and second we extend the class of spaces X we consider from Peano continua to a broad class of (metrisable) compacta – which we call *component-wise aligned compacta*.

To justify our first generalisation, recall that there is a standard fuzziness in the transition between combinatorial and topological graphs in the sense that degree-two vertices in combinatorial graphs are disregarded in the corresponding topological graph. This fuzziness is even more pronounced in the case of graph-like spaces: note that for example, both $V = \{0, 1\}$ and V the middle third Cantor set can function as vertex set of a graph-like continuum homeomorphic to the unit interval I . In this chapter, we set up the language for eliminating this imprecision, for the following reason: if $H = (V_H, E_H)$ and $G = (V_G, E_G)$ are combinatorial graphs such that H is a subgraph of G , then their combinatorial structures are naturally aligned in the sense that $V_H \subseteq V_G$ and $E_H \subseteq E_G$. However, viewing H and G as topological spaces, the free arcs of H might be strict supersets of the free arcs of G , with the undesirable consequence that $E(H)$ might not be a subset of $E(G)$.

DEFINITION 14.5.2 (Admissible vertex set). A compact subset $V \subseteq X$ of a Peano continuum X is an *admissible vertex set* provided that $\mathfrak{G}(X) \subseteq V$ and $V \setminus \mathfrak{G}(X)$ is zero-dimensional. For an admissible vertex set V , the space $X \setminus V$ is homeomorphic to a disjoint sum of open intervals, which we call the edges of X associated with V , written $E(X, V)$.

This definition is equivalent to saying that $\mathfrak{G}(X)$ is a subset of V , and for every free arc e of X , we have that \bar{e} is a graph-like space homeomorphic to an interval with zero-dimensional vertex set $(V \cap \bar{e})$.

For a Peano continuum X with admissible vertex set V , the edges $E(X, V)$ are the connected components of $X \setminus V$. Since $\mathfrak{G}(X) \subseteq V$ and V is closed, it follows that every edge is homeomorphic to an open interval. Moreover, if X is a Peano graph (so $E(X)$ is dense in X), then also (X, V) is a Peano graph in the sense that the edges $E(X, V)$ are dense in X . Moreover, we may generalise the notion of edge cuts from $(X, \mathfrak{G}(X))$ to (X, V) : an edge cut of (X, V) is the set of edges crossing a clopen partition $V = A \oplus B$. It is straightforward to check that all results about edge cuts from Section 14.1.3.1 still apply in this slightly generalised setting. Finally, we also extend Definition 14.3.1 of a standard subspace to this generalised setting, and call a subspace $Y \subseteq X$ *standard in (X, V)* if for every $e \in E(X, V)$, the fact $e \cap Y \neq \emptyset$ implies $e \subseteq Y$.

LEMMA 14.5.3. *Let $V \subseteq X$ be an admissible vertex set of a Peano continuum X . Then X satisfies the even-cut condition if and only if (X, V) does.*

PROOF. Note that the graph-like continuum $(X, V)_\sim$ is a subdivision of the graph-like continuum X_\sim (see the discussion in Section 14.5.3.1). In particular, they are homeomorphic. Thus, X has the even-cut property if and only if X_\sim is Eulerian if and only if $(X, V)_\sim$ is Eulerian if and only if (X, V) has the even-cut property, where the first and last equivalence follows from [70] (and see also the discussion leading up to Conjecture 14.1.5). \square

LEMMA 14.5.4. *If X is a Peano continuum and $V \subseteq X$ an admissible vertex set for X , then any non-trivial Peano subcontinuum $Y \subseteq X$ satisfying the even-cut condition is standard in (X, V) .*

PROOF. Note first that if Y satisfies the even-cut condition, then any free arc of Y lies on a simple closed curve of Y (cf. [70, Lemma 16]), and second, that any simple closed curve in X is necessarily a standard subspace of X (cf. [70, Lemma 5]). \square

14.5.2.2. *Combinatorial alignment.* To facilitate comparing edges across different spaces, from now on we will work with admissible vertex sets instead of ground sets.

DEFINITION 14.5.5 (Combinatorial alignment). Suppose that $Y \subseteq X$ are Peano continua, and that V_X and V_Y are admissible vertex sets for X and Y respectively. We say that (Y, V_Y) is *combinatorially aligned* in (X, V_X) if for every $e \in E(Y, V_Y)$, either $e \in E(X, V_X)$ or $e \subseteq V_X$. In this situation, write $E(Y, V_Y) = E_Y^{\text{real}} \sqcup E_Y^{\text{fake}}$ with $E_Y^{\text{real}} := E_Y \cap E(X, V_X)$

for the bipartition into real and fake edges. Finally, we say a combinatorially aligned continuum $(Y, V_Y) \subseteq (X, V_X)$ is *faithfully* aligned if $E(Y, V_Y) \subseteq E(X, V_X)$, i.e. if $E_Y^{\text{fake}} = \emptyset$.

As an example for combinatorial alignment, consider again the two simple closed curves C_1 and C_2 inside the hyperbolic tree Y from Figure 14.4 in Chapter 14.2. In both cases, the red simple closed curves enter and leave the hyperbolic boundary circle fairly often, so need to be subdivided accordingly, in order to ensure that their combinatorial structure matches up. Note further that $\mathfrak{G}(Y) \cap C_1$ is not an admissible vertex set for C_1 , as free arcs in C_1 intersect $\mathfrak{G}(Y)$ in non-trivial intervals.

LEMMA 14.5.6. *Suppose X is a Peano continuum and $V \subseteq X$ an admissible vertex set for X . Suppose $Y \subseteq X$ is a standard Peano subcontinuum. Then there is an admissible vertex set W for Y such that (Y, W) is combinatorially aligned in (X, V) .*

PROOF. Consider an edge $e \in E(Y, \mathfrak{G}(Y))$, that is to say, a free arc in Y . We show that we can subdivide \bar{e} by a compact zero-dimensional vertex set W_e such that every segment of $\bar{e} \setminus W_e$ is either an edge of (X, V) or is completely contained in V .

Consider $I_e := \{f \in E(X, V) : f \cap e \neq \emptyset\} = \{f \in E(X, V) : f \subseteq e\}$, by the fact that Y is standard in (X, V) . So I_e is a collection of disjoint open intervals on e . Define $W_e = \{e(0), e(1)\} \cup \overline{\bigcup I_e} \setminus \bigcup I_e$. It is easy to verify that W_e is as desired.

Finally, let $W := \mathfrak{G}(Y) \cup \bigcup \{W_e : e \in E(Y, \mathfrak{G}(Y))\}$. Since $\{W_e : e \in E(Y, \mathfrak{G}(Y))\}$ is a zero-sequence of closed sets all intersecting the closed set $\mathfrak{G}(Y)$, it follows from standard arguments (see, for example, the proof of [158, A.11.6]) that W is closed in Y , hence compact. By the sum theorem of dimension, [62, Thm. 1.5.2], $W \setminus \mathfrak{G}(Y) \subseteq \bigcup \{W_e : e \in E(Y, \mathfrak{G}(Y))\}$ is zero-dimensional, and so W is admissible. \square

COROLLARY 14.5.7. *Suppose X is a Peano continuum and $V \subseteq X$ an admissible vertex set for X . Suppose $Y \subseteq X$ is a non-trivial Peano subcontinuum satisfying the even-cut condition. Then there is an admissible vertex set W for Y such that (Y, W) is combinatorially aligned in (X, V) .*

PROOF. Combine Lemmas 14.5.4 and 14.5.6. \square

Finally, we prove a lemma giving a necessary condition when the even-cut condition is preserved under unions. This lemma can be seen as the dual statement to Lemma 14.1.14. A word of explanation and warning about the term ‘edge-disjoint’. Given a Peano continuum (X, V) and two combinatorially aligned subspaces (Y, V_Y) and (Z, V_Z) of X , we say that (Y, V_Y) and (Z, V_Z) are *edge-disjoint*, or more precisely *$E(X)$ -edge-disjoint*, if $E_Y^{\text{real}} \cap E_Z^{\text{real}} = \emptyset$, that is to say if each edge of (X, V) is contained in at most one of Y or Z . In particular, note it may happen that *fake* edges of Y and Z meet non-trivially.

LEMMA 14.5.8. *Let $(X_n)_{n \in \mathbb{N}}$ be a zero-sequence of non-trivial $E(P)$ -edge disjoint Peano subcontinua of a Peano continuum P such that $P = \bigcup_{n \in \mathbb{N}} X_n$. If each X_n satisfies the even-cut condition, then so does P .*

PROOF. By Corollary 14.5.7, we may assume without loss of generality that each X_n is combinatorially aligned with $(P, \mathfrak{G}(P))$. Since the X_n are pairwise $E(P)$ -edge disjoint, the sets in $\{E^{real}(X_n) : n \in \mathbb{N}\}$ are pairwise disjoint. We claim that

$$(32) \quad E(P) = \bigsqcup E^{real}(X_n).$$

Well, \supseteq is immediate from the definition of being combinatorially aligned. For the reverse direction, consider any edge $e \in E(P)$. Since $P = \bigcup_{n \in \mathbb{N}} X_n$ we may assume without loss of generality that $e \cap X_0 \neq \emptyset$, and so $e \subseteq X_0$, and so e has non-trivial intersection with an edge $e' \in E(X_0)$. But since X_0 was combinatorially aligned with P , it follows that $e = e'$.

Next, note that quite similarly, one obtains

$$(33) \quad \mathfrak{G}(X_n) \subseteq \mathfrak{G}(P)$$

for all $n \in \mathbb{N}$. Indeed, the previous argument shows that if x is an interior point of some edge $e \in E(P)$ and $x \in X_n$ then $e \in E(X_n)$.

Now in order to show that also P satisfies the even-cut condition, consider an arbitrary separation $A \oplus B$ of $\mathfrak{G}(P)$. Our task is to show that $E_P(A, B)$ is even. First, note that by (33), the separation $A \oplus B$ induces separations of $\mathfrak{G}(X_n)$ for each $n \in \mathbb{N}$. Moreover, since $|E_P(A, B)|$ is finite, it follows from (32) that there is $N \in \mathbb{N}$ such that

$$E_P(A, B) = E_{X_1}^{real}(A, B) \sqcup E_{X_2}^{real}(A, B) \sqcup \cdots \sqcup E_{X_N}^{real}(A, B).$$

Next, we claim that $E_{X_n}^{real}(A, B) = E_{X_n}(A, B)$ for all $n \in \mathbb{N}$. Indeed, since any fake edge $d \in E(X_n)$ is a subset of $\mathfrak{G}(P)$, by the property of being combinatorially aligned, it follows from d 's connectedness that d is contained completely on one side of the separation $A \oplus B$ of $\mathfrak{G}(P)$, and so $d \notin E_{X_n}(A, B)$, establishing the claim. Thus, we have

$$E_P(A, B) = E_{X_1}(A, B) \sqcup E_{X_2}(A, B) \sqcup \cdots \sqcup E_{X_N}(A, B),$$

and so $E_P(A, B)$ is the disjoint union of finitely many sets of even cardinality, and hence is an even edge cut. (Recall that by Lemma 14.5.3, the even-cut property is independent of the choice of admissible vertex sets.) Since $E_P(A, B)$ was arbitrary, we have established that P satisfies the even-cut condition. \square

14.5.2.3. *Combinatorially aligned spanning trees.* From Lemma 14.4.9 we know that in a Peano continuum X , for every zero-dimensional compact set $Y \subseteq \mathfrak{G}(X)$, there exists a standard graph-like continuum $Z \subseteq X$ with $Y \subseteq Z$. Suppose V is an admissible vertex set of X . Then the same proof shows that for every zero-dimensional compact set $Y \subseteq V$, there exists a standard graph-like continuum $Z \subseteq (X, V)$ with $Y \subseteq Z$.

A natural question is whether there also is a faithfully aligned graph-like continuum $Z = (V_Z, E_Z)$ spanning Y . To see that this is not always possible, consider a Peano graph X consisting of a dense zero-sequence of loops attached to ground space I . If $Y = \{0, 1\} \subseteq I = \mathfrak{G}(X)$ say, then it is not possible to find a graph-like continuum

$Z = (V_Z, E_Z)$ with $Y \subseteq Z$ and $E_X \subseteq E(X)$. However, if we only insist on combinatorially aligned, then the answer is in the affirmative.

LEMMA 14.5.9. *Suppose X is a Peano continuum and $V \subseteq X$ an admissible vertex set for X . For every zero-dimensional compact set $Y \subseteq V$, there exists a combinatorially aligned graph-like tree $T = (V_T, E_T)$ such that $Y \subseteq V_T$.*

PROOF. By Lemma 14.4.9, there exists at least one standard graph-like continuum in X covering Y . Take an inclusion-minimal such graph-like continuum T – by Lemma 14.4.2, this will be a standard graph-like tree. By Lemma 14.5.6, for the standard subspace T there is an admissible vertex set V_T such that (T, V_T) is combinatorially aligned with (X, V) . Note that in this case we necessarily have $Y \subseteq V_T$. \square

14.5.2.4. *Component-wise aligned compacta and sparse edge sets.* We now come to the second of our extensions where we extend the class of space we consider from Peano continua to so-called component-wise aligned compacta. Observe that the ground space $\mathfrak{G}(X) := X - E(X)$ defined as the complement of all free arcs is well-defined for an arbitrary (metrisable) compactum X .

DEFINITION 14.5.10. A compact space X is said to be *component-wise aligned* if the components of X form a null-family of Peano continua, and $V_Y := \mathfrak{G}(X) \cap Y$ is an admissible vertex set for every component Y of X .

For a component-wise aligned compactum X , note that by definition, we have $E(X) = \bigsqcup \{E(Y, V_Y) : Y \text{ a component of } X\}$. In particular, we have $\mathfrak{G}(X) = \bigcup V_Y$. Next, the definition of an admissible vertex set generalises naturally to component-wise aligned compacta X : $V \subseteq X$ is admissible if $\mathfrak{G}(X) \subseteq V$ and $V \setminus \mathfrak{G}(X)$ is zero-dimensional. As before, this allows us to define edge-cuts for (X, V) in terms of edges crossing a clopen partition of V for all component-wise aligned compacta X and admissible vertex sets V of X . It is straightforward to check that all results about edges and edge-cuts from Section 14.1.3.1 still apply in this slightly generalised setting. In particular, it follows from the fact that each $E(Y, V_Y)$ is a zero-sequence and the fact that the components Y of X form a null-family, that $E(X)$ is a zero-sequence, and so all edge-cuts in a component-wise aligned compactum are finite.

LEMMA 14.5.11. *A component-wise aligned compactum has the even cut property if and only if every component of it has the even cut property.*

PROOF. The forward implication follows as in Lemma 14.5.8.

Conversely, suppose that X is a component-wise aligned compactum which has the even-cut property and let Y be a component of X . So let (A, B) be a closed partition of $\mathfrak{G}(Y)$ and consider the corresponding finite edge cut $D = E_Y(A, B)$. Then $X[A] = Y[A]$ and $X[B] = Y[B]$ are disjoint compact subsets of $X - D$, and each a union of components

of $X - D$. By the Šura-Bura Lemma, there is a clopen partition $U \oplus W$ of $X - D$ such that $X[A] \subseteq U$ and $X[B] \subseteq W$. But this means that $D = E_X[U \cap \mathfrak{G}(X), W \cap \mathfrak{G}(X)]$, and so D is even by assumption on X . \square

Finally, let us see three natural examples of component-wise aligned compacta X .

LEMMA 14.5.12. *Every graph-like compactum is component-wise aligned.*

PROOF. The fact that the components of a graph-like compactum form a null-sequence is tantamount to saying that graph-like continua are *finitely Souslian*, which is well-known, cf. [70, §2.2]. Moreover, since the ground-space of a compact graph-like continuum is zero-dimensional, each V_Y is zero-dimensional, and it follows readily that (Y, V_Y) is a graph-like continuum with vertex set V_Y . \square

LEMMA 14.5.13. *Every locally connected compactum is component-wise aligned.* \square

Recall that an edge set is sparse if it induces a graph-like subspace.

LEMMA 14.5.14. *Let X be a Peano continuum with admissible vertex set V , and $F \subseteq E(X, V)$ be a sparse edge set. Then $Y = X - F$ is a component-wise aligned compactum. More precisely:*

- (1) V is an admissible vertex set for Y , and (Y, V) is faithfully aligned in (X, V) , and
- (2) for every component Z of Y , we have that (Z, V_Z) for $V_Z := V \cap Z$ is faithfully aligned in (Y, V) , and hence in (X, V) .

PROOF. (1) Clearly, we have $\mathfrak{G}(Y) \subseteq \mathfrak{G}(X) \subseteq V$. Hence, it remains to show that $V \cap \bar{e}$ is compact zero-dimensional for every $e \in E(Y)$. Suppose not. Then there is a free arc $e \in E(Y)$ such that $\bar{e} \cap V$ is not zero-dimensional, so there exists a non-trivial subarc $\alpha \subseteq e \cap V$. Since F is sparse, $\bar{F} \cap V$ is zero-dimensional, $\alpha \setminus \bar{F}$ is an open subset of X consisting of intervals. But then any such interval is open in X but completely contained in V , a contradiction that V was admissible for X .

In particular, $E(Y, V) = E(X, V) \setminus F$, and hence (Y, V) is faithfully aligned in (X, V) .

(2) Let Z be a component of Y . The argument that $V_Z = V \cap Z$ is an admissible vertex set for Z is analogous to the previous case. To see that each (Z, V_Z) is faithfully aligned in (Y, V) , consider an edge $e \in E(Z, V_Z)$. We need to show that e is open in Y . Otherwise, there is a sequence of points $z_n \in Y \setminus Z$ such that $z_n \rightarrow z \in e$. Without loss of generality, we may assume that $z_n \in Z_n$ is contained in components Z_n of Y which are pairwise distinct. Let $x_n \in V \cap Z_n$ arbitrary. Since by Lemma 14.4.8(i) the non-trivial components of Y form a zero-sequence, it follows that $x_n \rightarrow z$ as well. However, since $z \notin V$, this contradicts the fact that V is closed. \square

14.5.2.5. *Circle decompositions.* Recall that the edge set of a Peano continuum X can be decomposed into edge-disjoint circles if there is a collection of edge-disjoint copies of

S^1 contained in X such that each edge of X is contained in precisely one such circle. We stress that this collection of copies of S^1 is not required to cover all of X , as this may be impossible even for graph-like continua, see [70, Example 4]. This example also shows that any two circles in such a circle decomposition may be disjoint in X .

Applying the results previously obtained in this section, we are now ready to prove the following result announced in Section 14.1.2.2 of the introduction:

THEOREM 14.5.15. *A Peano continuum has the even-cut property if and only if its edge set can be decomposed into edge-disjoint circles.*

PROOF. Our proof generalises the corresponding proof for countable graphs due to Nash-Williams [123]. For the reverse implication, let $\{S_n : n \in \mathbb{N}\}$ be a collection of edge-disjoint simple closed curves in X together covering all edges of X , each of which we may assume to be combinatorially aligned in X by Corollary 14.5.7. Then each S_n satisfies the even-cut condition, and the assertion now follows as in the proof of Lemma 14.5.8.

For the forward implication, fix an enumeration of the edge set of X which is possible by Lemma 14.1.11(c). We will find the circle decomposition recursively in countably many steps. Suppose inductively that we have already selected edge-disjoint, combinatorially aligned simple closed curves S_1, \dots, S_n in X so that the first n edges in our enumeration of $E(X)$ are covered. Since $F_n = \bigcup_{i \in [n]} E^{\text{real}}(S_i)$ is sparse, the space $X - F_n$ is a component-wise aligned compactum by Lemma 14.5.14. Now consider the first edge e in our enumeration of $E(X)$ not already covered by the previously selected simple closed curves (if there is no such edge, we are done). Otherwise, e is an edge of some faithfully aligned component Z of $X - F_n$. Since each S_i for $i \in [n]$ meets each edge cut of X in an even number of edges, it follows that $X - F_n$ has the even-cut property, and hence so does the Peano continuum Z by Lemma 14.5.11. Therefore, removing e does not disconnect Z , and we may select an $e(0) - e(1)$ -arc α_e in $Z - e$. Then $S_{n+1} = \alpha_e \cup e$ is a simple closed curve covering e , which we may assume to be combinatorially aligned in X by Corollary 14.5.7. Moreover, S_{n+1} is edge-disjoint to all previously selected simple closed curves, completing the induction step. After countably many steps no uncovered edges of X remain, and we have found a circle decomposition of X . \square

14.5.3. Ensuring the Even-Cut Condition.

14.5.3.1. *Inverse limit representations of graph-like compacta.* In this section, we briefly recall inverse limit techniques to deal with graph-like compacta and the even-cut condition from [70]. For an extensive discussion of inverse limits of finite multi-graphs, the reader may consult [54, §8.8] and [70].

For general background on inverse limits of compact Hausdorff spaces over directed sets, see [63, §2.5 and 3.2.13ff]. For an introduction to inverse limit sequences, that is to say, inverse limits where the underlying directed set is $(\mathbb{N}, <)$, see [121, Chapter II].

Let X be a component-wise aligned compactum with admissible vertex set V . By subdividing edges once, if necessary, we may assume that every edge of X has two distinct endpoints in V , so that X is *simple*. A *clopen partition* of V is a partition $\mathcal{U} = \{U_1, U_2, \dots, U_n\}$ of V into pairwise disjoint clopen sets. Write

$$E(\mathcal{U}) = \bigcup_{i \in [n]} E(U_i, V \setminus U_i)$$

for the (finite) set of all cross edges of the finite partition \mathcal{U} . Recall that $X[U_i]$ denotes the space U_i together with all edges from X that have both their end points in U_i .

Next let $\Pi = \Pi(V)$ be the set of all clopen partitions of V . The refinement relation naturally turns (Π, \preceq) into a directed set. Now given (X, V) and $\mathcal{U} \in \Pi(V)$, the *multigraph associated with \mathcal{U}* for some $\mathcal{U} \in \Pi$ is the finite graph $X_{\mathcal{U}}$ with vertex set \mathcal{U} and edge set $E(\mathcal{U})$ of all cross edges of the finite partition with the natural edge-vertex incidence. Formally, we set $X_{\mathcal{U}} = X / \{X[U] : U \in \mathcal{U}\}$. If $\pi_{\mathcal{U}} : X \rightarrow X_{\mathcal{U}}$ denotes the quotient mapping from X to the multigraph associated with \mathcal{U} , then $\pi_{\mathcal{U}}$ is a contraction map (however, if some $X[U_i]$ is not connected, then $\pi_{\mathcal{U}}$ is not an edge-contraction map).

Whenever $p \geq q \in \Pi(V)$, there are natural bonding maps $f_{pq} = \pi_q \circ \pi_p^{-1} : X_p \rightarrow X_q$. These maps send vertices of X_p to the vertices of X_q that contain them as subsets; they are the identity on the edges of X_p that are also edges of X_q ; and they send any other edge of X_p to that dummy vertex in X_q containing both its endpoints. In other words, each bonding map is a contraction map. Also, these maps are compatible in the inverse limit sense (whenever $p \geq q \geq r$ then $f_{pr} = f_{pq} \circ f_{qr}$), and hence $(X_p : p \in \Pi)$ forms an inverse system.

We now have the following facts (compare to [70, Theorem 13].)

- For any component-wise aligned compactum X , we have $X_{\sim} \cong \varprojlim (X_p : p \in \Pi)$.
- X (or equivalently X_{\sim}) satisfies the even-cut condition if and only if every X_p satisfies the even-cut condition if and only if every X_p is an even graph.

Indeed, to see this, note that for any admissible vertex set V of X there is a natural surjection $f : X \rightarrow Y := \varprojlim (X_p : p \in \Pi(V))$ defined by $f(x) := (\pi_p(x) : p \in \Pi(V))$. By [63, 3.2.11], it follows that Y is homeomorphic to the quotient $X / \{f^{-1}(y) : y \in Y\}$. But the non-trivial fibres of f correspond precisely to the non-trivial components of $\mathfrak{G}(X)$, and hence $X_{\sim} \cong \varprojlim (X_p : p \in \Pi)$ as desired.

We conclude this brief recap with an alternative description for component-wise aligned compacta X with only finitely many components (which is equivalent to saying they are locally connected). So let X be a locally connected compactum, and V an admissible vertex set for X . Let $\mathcal{E} = ([E(X, V)]^{<\infty}, \subseteq)$ denote the collection of finite edge sets of (X, V) , directed by inclusion. For $F \in \mathcal{E}$, the space $X - F$ has finitely many components, listed as say $V_F = \{C_1, \dots, C_k\}$ by Lemma 14.1.11. The *contraction of X onto F* , denoted by $X.F$, is the finite multi-graph with vertex set V_F and edge set F , where an edge in

F goes between those components in V_F that contain its end points in X . Formally, $X.F = X/V_F$ is defined as the topological quotient of X into the finitely many closed sets of V_F and points of $\bigcup F$. Note that if $\pi_F: X \rightarrow X.F$ denotes the quotient mapping from X to the multigraph $X.F$, then π_F is an edge-contraction map. The notation $X.F$ is taken from the same concept in matroid theory, see for example [133, Chapter 3]. Contrary to the graphs $X_{\mathcal{U}}$ from above, the graphs $X.F$ may also contain loops.

- For any locally connected compactum X , we have $X_{\sim} = \varprojlim (X.F: F \in \mathcal{E})$.
- X (or equivalently X_{\sim}) satisfies the even-cut condition if and only if every $X.F$ satisfies the even-cut condition if and only if every $X.F$ is an even graph.

The proof of the first fact can be derived from the previous inverse limit description as follows: if X is locally connected, and V an admissible vertex set for U , then pick a cofinal, refining sequence $(\mathcal{U}_n: n \in \mathbb{N}) \subseteq \Pi(V)$ such that $X[U]$ is connected for all $U \in \mathcal{U}_n$ and $n \in \mathbb{N}$. Then $(E(\mathcal{U}_n): n \in \mathbb{N})$ is cofinal in \mathcal{E} , and furthermore, it is clear from the definitions that $X_{\mathcal{U}_n} = X.E(\mathcal{U}_n)$ and that the bonding maps agree. Thus, using the fact that inverse limits of cofinal subsystems agree, it follows that for locally connected compacta X , we have

$$X = \varprojlim (X_p: p \in \Pi) = \varprojlim (X_{\mathcal{U}_n}: n \in \mathbb{N}) = \varprojlim (X.E(\mathcal{U}_n): n \in \mathbb{N}) = \varprojlim (X.F: F \in \mathcal{E}).$$

When X is a locally connected compactum, and $E(X) = \{e_1, e_2, \dots\}$ is any enumeration of its edges, then for $E_n = \{e_i: i \in [n]\}$, we obviously have that $(E_n: n \in \mathbb{N})$ is cofinal in \mathcal{E} . Hence, also $\varprojlim (X.E_n: n \in \mathbb{N})$ is a compact graph-like space homeomorphic to X_{\sim} .

14.5.3.2. *Inverse limits and sparse edge sets.* It will be important to understand how the even-cut condition changes when deleting or adding certain edge sets. For this, we shall need the following lemmas, which say that the inverse limit operation commutes with deletion of edges.

LEMMA 14.5.16. *Let X be a Peano continuum with admissible vertex set V , and $E(X, V) = \{e_1, e_2, \dots\}$ be any enumeration of its edges. For sparse $F \subseteq E(X, V)$ write $F_n := F \cap E_n$. Then $(X - F)_{\sim} = \varprojlim ((X.E_n) - F_n)$. In particular, if F is such that each $(X.E_n) - F_n$ is an even graph, then $X - F$ is a component-wise aligned compactum that has the even-cut property.*

PROOF. Consider a sparse edge set $F \subseteq E(X, V)$. By Lemma 14.5.14 we know that $Y = X - F$ is a component-wise aligned compactum with admissible vertex set V . Now for any $D \in \mathcal{E}$, let us write $F_D := F \cap D$ (so $F_n = F_{E_n}$) and consider the inverse limit $\mathcal{Y} = \varprojlim (X.D - F_D: D \in \mathcal{E})$. Now clearly, $(E_n: n \in \mathbb{N})$ is cofinal in \mathcal{E} , and we have $\mathcal{Y} = \varprojlim ((X.E_n) - F_n)$.

At the same time, for any cofinal sequence $(\mathcal{U}_n: n \in \mathbb{N})$ for $\Pi(V)$ we have $\mathcal{Y} = \varprojlim (X_{\mathcal{U}_n} - F_{E(\mathcal{U}_n)}: n \in \mathbb{N})$. However, given any clopen partition $\mathcal{U} \in \Pi(V)$, we have

$Y_{\mathcal{U}} = X_{\mathcal{U}} - F_{E(\mathcal{U})}$. Therefore, we have

$$Y_{\sim} = \varprojlim (Y_{\mathcal{U}_n} : n \in \mathbb{N}) = \varprojlim (X_{\mathcal{U}_n} - F_{E(\mathcal{U}_n)} : n \in \mathbb{N}) = \mathcal{Y} = \varprojlim ((X.E_n) - F_n),$$

and the first assertion of the lemma is proven.

The second part now follows now from the previous discussion about inverse limits and the even-cut property: if $(X.E_n) - F_n$ is even for each $n \in \mathbb{N}$, then \mathcal{Y} , and hence Y_{\sim} , have the even-cut property, too. \square

14.5.3.3. *Bipartite Peano partitions.* Recall Definition 14.3.7 for the definition of an intersection graph.

DEFINITION 14.5.17 (Bipartite Peano cover, zero-dimensional overlap). A Peano cover / partition \mathcal{U} is called *bipartite*, if its intersection graph $G_{\mathcal{U}}$ is bipartite.

For a bipartite Peano cover \mathcal{U} we also write $\mathcal{U} = \{K_1, K_2, \dots, K_{\ell}, U_1, U_2, \dots, U_k\}$ and mean that the K 's form one partition class, and the U 's form the other partition class of the bipartite graph $G_{\mathcal{U}}$. Even briefer, we say that (K, U) forms a bipartite Peano cover of some Peano continuum X if $X = K \cup U$ and both K and U are locally connected compacta (note that this is indeed a bipartite cover).

Finally, a bipartite Peano cover (K, U) is said to have *zero-dimensional overlap* if $K \cap U$ is zero-dimensional.

LEMMA 14.5.18. *Let X be a Peano continuum with admissible vertex set V . Then for every $\varepsilon > 0$ there is finite edge set $F \subseteq E(X, V)$ such that for each component D of $X - F$ there is a component C of V with $D \subseteq B_{\varepsilon}(C)$.*

PROOF. Suppose for a contradiction the assertion is false for some $\varepsilon > 0$. Enumerate $E(X, V) = \{e_1, e_2, e_3, \dots\}$ and let $F_n = \{e_1, \dots, e_n\}$. Then for each $n \in \mathbb{N}$, there is at least one *bad* component D of $X - F_n$ for which there is no component C of V with $D \subseteq B_{\varepsilon}(C)$. Further, every bad component of $X - F_{n+1}$ is contained in a bad component of $X - F_n$. Since $X - F_n$ has only finitely many components, Lemma 14.1.11, it follows from Königs Infinity Lemma [54, Lemma 8.1.2] that there is a decreasing sequence $(D_n : n \in \mathbb{N})$ of bad components D_n of $X - F_n$.

Since $\bigcup_n F_n = E(X, V)$, it follows that $C := \bigcap_n D_n$ is a component of V . However, since all D_n are closed in X and $\bigcap_n D_n \subseteq B_{\varepsilon}(C)$, it follows from topological compactness that there is $N \in \mathbb{N}$ with $D_N \subseteq B_{\varepsilon}(C)$, contradicting that D_N was bad. \square

THEOREM 14.5.19. *Let X be a Peano continuum, and suppose that $X = K \cup U$ such that $K = K_1 \oplus K_2 \oplus \dots \oplus K_{\ell}$ consists of finitely many Peano components and the non-trivial components of U form a zero-sequence of Peano continua U_1, U_2, \dots . Suppose further that every edge of K intersects at most one U_i . Let V be an admissible vertex set of K .*

Then for every $\varepsilon > 0$ there is $N \in \mathbb{N}$ such that $K' = K \cup \bigcup_{n>N} U_n$ admits a finite edge set $F_K \subseteq E(K, V)$ so that for each component D' of $K' - F_K$ there is a component C of V with $D \subseteq B_\varepsilon(C)$.

PROOF. Apply Lemma 14.5.18 to find $F_K \subseteq E(K, V)$ finite such that components of $K - F_K$ are $\varepsilon/2$ -close to V . The components of $K - F_K$ are finitely many disjoint closed subsets of X , so some pair has minimal distance from each other. Denote that minimal distance by $\delta > 0$. Let $\eta := \min \{\varepsilon/2, \delta/3\}$.

Now choose $N \in \mathbb{N}$ large enough such that $\text{diam}(U_n) < \eta$ and $U_n \cap (\bigcup F_K) = \emptyset$ for all $n \geq N$. We claim that N is as desired. First, note that since X is connected, every U_n has non-empty intersection with K . Therefore, it follows that $K' = K \cup \bigcup_{n>N} U_n$ still has at most ℓ components, which are all Peano by Lemma 14.1.13.

Moreover, any two components of $K' - F_K$ have, by choice of η and N , distance at least $\delta - 2\eta > 0$. In particular, no two components of $K - F_K$ fuse together by adding $\bigcup_{n>N} U_n$. Hence, for any component D' of $K' - F_K$ there is a component D of $K - F_K$ such that $D' \subseteq B_\eta(D)$. And by choice of F_K , there is a component C of V such that $D \subseteq B_{\varepsilon/2}(C)$. Thus, $D' \subseteq B_{\eta+\varepsilon/2}(C) \subseteq B_\varepsilon(C)$, which completes the proof. \square

14.5.3.4. *Modifying Peano partitions with zero-dimensional boundaries.* Consider a Peano graph X for which we have a bipartite Peano partition (K, U) with zero-dimensional overlap. In this subsection, we demonstrate how to modify the elements of K and U to obtain a new bipartite partition K', U' as to guarantee that the resulting K', U' satisfy the even-cut condition. Moreover, we will do these changes so that K' and U' are arbitrarily close to the original K and U .

THEOREM 14.5.20. *Let X be a Peano continuum satisfying the even-cut condition that has a bipartite Peano partition $\mathcal{U} = (K, U)$ with zero-dimensional overlap. Then for every $\varepsilon > 0$ there is a bipartite Peano cover $\mathcal{U}' = (K', U')$ such that*

- (A1) $K \subseteq K'$ and $U' \subseteq U$,
- (A2) *there is a finite edge set $F_K \subseteq E(K')$, so that each component of $K' - F_K$ either has diameter $< \varepsilon$ or is ε -close to a component of $\mathfrak{G}(K)$, and*
- (A3) *all elements of \mathcal{U}' satisfy the even-cut condition.*

PROOF. Since $K \cap U$ is compact zero-dimensional, the set $V := \mathfrak{G}(X) \cup (K \cap U)$ is an admissible vertex set for X . Then every element of \mathcal{U} with the naturally induced admissible vertex set is faithfully aligned with (X, V) . Write $K = K_1 \oplus K_2 \oplus \dots \oplus K_\ell$ and $U = U_1 \oplus U_2 \oplus \dots \oplus U_k$ for the Peano components of the two sides (K, U) . Since $U_i \cap K \subseteq U_i$ is zero-dimensional and contained in the vertex set of U_i for each $i \in [k]$, by Lemma 14.5.9 there are combinatorially aligned graph-like trees $T_i \subseteq U_i$ with $U_i \cap K \subseteq V(T_i)$. Define $T = \bigcup_{i \in [k]} T_i$, a graph-like forest with k components. Note that T is combinatorially

aligned with (X, V) but may contain fake edges (edges contained in the ground space of X). However, as $T \cap K = U \cap K \subseteq V(T)$, no edge of T intersects K .

In order to arrange for (A3), our aim is to find a subset $F \subseteq E(T)$ such that by adding F to K , denoted by $K + F := K \cup T[F]$, and removing $F^{\text{real}} = F \cap E_T^{\text{real}}$ from U , denoted by $U - F^{\text{real}}$, we obtain an edge-disjoint cover $\{K + F, U - F^{\text{real}}\}$ of X such that both sides satisfy the even-cut condition. In order to find this set F , we use logical compactness as follows. First, let $E(X) \cup E(T) = \{e_1, e_2, e_3, \dots\}$ be an enumeration of the countably many edges of (X, V) together with the fake edges of T . Put $E_n := \{e_1, \dots, e_n\}$. Define $K^* = K \cup T$, which is a Peano continuum. Now define (using the notation $Y.F := Y.(E(Y) \cap F)$, called *contracting onto F* , as introduced in Section 14.5.3.1 above)

$$X_n := X.E_n, \quad K_n^* := K^*.E_n, \quad K_n := K.E_n, \quad U_n := U.E_n, \quad \text{and} \quad S_n := T.E_n.$$

We reiterate that not all edges of E_n are edges of X . So $X_n - E_n$ stands for $X - (E_n \cap E(X))$, $X_n = X.E_n$ stands for $X.(E_n \cap E(X))$, and so $E(X_n) = E_n \cap E(X)$, and similarly in the other cases. By the results from Section 14.5.3.1, we have $X_\sim = \varprojlim X_n$, and similarly in the other cases. Note also that since X is connected and satisfies the even-cut condition, every finite graph X_n is Eulerian.

DEFINITION 14.5.21. Let $\kappa: K \rightarrow K^*$, $\sigma^*: T \rightarrow K^*$ and $\sigma: T \rightarrow U$ be the (injective) inclusion maps. For every $n \in \mathbb{N}$, let π_n be the (surjective) projection maps corresponding to the operation of contracting onto the edge set E_n , and define

- $\kappa_n := \pi_n \circ \kappa \circ \pi_n^{-1}: K_n \rightarrow K_n^*$,
- $\sigma_n^* := \pi_n \circ \sigma^* \circ \pi_n^{-1}: S_n \rightarrow K_n^*$, and
- $\sigma_n := \pi_n \circ \sigma \circ \pi_n^{-1}: S_n \rightarrow U_n$.

We may visualise these maps in a commuting diagram as follows:

$$\begin{array}{ccccccc}
 K & \xrightarrow{\kappa} & K^* & \xleftarrow{\sigma^*} & T & \xrightarrow{\sigma} & U \\
 \downarrow \pi_n & & \downarrow \pi_n & & \downarrow \pi_n & & \downarrow \pi_n \\
 K_n & \xrightarrow{\kappa_n} & K_n^* & \xleftarrow{\sigma_n^*} & S_n & \xrightarrow{\sigma_n} & U_n
 \end{array}$$

LEMMA 14.5.22. *The following facts about the above diagram are true:*

- (1) *The maps κ_n , σ_n^* and σ_n are well-defined (i.e. single valued) contraction maps, and the diagram commutes.*
- (2) *$\kappa_n \upharpoonright E(K_n)$, $\sigma_n^* \upharpoonright E(S_n)$ and $\sigma_n \upharpoonright E^{\text{real}}(S_n)$ act as identity, whereas $\sigma_n(E^{\text{fake}}(S_n)) \subseteq V(U_n)$,*
- (3) *$\kappa(K)$ and $\sigma^*(T)$ form a decomposition of K^* into connected subgraphs, and hence $\kappa_n(K_n)$ and $\sigma_n^*(S_n)$ form a decomposition of K_n^* into connected subgraphs,*
- (4) *If $P \subseteq T$ is a standard arc with end-vertices a and b , then*
 - *$Q = \pi_n(P)$ forms a path in S_n with edge set $F := E(P) \cap E(S_n)$,*

- $\sigma_n^*(Q)$ forms a trail²⁴ in K_n^* with edge set F from $\pi_n(\sigma^*(a))$ to $\pi_n(\sigma^*(b))$,
- $\sigma_n(Q)$ forms a trail in U_n with edge set F_n^{real} from $\pi_n(\sigma(a))$ to $\pi_n(\sigma(b))$.

PROOF. (1) and (2). To see that κ_n is a well-defined contraction map and acts as identity on $E(K)$, note that since $E_n \cap E(K) \subseteq E_n \cap E(K^*)$, it follows that every edge $e \in E(K)$, we have $\pi_n^{-1}(e) = e$, and hence $\kappa_n(e) = \pi_n \circ \kappa \circ \pi_n^{-1}(e) = e$. For a vertex $v \in V(K_n)$, note that by definition $\pi_n^{-1}(v)$ is a connected component of $K - E_n$. Hence, $\kappa(\pi_n^{-1}(v))$ is a connected subspace of $K^* - E_n$, and hence belongs to a connected component of $K^* - E_n$. Thus, $\pi_n(\kappa(\pi_n^{-1}(v))) = \kappa_n(v)$ is a vertex of K_n^* .²⁵ The proof for σ_n^* is the same. The third case of σ_n is almost the same, with the difference that while σ_n is the identity on real edges of S_n , for every fake edge e of S_n , we have $\sigma(\pi_n^{-1}(e)) \subseteq \mathfrak{G}(U)$, and hence belongs to a connected component of $U - E_n$, so $\sigma_n(e) = \pi_n(\sigma(\pi_n^{-1}(e))) \in V(U_n)$.

Next, assertion (3) is clear by construction and the fact that $\kappa_n \upharpoonright E(K_n)$, $\sigma_n^* \upharpoonright E(S_n)$ act as identity. Finally, (4) follows from the fact that since all maps are contraction maps, trails get mapped to trails. \square

Let us call a subset $F_n \subseteq E(S_n)$ *semi-good* if $U_n - \sigma_n(F_n) = U_n - F_n^{\text{real}}$ is an even subgraph of U_n . A semi-good set is called *good*, if also $\kappa(K_n) + \sigma_n^*(F_n) = K_n^*[E(K_n) \cup F_n]$ is an even subgraph of K_n^* .

Main claim: *For each $n \in \mathbb{N}$ there exists at least one good subset of $E(S_n)$.*

We will prove our main claim in two steps, first constructing a semi-good set, which we modify in a second step to a good set.

Step 1: There exists a semi-good subset $F'_n \subseteq E(S_n)$. To see this, note that each graph U_n has precisely k connected components, and by the handshaking lemma, the number of odd-degree vertices of U_n inside each component is even, so come in pairs. Let \approx denote the corresponding equivalence relation, where each equivalence class consists of one such pair. Now for each vertex $u \in V(U_n)$, the preimage $\pi_n^{-1}(u)$ induces a clopen subset of the vertex set $V \cap U$ of U . If u has odd degree, then necessarily $\pi_n^{-1}(u) \cap K \neq \emptyset$, as otherwise the edge-cut of $\pi_n^{-1}(u)$ induced in U equals the edge-cut of $\pi_n^{-1}(u)$ induced in X , contradicting the even-cut property of X . By construction of T , there is a point $v_u \in \pi_n^{-1}(u) \cap K \cap V(T)$, and this point must satisfy $u = \pi_n(\sigma(v_u))$. Next, for each pair $u \approx u'$ of odd-degree vertices of U_n , v_u and $v_{u'}$ lie in the same connected component of T , so there exists a unique path $P_{v_u, v_{u'}}$ in T from v_u to $v_{u'}$. By Lemma 14.5.22(4), if we let $Q_{u, u'} = \pi_n(P_{v_u, v_{u'}})$ be the corresponding path in S_n , then $\sigma_n(Q_{u, u'})$ is a trail in U_n from $\sigma_n(\pi_n(v_u)) = \pi_n(\sigma(v_u)) = u$ to $\sigma_n(\pi_n(v_{u'})) = \pi_n(\sigma(v_{u'})) = u'$, where the respective first equalities hold since the above diagram commutes, and the respective second equalities

²⁴A trail is a walk without repeated edges

²⁵Note, however, that distinct vertices $v \neq v' \in V(K_n)$ may be mapped onto the same vertex in $V(K_n^*)$, as $\pi_n^{-1}(v)$ and $\pi_n^{-1}(v')$ are distinct components of $K - E_n$, but as subspaces might belong to the same component of $K^* - E_n$.

hold by choice of v_u and $v_{u'}$. In particular, all vertices, apart from the end-vertices have even degree in that trail. Define $F'_n := \sum_{u \approx u'} E(Q_{u,u'})$. Then $\sigma_n(F'_n) = \sum_{u \approx u'} \sigma_n(Q_{u,u'})$ is the mod-2 sum over these trails, and so it is precisely the odd degree vertices of U_n that have odd parity in $U_n[\sigma_n(F'_n)]$. Thus, $U_n - \sigma_n(F'_n)$ is an even graph, and so F'_n is semi-good.

Step 2: There exists a good subset $F_n \subseteq E(S_n)$. First, fix a semi-good subset $F'_n \subseteq E(S_n)$, let $F_n^{\complement} = E(S_n) \setminus F'_n$ and define $K'_n = K_n^* - \sigma_n^*(F_n^{\complement})$ and $U'_n = U_n - \sigma_n(F'_n)$. As before, for each vertex $k \in V(K'_n) = V(K_n^*)$, the set $\pi_n^{-1}(k)$ is a connected component of $K_n^* - E_n$, and hence a subcontinuum of $X - E_n$. Similarly, for each vertex $u \in V(U'_n) = V(U_n)$, the set $\pi_n^{-1}(u)$ is a connected component of $U - E_n$, and hence also a subcontinuum of $X - E_n$. Hence, for $\mathcal{U} = \{\pi_n^{-1}(v) : v \in V(K_n^*) \sqcup V(U_n)\}$ we may consider the intersection graph $G = G_{\mathcal{U}}$ of \mathcal{U} in $X - E_n$. For ease of notation, relabel

$$V(G) = V(K_n^*) \sqcup V(U_n) \quad \text{and} \quad E(G) = \{vw : \pi_n^{-1}(v) \cap \pi_n^{-1}(w) \neq \emptyset\}.$$

Observe that G is a bipartite graph with vertex bipartition $V(G) = V(K_n^*) \sqcup V(U_n)$, as whenever $k \neq k'$ are distinct vertices in K_n^* , then $\pi_n^{-1}(k)$ and $\pi_n^{-1}(k')$ are distinct components of $K_n^* - E_n$, and hence do not intersect, and similarly for $u \neq u' \in V(U_n)$.

SUBCLAIM 1. *Whenever $ku \in E(G)$, then $\pi_n^{-1}(k) \cap \pi_n^{-1}(u) \cap V(T) \neq \emptyset$.*

PROOF OF SUBCLAIM 1. Since $K_n^* \cap U = (K \cap U) \cup T$, the fact that $ku \in E(G)$ implies $\pi_n^{-1}(k) \cap \pi_n^{-1}(u) \subseteq (K \cap U) \cup T$. Since $K \cap U \subseteq V(T)$, we only have to consider the case where $\pi_n^{-1}(k) \cap \pi_n^{-1}(u)$ intersect in an edge e of $E(T)$, in which case $e(0), e(1) \in \pi_n^{-1}(k) \cap \pi_n^{-1}(u) \cap V(T)$, as $\pi_n^{-1}(k)$ and $\pi_n^{-1}(u)$ are standard subcontinua, and if e is a fake edge, then $\bar{e} \subseteq \mathfrak{G}(U)$, so contained in a single component of $U - E_n$. ◇

Next, for every connected component C of the graph G , the set $\bigcup \pi_n^{-1}(C)$ is a subspace of $X - E_n$. Write $\mathcal{C}(G) := \{\bigcup \pi_n^{-1}(C) : C \text{ a connected component of } G\}$.

SUBCLAIM 2. *We have $\{\pi_n^{-1}(x) : x \in V(X_n)\} = \mathcal{C}(G)$.*

PROOF OF SUBCLAIM 2. This will follow once we show that $\mathcal{C}(G)$ forms a partition of $X - E_n$ into subcontinua. First, each $\pi_n^{-1}(C)$ is a subcontinuum of $X - E_n$. This follows easily by induction on $|C|$, since for every edge $ku \in E(G)$, the two subcontinua $\pi_n^{-1}(k)$ and $\pi_n^{-1}(u)$ intersect by definition, so $\pi_n^{-1}(k) \cup \pi_n^{-1}(u)$ is again a subcontinuum. Next, for components $C \neq C'$ of A , if $\bigcup \pi_n^{-1}(C) \cap \bigcup \pi_n^{-1}(C') \neq \emptyset$, there would be $v \in C$ and $w \in C'$ such that $\pi_n^{-1}(v) \cap \pi_n^{-1}(w) \neq \emptyset$, and so $vw \in E(G)$, contradicting that v and w belong to distinct components of G . Finally, $X - E_n \subseteq (K_n^* - E_n) \cup (U_n - E_n)$ yields that $\bigcup \pi_n^{-1}(V(G)) = X - E_n$. ◇

Now a component C of G can be viewed as a single vertex of X_n , and hence induces an edge cut in X_n . Similarly, by the nature of G , a component C also induces edge cuts in

K'_n and in U'_n : write $E_{K'_n}(C, C^{\mathbb{L}})$ as shorthand for the edge cut of K'_n with sides $V(K'_n) \cap C$ versus $V(K'_n) \setminus C$.

SUBCLAIM 3. *We have $E_{X_n}(C, C^{\mathbb{L}}) = E_{K'_n}(C, C^{\mathbb{L}}) \sqcup E_{U'_n}(C, C^{\mathbb{L}})$ for any component C of G , and hence $E_{K'_n}(C, C^{\mathbb{L}})$ is always even.*

PROOF OF SUBCLAIM 3. To see this claim, note that $E_{K'_n}(C, C^{\mathbb{L}})$ cannot contain fake edges of T , as any such edge lies in $\mathfrak{G}(U)$, contradicting that C is a component of A . Hence, all edge cuts are subsets of $E(X_n)$. The equality of sets now follows from that fact that K'_n and U'_n are $E(X_n)$ -edge-disjoint, and together cover all edges of X_n . Now since X_n and U'_n were even graphs by assumption, and so have the even-cut property, it follows that $E_{K'_n}(C, C^{\mathbb{L}})$ is even for every component C of A . \diamond

To complete the proof of the second step, and hence of our main claim, note that by Subclaim 3 and the handshaking lemma, for any connected component C of G , the number of vertices of K_n^* which have odd-degree in K'_n in C is always even. Hence, we can pair up odd degree vertices of K'_n such that for every pair $k \approx k'$ there is a path $A_{k,k'}$ in G say with vertices $k_0 u_1 k_1 u_1 \dots u_{j-1} k_j$ where $k = k_0$, $k' = k_j$, $k_i \in V(K_n^*)$, $u_i \in V(U_n)$ and edges $\{k_0 u_1, u_1 k_1, k_1 u_2, \dots, u_{j-1} k_j\} \subseteq E(G)$, using that G is bipartite. By Subclaim 1, for every $i \in [j]$ we may pick a point $a_i \in \pi_n^{-1}(k_{i-1}) \cap \pi_n^{-1}(u_i) \cap V(T)$ and a point $b_i \in \pi_n^{-1}(u_i) \cap \pi_n^{-1}(k_i) \cap V(T)$, and let P_i be the unique path from a_i to b_i in the forest T , which exists as $\pi_n^{-1}(u_i)$ is contained in a unique component of U .

Now arguing as in Step 1, if we let $Q_i = \pi_n(P_i)$ be the corresponding path in S_n , then $\sigma_n(Q_i)$ is a trail in U_n from $\pi_n(\sigma(a_i)) = u_i$ to $\pi_n(\sigma(b_i)) = u_i$, i.e. $\sigma_n(Q_i)$ is a closed trail, so all vertices of U_n in $\sigma_n(Q_i)$ have even degree. Hence, $\sum_{i \in [j]} \sigma_n(Q_i)$ is an even subgraph of U_n . At the same time, however, every $\sigma_n^*(Q_i)$ is a trail in K_n^* from $\pi_n(\sigma^*(a_i)) = k_{i-1}$ to $\pi_n(\sigma^*(b_i)) = k_i$, and so $\sum_{i \in [j]} \sigma_n^*(Q_i)$ induces a subgraph in K_n^* in which all vertices, apart from $k = k_0$ to $k' = k_n$ have even degree. Thus, if we let $F_{k,k'} = \sum_{i \in [j]} E(Q_i)$, then $\sigma_n(F_{k,k'})$ is an even subgraph of U_n , and in the subgraph induced by $\sigma_n^*(F_{k,k'})$ in K_n^* , all vertices have even parity apart from precisely k and k' . Hence, $F_n := F'_n + \sum_{k \approx k'} F_{k,k'}$ is a good subset $F_n \subseteq E(S_n)$. This completes the proof of Step 2.

Recall that we set out to show the existence of a set $F \subseteq E(T)$ such that by adding F to K and removing $F^{\text{real}} = F \cap E_T^{\text{real}}$ from U , we obtain an edge-disjoint cover $\{K + F, U - F^{\text{real}}\}$ of X such that both sides satisfy the even-cut condition. We will now obtain such a set F from the good edge sets of $E(S_n)$ as follows. Since $E(S_n)$ is finite, each $E(S_n)$ has only finitely many good subsets. Moreover, since $U_n = U_{n+1}/e_{n+1}$ and $K_n^* = K_{n+1}^*/e_{n+1}$ are obtained by edge-contraction, even subgraphs of H_{n+1} and K_{n+1}^* restrict to even subgraphs of U_n and K_n^* . Thus, every good choice $F_{n+1} \subseteq E(S_{n+1})$ at step $n + 1$ induces a good choice $F_n = F_{n+1} \cap E(S_n)$ at step n . So by Königs Infinity Lemma [54, Lemma 8.1.2], there is a sequence of good sets $(F_n : n \in \mathbb{N})$ with $F_n \subseteq E(S_n)$ such

that $F_{n+1} \cap E(S_n) = F_n$ for all $n \in \mathbb{N}$. Now given such a sequence $(F_n : n \in \mathbb{N})$, define $F = \bigcup_{n \in \mathbb{N}} F_n \subseteq E(T)$ and claim that F is as desired, i.e. that $K + T[F]$ and $U - F^{\text{real}}$ have the even-cut property. Indeed, since $F^{\text{real}} \cap E(U_n) = F_n^{\text{real}}$ it follows from Lemma 14.5.16 that $(U - F^{\text{real}})_{\sim} = \varprojlim (U_n - F_n^{\text{real}})$ has the even-cut property. Hence, $U - F^{\text{real}}$ has the even-cut property. Similarly, also $K \cup T[F]$ has the even cut property, as $K_{\sim}[E(K) \cup F] = \varprojlim (K_n^*[E(K_n) \cup F_n])$ is the inverse limit of even graphs.

Moreover, since $K'' = K \cup T[F]$ satisfies the even-cut condition, every leaf of $T[F]$ must intersect K (as otherwise, there would be a vertex in $(K \cup T[F])_{\sim}$ of degree 1, contradicting the even-cut property), and hence $K \cup T[F]$ continues to have at most ℓ connected components. Moreover, since the non-trivial components of $T[F]$ form a zero-sequence of graph-like continua, Lemma 14.5.12, each of the ℓ components of $K \cup T[F]$ remains a Peano continuum, Lemma 14.1.13. Since F is sparse, $U'' = U - F^{\text{real}}$ is a component-wise aligned compactum such that every component of U'' is faithfully aligned in (X, V) , Lemma 14.5.14. By Lemma 14.5.11, each component of U'' satisfies the even-cut condition. To complete the proof of the theorem, we would like U'' to have only finitely many components. We rectify this problem by reassigning all but finitely many of these components of U'' back to K'' , without violating property (A2). Indeed, we may construct K' and U' as desired by applying Theorem 14.5.19 with ε , providing a finite edge set F_K as to satisfy (A3). Moreover, that by Lemma 14.5.8, this reassignment preserves the even-cut condition of K'' , and so K' and U' satisfy (A2). That it satisfies (A1) is clear from construction, since we only ever added edge sets to K . \square

14.5.4. Eulerian Decompositions of One-Dimensional Peano Continua.

14.5.4.1. The decomposition theorem.

THEOREM 14.5.23 (2nd decomposition theorem). *Every one-dimensional Peano continuum $X \subseteq [0, 1]^3$ with admissible vertex set V satisfying the even-cut condition admits a Peano cover $\{X_1, \dots, X_s\}$ into edge-disjoint standard connected, combinatorially aligned Peano subgraphs with edge sets V_i each satisfying the even-cut condition, and for each $i \in [s]$ there is a edge vertex set $F_i \subseteq E(X_i, V_i)$ such that every component C of $X_i - F_i$ either satisfies $C \subseteq [0, \frac{2}{3}] \times [0, 1] \times [0, 1] \subseteq [0, 1]^3$ or $C \subseteq [\frac{1}{3}, 1] \times [0, 1] \times [0, 1] \subseteq [0, 1]^3$.*

Our proof relies crucially on the fact that one-dimensional Peano continua have exceptionally nice Peano partitions (Def. 14.2.6) that reflect properties of dimension, announced by Bing in [26, Theorem 11] and used crucially by Andersen as a step towards the topological characterisation of the Menger universal curve in [7, 8]. See also [117] for a detailed account, including a published proof in the one-dimensional case.

THEOREM 14.5.24 ([117, Theorem 2.9]). *A one-dimensional Peano continuum admits a decreasing sequence of $1/n$ Peano partitions $\{\mathcal{U}_n : n \in \mathbb{N}\}$ with zero-dimensional boundaries.*²⁶

PROOF OF THEOREM 14.5.23. For $i \in [3]$ let $\pi_i : [0, 1]^3 \rightarrow [0, 1]$ denote the projection map from the cube onto the i th coordinate. Let $\varepsilon = 1/6$. Pick an ε -brick-partition \mathcal{U} of X with zero-dimensional boundaries as in Theorem 14.5.24, and let $\mathcal{U}_u \subseteq \mathcal{U}$ be the subcollection $\mathcal{U}_u = \{U \in \mathcal{U} : U \cap \pi_1^{-1}[2/3, 1] \neq \emptyset\}$ and let $\mathcal{U}_\ell := \mathcal{U} \setminus \mathcal{U}_u$. Next, let $K = \bigcup \mathcal{U}_u$, and similarly let $U = \bigcup \mathcal{U}_\ell$, giving rise to a bipartite Peano partition $\mathcal{U} = (K, U)$ of X with zero-dimensional overlap by the sum theorem of dimension, [62, Thm. 1.5.2]. Apply Theorem 14.5.20 to \mathcal{U} with $\varepsilon = 1/3$ to obtain a bipartite Peano cover $\mathcal{U}' = (K', U')$ of X with properties (A1), (A2) and (A3) of Theorem 14.5.20. For later use, let F_K denote the finite edge set of K' witnessing (A2). We claim that \mathcal{U}' is as desired.

Clearly, by construction and property (A3), \mathcal{U}' is a finite decomposition of X into edge-disjoint standard Peano subgraphs each satisfying the even-cut condition. To see the first bullet point, note that by (A1), $U' \subseteq U$ and so every component of U' is contained in a component of U , which by construction was almost contained in $[0, \frac{2}{3}] \times [0, 1] \times [0, 1]$.

Lastly, we claim that F_K from (A2) is a witness for the second bullet point. Indeed, any component C of $K' - F_K$ either has diameter $\text{diam}(C) \leq \varepsilon < 1/3$, in which case we have trivially

$$C \subseteq [0, \frac{2}{3}] \times [0, 1] \times [0, 1] \subseteq [0, 1]^3 \quad \text{or} \quad C \subseteq [\frac{1}{3}, 1] \times [0, 1] \times [0, 1] \subseteq [0, 1]^3,$$

or C is contained in $B_\varepsilon(D)$ for some component D of $\mathfrak{G}(K)$. In this case, since by construction, we have $D \subseteq K \subseteq [\frac{2}{3} - \varepsilon, 1] \times [0, 1] \times [0, 1]$, the fact $C \subseteq B_\varepsilon(D)$ implies that

$$C \subseteq [\frac{2}{3} - 2\varepsilon, 1] \times [0, 1] \times [0, 1] = [\frac{1}{3}, 1] \times [0, 1] \times [0, 1],$$

completing the proof.

Note that by Corollary 14.5.7, given (X, V) we may pick admissible vertex sets for K and U such that they are combinatorially aligned with (X, V) . □

14.5.4.2. *Eulerian decompositions of one-dimensional Peano continua.* In this section we finally prove Theorem 14.5.23. Let us fix a one-dimensional Peano continuum X which satisfies the even-cut condition. By Nöbling’s embedding theorem [62, 1.11.4], every one-dimensional continuum embeds into the unit cube $[0, 1]^3$, and so for our purposes we may assume that X is given as a subspace $X \subseteq [0, 1]^3$. The goal is to show how the decomposition theorem may be used to construct an approximating sequence of Eulerian decompositions for X , thereby implying the Eulerianity conjecture for all one-dimensional Peano continua.

²⁶The Theorem proved in [117, Thm. 2.9] is stronger, but we shall not need these additional properties.

First, recall that by [62, Thm. 1.8.13], since X is one-dimensional, the complement of X in $[0, 1]^3$ is connected, and since it is open, it must then be path-connected. Therefore, given $X \subseteq [0, 1]^3$, we may add any finite set of edges between specified points of X in 3-space to obtain a Peano continuum X' such that $X \subseteq X' \subseteq [0, 1]^3$.

DEFINITION 14.5.25 (Truncation). Let $\mathcal{D} = (G, \eta)$ be a decomposition of a Peano continuum X , and let $v \in V(G)$. The *truncation of \mathcal{D} to v* , denoted by $\tau(v)$, is a Peano continuum with $\tau(v) \supseteq \eta(v)$ with additional edges $E(\tau(v)) \setminus E(\eta(v)) = \{e \in E(G) : e \sim v\}$ and ground set

$$\mathfrak{G}(\tau(v)) = \begin{cases} \mathfrak{G}(\eta(v)) & \text{if } E_G(v, G - v) = \emptyset, \\ \mathfrak{G}(\eta(v)) \oplus \{\star\} & \text{otherwise,} \end{cases}$$

where vertex-edge incidences for the new edges are given by

$$e_\tau(i) = \begin{cases} (\eta(e))(i) & \text{if } e(i) = v \\ \star & \text{otherwise.} \end{cases}$$

for $e \sim v$ in G and $i \in \{0, 1\}$.

Truncating means first contracting the subgraph $G[V(G - v)]$ to a single vertex \star , and then blowing up the ‘vertex’ v to its associated tile $\eta(v)$, connecting all edges previously incident with v in G to their correct endpoints in $\eta(v)$. The case distinction ensures that if \star was isolated, it is to be disregarded (there might still be loops attached to v in G).

From the above discussion we deduce the next lemma.

LEMMA 14.5.26. *Let $\mathcal{D} = (G, \eta)$ be a decomposition of a Peano continuum X . A truncation $\tau(v)$ is always a connected Peano graph, and if $\eta(v) \subseteq [0, 1]^3$, then we may always assume that $\eta(v) \subseteq \tau(v) \subseteq [0, 1]^3$ for all $v \in V(\Gamma)$.*

As announced, let us see how the Decomposition Theorem 14.5.23 can be used to construct an approximating sequence of Eulerian decompositions. For an example of an approximating sequence of Eulerian decompositions that satisfies property (E9) in the next proof, consider once more the hyperbolic 4-regular tree from Figure 14.6 in Chapter 14.3.

PROOF OF THEOREM 14.5.1. We construct a sequence $((G_n, \eta_n) : n \in \mathbb{N})$ of Eulerian decompositions for X with $(G_0, \eta_0) \preceq (G_1, \eta_1) \preceq \dots$ by recursion on n , such that each Eulerian decomposition (G_n, η_n) satisfies, besides its usual properties (E1)–(E7) from Definition 14.3.2, the following extra two requirements:

- (E8) each tile $\eta_n(v)$ is combinatorially aligned with X ,
- (E9) each truncation $\tau_n(v)$ satisfies the even-cut condition for all vertices v of (G_n, η_n) ,
- (E10) for every vertex v of (G_n, η_n) , the tile $\eta_n(v)$ is contained in a cube I_v with

$$\eta_n(v) \subseteq I_v = I_v^1 \times I_v^2 \times I_v^3 \subseteq [0, 1]^3$$

such that for $r = n \pmod{3}$ we have

$$\text{diam}(I_v^k) = \begin{cases} \left(\frac{2}{3}\right)^{\lfloor n/3 \rfloor + 1} & \text{if } k \leq r \\ \left(\frac{2}{3}\right)^{\lfloor n/3 \rfloor} & \text{otherwise.} \end{cases}$$

For the base case, we can choose the trivial decomposition. So suppose for some $n \in \mathbb{N}$ we have an Eulerian decomposition (G_n, η_n) with properties (E8), (E9) and (E10), and write $E(G_n) = F_n \sqcup D_n$ for the implicit partition into displayed and dummy edges. Our task is to construct an Eulerian decomposition (G_{n+1}, η_{n+1}) with properties (E8), (E9) and (E10), so that (G_{n+1}, η_{n+1}) extends (G_n, η_n) . In order to satisfy (E10) at step $n + 1$, it is clear that we have to cut our tiles apart along the unique coordinate $i \in \{1, 2, 3\}$ where $n + 1 = 3m + i$ for some $m \in \mathbb{N}$; without loss of generality, we may assume in the following that $i = 1$.

Consider $v \in V(G_n)$. For ease of notation, we rescale affinely in all coordinates so that $I_v = [0, 1]^3$. By Lemma 14.5.26, we may assume that $\eta(v) \subseteq \tau_n(v) \subseteq [0, 1]^3$. Then in combination with property (E8) and (E9), we are allowed to apply Theorem 14.5.23 to the truncation $\tau_n(v)$ and obtain a finite Peano cover

$$\mathcal{S}_v = \{X_1, X_2, \dots, X_{s(v)}\}$$

of $\tau_n(v)$ such that

- (i) the elements are pairwise edge-disjoint,
- (ii) each element satisfies the even-cut condition,
- (iii) each element is combinatorially aligned $\tau_n(v)$,
- (iv) for each $i \in [s(v)]$ there is a finite edge set $F_i \subseteq E(X_i)$ such that every component C of $X_i - F_i$ either satisfies $C \subseteq [0, \frac{2}{3}] \times [0, 1] \times [0, 1] \subseteq [0, 1]^3$ or $C \subseteq [\frac{1}{3}, 1] \times [0, 1] \times [0, 1] \subseteq [0, 1]^3$.

Write $E_v = E(\tau_n(v)) \setminus E(\eta_n(v))$ for the ‘artificial’ edges of $\tau_n(v)$. Write $F'_i = F_i \setminus E_v$, $F_v := \bigcup_{i \in [s(v)]} F'_i$, and let us write $X_{i_1}, \dots, X_{i_{l_i}}$ for the finitely many components of $X_i - (E_v \cup F'_i)$ other than \star (Lemma 14.1.11). Let us write \mathcal{V}_v for the collection of all these X_{i_k} . We have obtained a decomposition $\mathcal{P}_v = \mathcal{V}_v \cup F_v$ of $\eta(v)$ into edge disjoint standard subspace \mathcal{V}_v and newly displayed edges F_v .²⁷ Repeat this procedure for each $v \in V(G_n)$.

Our next task is to turn these partitions into an Eulerian decomposition (G_{n+1}, η_{n+1}) of X . For this, we first define an auxiliary decomposition (G'_{n+1}, η'_{n+1}) , where the underlying graph G'_{n+1} has vertex and edge set $E(G'_{n+1}) := F_{n+1} \sqcup D_n$ as follows:

- $V(G_{n+1}) := \bigsqcup_{v \in V(G_n)} \mathcal{V}_v$ and
- $F_{n+1} := F_n \sqcup \bigsqcup_{v \in V(G_n)} F_v$.

For the map η'_{n+1} we take the natural candidate: for $e \in F_n \cup D_n$, define $\eta'_{n+1}(e) := \eta_n(e)$. And for $x \in \mathcal{P}_v$ (vertices and newly displayed edges alike) define $\eta'_{n+1}(x) = x$. Next, note

²⁷Note that some X_{i_k} is allowed to consist of a single edge, which does not count as being displayed.

that the map $\varrho'_n: G'_{n+1} \rightarrow G_n$ defined by $\varrho'_n \upharpoonright (F_n \cup D_n) := \text{id}$ and $\varrho'^{-1}_n(v) := \mathcal{P}_v$ is a surjective map satisfying (Q1) and (Q2) of a contraction map, cf. Definition 14.3.13. As our next step, we need to define vertex-edge-incidences for G'_{n+1} so that

- (a) (E6) and (E7) are satisfied, i.e. (G'_{n+1}, η'_{n+1}) is indeed a decomposition of X according to Definition 14.3.2,
- (b) (Q3) and (Q4) are satisfied for ϱ'_n , i.e. ϱ'_n is a contraction map from G'_{n+1} to G_n according to Definition 14.3.13, and so that
- (c) ϱ'_n is η -compatible according to Definition 14.3.14.

So let us consider an arbitrary edge $f \in E(G'_{n+1})$. Suppose first that $f \in F_n \cup D_n$. Then $f \in E(G_n)$ where it is incident to $f_{G_n}(0) = v$ and $f_{G_n}(1) = w$ say (not necessarily distinct). In order to define $f_{G'_{n+1}}(0)$, note that $f \in \tau_n(v)$, and hence there is a unique $X_i \in \mathcal{S}_v$ with $f \in E(X_i)$. Since $f \in E_v$, there is a unique component X_{ik} of $X_i - (E_v \cup F'_i)$ such that $f(0) \in X_{ik}$, and so we may define $f_{G'_{n+1}}(0) := X_{ik}$. This assignment satisfies (E6) or (E7) respectively by construction, as well as (Q3). Suppose next that $f \in F_{n+1} \setminus F_n$. By definition of F_{n+1} , there is a unique $v \in V(G_n)$ such that $f \in F_v$. This means in turn, that $f \in E(X_i)$ for some $X_i \in \mathcal{S}_v$, and so there are unique components $X_{i,k}, K_{ij}$ of $X_i - (E_v \cup F'_i)$ such that $f(0) \in X_{i,k}$ and $f(1) \in X_{i,j}$. Hence, by defining $f_{G'_{n+1}}(0) = X_{i,k}$ and $f_{G'_{n+1}}(1) = X_{i,j}$, we see that this assignment satisfies (E6) as well as (Q4). Hence, we have verified (a) and (b), and now that ϱ'_n is indeed a contraction map, it is clear that it also is η -compatible, for we have

$$\eta_n(x) = \bigcup \{ \eta'_{n+1}(y) : y \in \varrho'^{-1}_n(v) \}$$

for all $x \in V(G_n) \cup E(G_n)$ by construction.

This completes the construction of G'_{n+1} and $\varrho'_n: G'_{n+1} \rightarrow G_n$. Next, we claim that every vertex in G'_{n+1} has even degree: indeed, for every vertex v of G'_{n+1} with corresponding tile $\eta'_n(v) = X_{ik}$ with $X_{ik} \subseteq X_i \in \mathcal{S}_{\varrho'_n(v)}$, we have that the edges $E_{G'_{n+1}}(v)$ incident with v in G'_{n+1} correspond precisely to the edges in $(E_v \cup F_v) \cap E(X_i)$ incident with the component X_{ik} . However, since X_i satisfies the even-cut condition by (ii), it follows that this is an even number of edges, and hence that v has even degree in G'_{n+1} .

For later use, note that it follows from (iv) that (G'_{n+1}, η'_{n+1}) satisfies (E10). Moreover, (G'_{n+1}, η'_{n+1}) also satisfies (E9): indeed, for every $w \in V_{n+1}$ with $\eta'_{n+1}(w) \subseteq X_i \in \mathcal{S}_v$ it is easy to verify that $\tau'_{n+1}(w)$ is a contraction of X_i ; since X_i satisfied the even-cut condition by (ii), so does $\tau'_{n+1}(w)$.

To turn G'_{n+1} into the final Eulerian multi-graph G_{n+1} , we now generously add parallel dummy edges in $D_{n+1} \setminus D_n$ in order to make the graph connected,²⁸ making sure that (E7), (Q4) and (Q5) hold for these new dummy edges. Indeed, to achieve connectedness of G_{n+1} is it sufficient, since G_n was connected, to arrange for (Q5), i.e. to show that ϱ_n

²⁸While dummy edges are introduced in parallel pairs when they emerge for the first time in G_{n+1} , we do not (and cannot) require them to remain parallel in G_{n+2} .

has connected fibres. Towards this, recall that every $\eta_n(v)$ for $v \in V(G_n)$ was connected by definition. Let \mathcal{U}_v be the family of components of $\{Y - E_v : Y \in \mathcal{S}_v\}$. Then \mathcal{U}_v is a finite family of continua covering $\eta_n(v)$, and hence its intersection graph $G_{\mathcal{U}_v}$ on $\eta_n(v)$ is connected. Pick a spanning tree T_v for $G_{\mathcal{U}_v}$. For every edge $g = ab \in E(T_v)$ pick a point $x_g \in a \cap b \neq \emptyset$ in the overlap of the corresponding sets and then add two parallel dummy edges d^1, d^2 to G_{n+1} with associated point $\eta_{n+1}(d^1) = x_g = \eta_{n+1}(d^2)$ and incidences so that $d^1(0) = d^2(0) \subseteq a$ and $d^1(1) = d^2(1) \subseteq b$.

Then it is clear that G_{n+1} is connected, and since we added new dummy edges in pairs, G_{n+1} is still even. Thus, we have verified that G_{n+1} is Eulerian, and so (G_{n+1}, η_{n+1}) is an Eulerian decomposition of X extending (G_n, η_n) and satisfying (E10). Finally, it remains to check that also (E9) holds true for (G_{n+1}, η_{n+1}) . But this now follows easily from the fact that (G'_{n+1}, η'_{n+1}) satisfied (E9): indeed, since new dummy edges only occur in pairs, it follows that for every $w \in V(G_{n+1}) = V(G'_{n+1})$, the truncations τ_{n+1} and $\tau'_{n+1}(w)$ differ only by a finite family of edges, which come in parallel pairs between \star and (pairwise) the same point on the ground set on $\eta_{n+1}(w)$. It is clear that the even-cut condition is unaffected by these changes.

But now, since (E10) implies that $w(G_n, \eta_n) \leq (\frac{2}{3})^{\lfloor n/3 \rfloor} \rightarrow 0$, it follows that (A1) and (A2) of Definition 14.3.5 are satisfied, i.e. $((G_n, \eta_n) : n \in \mathbb{N})$ is an approximating sequence of Eulerian decompositions for X . This completes the proof. \square

14.5.5. Outlook. The techniques introduced in this chapter for one-dimensional continua lead to an abstract framework and to a technical conjecture, the truth of which implies the truth of the Eulerianity conjecture.

DEFINITION 14.5.27. The *core-size* of a Peano continuum X is the real number $\text{core}(X) = \sup \{\text{diam}(C) : C \text{ a connected component of } \mathfrak{G}(X)\}$. For a collection of Peano continua \mathcal{U} , we write $\mathfrak{G}\text{-mesh}(\mathcal{U}) = \sup \{\text{core}(X) : X \in \mathcal{U}\}$.

DEFINITION 14.5.28. An *even-cut decomposition* of a Peano continuum X is a finite cover \mathcal{U} of X consisting of edge-disjoint standard subcontinua each of which has the even-cut property. A class \mathcal{C} of Peano continua is *closed under even-cut decompositions* if every $X \in \mathcal{A}$ satisfies the even-cut property and admits even-cut decompositions \mathcal{U} of arbitrarily small $\mathfrak{G}\text{-mesh}(\mathcal{U})$ such that $U \in \mathcal{C}$ for all $U \in \mathcal{U}$.

The results of this Chapter 14.5 can then summarised as follows:

THEOREM 14.5.29. *The class of all one-dimensional Peano continua with the even-cut property is closed under even-cut decompositions.* \square

THEOREM 14.5.30. *If \mathcal{C} is a class of Peano continua closed under even-cut decompositions, then the Eulerianity conjecture holds for every $X \in \mathcal{C}$.* \square

Indeed, Theorem 14.5.29 follows by iterative applications of Theorem 14.5.23, and Theorem 14.5.30 follows as in the proof of Theorem 14.5.1 above, noting that by Lemma 14.5.18, for every Peano continuum X and every $\varepsilon > 0$ there is a finite edge set $F \subseteq E(X)$ such that $\text{diam}(C) < \text{core}(X) + \varepsilon$ for every component C of $X - F$.

CONJECTURE 14.5.31. *The class \mathcal{C} of all Peano continua with the even-cut property is closed under even-cut decompositions.*

In other words, we conjecture that every Peano continuum X satisfying the even-cut condition admits, for every $\varepsilon > 0$, a finite cover \mathcal{U} of edge-disjoint standard subcontinua of X all satisfying the even-cut condition with $\mathfrak{G}\text{-mesh}(\mathcal{U}) < \varepsilon$.

By Theorem 14.5.30, the truth of Conjecture 14.5.31 implies the truth of Conjecture 14.1.2.

Bibliography

- [1] H. Abobaker and W. Charatonik. Hereditarily irreducible maps. *preprint*, 2019.
- [2] B. Alspach. Research problems. *Discrete Mathematics*, 50:115, 1984.
- [3] B. Alspach. Unsolved problem 4.5. *Annals of Discrete Mathematics*, 27:464, 1985.
- [4] B. Alspach. The wonderful Walecki construction. *Bull. Inst. Combin. Appl.*, 52:7–20, 2008.
- [5] B. Alspach, J. C. Bermond, and D. Sotteau. Decomposition into cycles I: Hamilton decompositions. In *Cycles and rays*, pages 9–18. Springer, 1990.
- [6] B. Alspach and C. D. Godsil. *Cycles in graphs*, volume 27. Elsevier, 1985.
- [7] R. D. Anderson. A characterization of the universal curve and a proof of its homogeneity. *Annals of Mathematics*, pages 313–324, 1958.
- [8] R. D. Anderson. One-dimensional continuous curves and a homogeneity theorem. *Annals of Mathematics*, pages 1–16, 1958.
- [9] T. Andreae. Bemerkung zu einem Problem aus der Theorie der unendlichen Graphen. In *Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg*, volume 46, page 91. Springer, 1977.
- [10] T. Andreae. Über eine Eigenschaft lokalfiniten, unendlicher Bäume. *Journal of Combinatorial Theory, Series B*, 27(2):202–215, 1979.
- [11] T. Andreae. On a problem of R. Halin concerning infinite graphs II. *Discrete Mathematics*, 29(3):219–227, 1980.
- [12] T. Andreae. On the Reconstruction of Locally Finite Trees. *Journal of Graph Theory*, 5(2):123–135, 1981.
- [13] T. Andreae. Simpler counterexamples to the edge-reconstruction conjecture for infinite graphs. *Journal of Combinatorial Theory, Series B*, 32(3):258 – 263, 1982.
- [14] T. Andreae. On reconstructing rooted trees. *Journal of Combinatorial Theory, Series B*, 62(2):183 – 198, 1994.
- [15] T. Andreae. On disjoint configurations in infinite graphs. *Journal of Graph Theory*, 39(4):222–229, 2002.
- [16] T. Andreae. Classes of locally finite ubiquitous graphs. *Journal of Combinatorial Theory, Series B*, 103(2):274–290, 2013.
- [17] S. Baldwin and R. E. Beaudoin. Countable dense homogeneous spaces under Martin’s axiom. *Israel Journal of Mathematics*, 65(2):153–164, 1989.
- [18] J. E. Baumgartner, A. Hajnal, and A. Mate. Weak saturation properties of ideals. pages 137–158. *Colloq. Math. Soc. János Bolyai*, Vol. 10, 1975.
- [19] M. Bell and K. Kunen. On the pi character of ultrafilters. *CR Math. Rep. Acad. Sci. Canada*, 3(6):351–356, 1981.
- [20] E. Berger and H. Bruhn. Eulerian edge sets in locally finite graphs. *Combinatorica*, 31(1):21–38, 2011.

- [21] J. C. Bermond. Hamiltonian decompositions of graphs, directed graphs and hypergraphs. *Annals of Discrete Mathematics*, 3:21–28, 1978.
- [22] J. C. Bermond, O. Favaron, and M. Maheo. Hamiltonian decomposition of Cayley graphs of degree 4. *Journal of Combinatorial Theory, Series B*, 46(2):142–153, 1989.
- [23] N. Biggs, E. K. Lloyd, and R. J. Wilson. *Graph Theory, 1736-1936*. Oxford University Press, 1986.
- [24] R. Bing. Complementary domains of continuous curves. *Fundamenta Mathematicae*, 36(1):303–318, 1949.
- [25] R. Bing. Partitioning a set. *Bulletin of the American Mathematical Society*, 55(12):1101–1110, 1949.
- [26] R. Bing. Partitioning continuous curves. *Bulletin of the American Mathematical Society*, 58(5):536–556, 1952.
- [27] B. Bollobás. *Graph theory*, volume 63 of *Graduate Texts in Mathematics*. Springer-Verlag, New York-Berlin, 1979. An introductory course.
- [28] J. A. Bondy and R. Hemminger. Reconstructing infinite graphs. *Pacific Journal of Mathematics*, 52(2):331–340, 1974.
- [29] J. A. Bondy and R. L. Hemminger. Graph reconstruction – a survey. *Journal of Graph Theory*, 1(3):227–268, 1977.
- [30] N. Bowler, J. Carmesin, and R. Christian. Infinite graphic matroids. *Combinatorica*, 38(2):305–339, 2018.
- [31] N. Bowler, C. Elbracht, J. Erde, P. Gollin, K. Heuer, M. Pitz, and M. Teegen. Ubiquity in graphs I: Topological ubiquity of trees. in preparation.
- [32] N. Bowler, C. Elbracht, J. Erde, P. Gollin, K. Heuer, M. Pitz, and M. Teegen. Ubiquity in graphs II: Ubiquity of graphs with non-linear end structure. in preparation.
- [33] N. Bowler, C. Elbracht, J. Erde, P. Gollin, K. Heuer, M. Pitz, and M. Teegen. Ubiquity in graphs III: Ubiquity of a class of locally finite graphs. in preparation.
- [34] N. Bowler, C. Elbracht, J. Erde, P. Gollin, K. Heuer, M. Pitz, and M. Teegen. Ubiquity in graphs IV: Ubiquity of graphs of bounded tree-width. in preparation.
- [35] N. Bowler, J. Erde, P. Heinig, F. Lehner, and M. Pitz. A counterexample to the reconstruction conjecture for locally finite trees. *Bull. Lond. Math. Soc.*, 49(4):630–648, 2017.
- [36] N. Bowler, J. Erde, P. Heinig, F. Lehner, and M. Pitz. Non-reconstructible locally finite graphs. *Journal of Combinatorial Theory, Series B*, 133:122–152, 2018.
- [37] N. Bowler, S. Geschke, and M. Pitz. Minimal obstructions for normal spanning trees. *Fund. Math.*, 241(3):245–263, 2018.
- [38] H. Bruhn, R. Diestel, M. Kriesell, R. Pendavingh, and P. Wollan. Axioms for infinite matroids. *Advances in Mathematics*, 239:18–46, 2013.
- [39] H. Bruhn and M. Stein. On end degrees and infinite cycles in locally finite graphs. *Combinatorica*, 27(3):269, 2007.
- [40] D. Bryant, S. Herke, B. Maenhaut, and B. Webb. On Hamilton decompositions of infinite circulant graphs. *arXiv preprint arXiv:1701.08506*, 2017.
- [41] W. Bula, J. Nikiel, and E. Tymchatyn. The Königsberg bridge problem for Peano continua. *Canadian Journal of Mathematics*, 46(6):1175–1187, 1994.
- [42] R. Buskirk, J. Nikiel, and E. Tymchatyn. Totally regular curves as inverse limits. *Houston J. Math.*, 18(3):319–327., 1992.
- [43] G. Călinescu, A. Dumitrescu, and J. Pach. Reconfigurations in graphs and grids. *SIAM Journal on Discrete Mathematics*, 22(1):124–138, 2008.
- [44] C. Capel. Inverse limit spaces. *Duke Mathematical Journal*, 21(2):233–245, 1954.

- [45] P. A. Catlin. Super-Eulerian graphs: a survey. *J. Graph Theory*, 16(2):177–196, 1992.
- [46] C. Champetier. Propriétés statistiques des groupes de présentation finie. *Adv. Math.*, 116(2):197–262, 1995.
- [47] J. J. Charatonik. Monotone mappings of universal dendrites. *Topology and its Applications*, 38(2):163–187, 1991.
- [48] G. L. Chia and Q. R. Yu. On the number of Hamilton cycles in cubic graphs. In *Proceedings of the Twenty-sixth Southeastern International Conference on Combinatorics, Graph Theory and Computing (Boca Raton, FL, 1995)*, volume 110, pages 13–32, 1995.
- [49] R. Christian, R. B. Richter, and B. Rooney. The planarity theorems of MacLane and Whitney for graph-like continua. *the electronic journal of combinatorics*, 17(1):12, 2010.
- [50] V. Chvátal. Tough graphs and hamiltonian circuits. *Discrete Math*, 5(3):215–228, 1973.
- [51] R. Diestel. The cycle space of an infinite graph. *Combinatorics, Probability and Computing*, 14(1-2):59–79, 2005.
- [52] R. Diestel. Locally finite graphs with ends: A topological approach, II. applications. *Discrete Mathematics*, 310(20):2750–2765, 2010.
- [53] R. Diestel. Locally finite graphs with ends: A topological approach, I. basic theory. *Discrete Mathematics*, 311(15):1423–1447, 2011.
- [54] R. Diestel. *Graph Theory*. Springer, 5th edition, 2016.
- [55] R. Diestel and D. Kühn. Graph-theoretical versus topological ends of graphs. *Journal of Combinatorial Theory, Series B*, 87:197–206, 2003.
- [56] R. Diestel and D. Kühn. On infinite cycles I. *Combinatorica*, 24(1):68–89, 2004.
- [57] R. Diestel and D. Kühn. On infinite cycles II. *Combinatorica*, 24(1):91–116, 2004.
- [58] R. Diestel and I. Leader. Normal spanning trees, Aronszajn trees and excluded minors. *Journal of the London Mathematical Society*, 63(1):16–32, 2001.
- [59] R. Diestel and P. Sprüssel. Locally finite graphs with ends: A topological approach, iii. fundamental group and homology. *Discrete Mathematics*, 312(1):21–29, 2012.
- [60] R. Diestel and R. Thomas. Excluding a countable clique. *Journal of Combinatorial Theory, Series B*, 76(1):41–67, 1999.
- [61] G. A. Dirac. In abstrakten Graphen vorhandene vollständige 4-Graphen und ihre Unterteilungen. *Math. Nachr.*, 22:61–85, 1960.
- [62] R. Engelking. *Dimension theory*. North-Holland Publishing Co., Amsterdam-Oxford-New York; PWN—Polish Scientific Publishers, Warsaw, 1978. North-Holland Mathematical Library, 19.
- [63] R. Engelking. *General topology*, volume 6 of *Sigma Series in Pure Mathematics*. Heldermann Verlag, Berlin, second edition, 1989.
- [64] R. C. Entringer and H. Swart. Spanning cycles of nearly cubic graphs. *Journal of Combinatorial Theory, Series B*, 29(3):303–309, 1980.
- [65] J. Erde and F. Lehner. Hamilton decompositions of infinite 4-regular Cayley graphs. In preparation.
- [66] J. Erde, F. Lehner, and M. Pitz. Hamilton decompositions of one-ended Cayley graphs. *Journal of Combinatorial Theory, Series B*, To appear.
- [67] P. Erdős, T. Grunwald, and E. Vazsonyi. Über Euler-Linien unendlicher Graphen. *Journal of Mathematics and Physics*, 17:59–75, 1938.
- [68] B. Espinoza, P. Gartside, M. Kovan-Bakan, and A. Mamatelashvili. Strong arcwise connectedness. *Houston Journal of Mathematics*, 43(2):577–610, 2017.
- [69] B. Espinoza, P. Gartside, and A. Mamatelashvili. n -arc connected spaces. *Colloquium Mathematicum*, 130:221–240, 2013.

- [70] B. Espinoza, P. Gartside, and M. Pitz. Graph-like compacta: Characterizations and eulerian loops. *arXiv preprint arXiv:1609.00933*, 2016.
- [71] B. Espinoza and E. Matsushashi. Arcwise increasing maps. *Topology and its Applications*, 190:74–92, 2015.
- [72] A. Fedeli and A. Le Donne. Eulerian paths and a problem concerning n -arc connected spaces. *Topology and its Applications*, 161:159–162, 2014.
- [73] J. Fisher. A counterexample to the countable version of a conjecture of Ulam. *Journal of Combinatorial Theory*, 7(4):364–365, 1969.
- [74] J. Fisher, R. L. Graham, and F. Harary. A simpler counterexample to the reconstruction conjecture for denumerable graphs. *Journal of Combinatorial Theory, Series B*, 12(2):203–204, 1972.
- [75] L. Fuchs. *Abelian groups*. Springer, 2015.
- [76] P. Gartside, A. Mamatelashvili, and M. Pitz. n -arc connected graphs. *arXiv preprint arXiv:1801.00179*, 2018.
- [77] P. Gartside and M. Pitz. Eulerian spaces. *arXiv preprint arXiv:1904.02645*, 2019.
- [78] P. Gartside and M. Pitz. n -arc and n -circle connected graph-like spaces. *Topology and its Applications*, 256:7–25., 2019.
- [79] A. Georgakopoulos. Connected but not path-connected subspaces of infinite graphs. *Combinatorica*, 27(6):683–698, 2007.
- [80] A. Georgakopoulos. Topological circles and Euler tours in locally finite graphs. *Electronic J. Comb.*, 16(1):40, 2009.
- [81] A. Georgakopoulos. Graph topologies induced by edge lengths. *Discrete Math.*, 311(15):1523–1542, 2011.
- [82] P. Gollin and K. Heuer. k -connected sets in infinite graphs. In preparation.
- [83] R. J. Gould. A look at cycles containing specified elements of a graph. *Discrete Math.*, 309(21):6299–6311, 2009.
- [84] R. Häggkvist and C. Thomassen. Circuits through specified edges. *Discrete Math.*, 41(1):29–34, 1982.
- [85] R. Halin. Über die Maximalzahl fremder unendlicher Wege in Graphen. *Mathematische Nachrichten*, 30(1-2):63–85, 1965.
- [86] R. Halin. Die Maximalzahl fremder zweiseitig unendlicher Wege in Graphen. *Mathematische Nachrichten*, 44(1-6):119–127, 1970.
- [87] R. Halin. A problem in infinite graph-theory. In *Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg*, volume 43, pages 79–84. Springer, 1975.
- [88] R. Halin. S -functions for graphs. *J. Geometry*, 8(1-2):171–186, 1976.
- [89] R. Halin. Miscellaneous problems on infinite graphs. *Journal of Graph Theory*, 35(2):128–151, 2000.
- [90] F. Harary. On the reconstruction of a graph from a collection of subgraphs. In *Theory of Graphs and its Applications (Proc. Sympos. Smolenice, 1963)*, pages 47–52, 1964.
- [91] F. Harary, A. J. Schwenk, and R. L. Scott. On the reconstruction of countable forests. *Publ. Inst. Math.*, 13(27):39–42, 1972.
- [92] O. G. Harrold, Jr. A note on strongly irreducible maps of an interval. *Duke Math. J.*, 6:750–752, 1940.
- [93] O. G. Harrold, Jr. A mapping characterization of Peano spaces. *Bull. Amer. Math. Soc.*, 48:561–566, 1942.
- [94] S. H. Hechler. Generalizations of almost disjointness, \mathfrak{c} -sets, and the Baire number of $\beta\mathbb{N} \setminus \mathbb{N}$. *General Topology and its Applications*, 8(1):93–110, 1978.

- [95] K. Heuer. Hamiltonicity in locally finite graphs: Two extensions and a counterexample. *The Electronic Journal of Combinatorics*, 25(3):3–13, 2018.
- [96] D. Hilbert. Über die stetige Abbildung einer Linie auf ein Flächenstück. *Math. Ann.*, 38(3):459–460, 1891.
- [97] F. Jaeger. A note on sub-Eulerian graphs. *J. Graph Theory*, 3(1):91–93, 1979.
- [98] M. Kapovich and B. Kleiner. Hyperbolic groups with low-dimensional boundary. *Ann. Sci. École Norm. Sup. (4)*, 33(5):647–669, 2000.
- [99] K.-I. Kawarabayashi. One or two disjoint circuits cover independent edges. Lovász-Woodall conjecture. *J. Combin. Theory Ser. B*, 84(1):1–44, 2002.
- [100] P. Kelly. A congruence theorem for trees. *Pacific J. Math.*, 7(1):961–968, 1957.
- [101] P. Knappe and M. Pitz. Circuits through prescribed edges. *arXiv preprint arXiv:1810.09323*, 2018.
- [102] D. Kornhauser, G. Miller, and P. Spirakis. Coordinating pebble motion on graphs, the diameter of permutation groups, and applications. In *Proceedings of the 25th Annual Symposium on Foundations of Computer Science, 1984*, SFCS '84, pages 241–250. IEEE Computer Society, 1984.
- [103] A. Kotzig. Every Cartesian product of two circuits is decomposable into two hamiltonian circuits. *Centre de Recherche Mathématiques, Montreal*, 1973.
- [104] J. Krasinkiewicz. On two theorems of dyer. In *Colloquium Mathematicum*, volume 2, pages 201–208, 1986.
- [105] I. Kríž and R. Thomas. The menger-like property of the tree-width of infinite graphs. *Journal of Combinatorial Theory, Series B*, 52(1):86–91, 1991.
- [106] K. Kunen. Set theory, volume 34 of studies in logic, 2011.
- [107] K. Kuratowski. *Topology*, volume 2. PWN–Polish Scientific Publishers, 1968.
- [108] H.-J. Lai. Eulerian subgraphs containing given edges. *Discrete Math.*, 230(1-3):63–69, 2001. Paul Catlin memorial collection (Kalamazoo, MI, 1996).
- [109] J. Lake. A problem concerning infinite graphs. *Discrete Mathematics*, 14(4):343–345, 1976.
- [110] R. Laver. On Fraïssé’s order type conjecture. *Ann. of Math.*, 93:89–111, 1971.
- [111] R. Laver. Better-quasi-orderings and a class of trees. *Studies in foundations and combinatorics*, 1:31–48, 1978.
- [112] F. Laviolette. Decompositions of infinite graphs. II. Circuit decompositions. *J. Combin. Theory Ser. B*, 94(2):278–333, 2005.
- [113] J. Liu. Hamiltonian decompositions of Cayley graphs on abelian groups. *Discrete Mathematics*, 131(1-3):163–171, 1994.
- [114] J. Liu. Hamiltonian decompositions of Cayley graphs on abelian groups of odd order. *Journal of Combinatorial Theory, Series B*, 66(1):75–86, 1996.
- [115] J. Liu. Hamiltonian decompositions of Cayley graphs on abelian groups of even order. *Journal of Combinatorial Theory, Series B*, 88(2):305–321, 2003.
- [116] L. Lovász. Problem 5. *Period. Math. Hungar.*, 4:82, 1973.
- [117] J. C. Mayer, L. G. Oversteegen, and E. D. Tymchatyn. The Menger curve. Characterization and extension of homeomorphisms of non-locally-separating closed subsets. *Dissertationes Math. (Rozprawy Mat.)*, 252:45, 1986.
- [118] K. Menger. Zur allgemeinen Kurventheorie. *Fundamenta Mathematicae*, 10(1):96–115, 1927.
- [119] B. Mohar. Hamiltonicity of infinite graphs. http://www.fmf.uni-lj.si/~mohar/Problems/P0703_HamiltonicityInfinite.html and http://www.openproblemgarden.org/op/infinite_uniquely_hamiltonian_graphs, accessed on 2019-04-19.

- [120] S. B. Nadler, Jr. Induced universal maps and some hyperspaces with the fixed point property. *Proc. Amer. Math. Soc.*, 100(4):749–754, 1987.
- [121] S. B. Nadler, Jr. *Continuum theory: an introduction*. CRC Press, 1992.
- [122] C. S. J. A. Nash-Williams. Abelian groups, graphs and generalized knights. In *Mathematical Proceedings of the Cambridge Philosophical Society*, volume 55, pages 232–238. Cambridge Univ Press, 1959.
- [123] C. S. J. A. Nash-Williams. Decomposition of graphs into closed and endless chains. *Proceedings of the London Mathematical Society*, 3(1):221–238, 1960.
- [124] C. S. J. A. Nash-Williams. Decomposition of graphs into two-way infinite paths. *Canad. J. Math.*, 15:479–485, 1963.
- [125] C. S. J. A. Nash-Williams. On well-quasi-ordering infinite trees. *Mathematical proceedings of the Cambridge philosophical society*, 61(3):697–720, 1965.
- [126] C. S. J. A. Nash-Williams. Reconstruction of locally finite connected graphs with at least three infinite wings. *Journal of Graph Theory*, 11(4):497–505, 1987.
- [127] C. S. J. A. Nash-Williams. Reconstruction of infinite graphs. *Discrete Mathematics*, 95(1):221–229, 1991.
- [128] C. S. J. A. Nash-Williams. Reconstruction of locally finite connected graphs with two infinite wings. *Discrete Mathematics*, 92(1):227–249, 1991.
- [129] J. Nikiel. Images of arcs—a nonseparable version of the Hahn-Mazurkiewicz theorem. *Fund. Math.*, 129(2):91–120, 1988.
- [130] J. Nikiel. Locally connected curves viewed as inverse limits. *Fundamenta Mathematicae*, 133(2):125–134, 1989.
- [131] G. Nöbling. Eine Verschärfung des n -Bein Satzes. *Fundamenta Mathematicae*, 18:23–28, 1931.
- [132] G. Nöbling. Reguläre Kurven als Bilder der Kreislinie. *Fund. Math.*, 20:30–46, 1933.
- [133] J. Oxley. *Matroid theory*, volume 21 of *Oxford Graduate Texts in Mathematics*. Oxford University Press, Oxford, second edition, 2011.
- [134] M. F. Pitz. Hamilton cycles in infinite cubic graphs. *Electron. J. Combin.*, 25(3):Paper 3.3, 2018.
- [135] N. Robertson and P. D. Seymour. Graph minors. II. Algorithmic aspects of tree-width. *J. Algorithms*, 7(3):309–322, 1986.
- [136] N. Robertson and P. D. Seymour. Graph minors. XX. Wagner’s conjecture. *Journal of Combinatorial Theory, Series B*, 92(2):325 – 357, 2004. Special Issue Dedicated to Professor W.T. Tutte.
- [137] N. Robertson and P. D. Seymour. Graph minors XXIII. Nash-Williams’ immersion conjecture. *J. Comb. Theory, Ser. B*, 100(2):181–205, 2010.
- [138] J. Roitman and L. Soukup. Luzin and anti-Luzin almost disjoint families. *Fundamenta Mathematicae*, 158:51–67, 1998.
- [139] B. Rothschild. The decomposition of graphs into a finite number of paths. *Canad. J. Math.*, 17:468–479, 1965.
- [140] G. Sabidussi. Infinite Euler graphs. *Canad. J. Math.*, 16:821–838, 1964.
- [141] P. Scott and T. Wall. Topological methods in group theory. In *Homological group theory (Proc. Sympos., Durham, 1977)*, volume 36, pages 137–203, 1979.
- [142] J. Sheehan. The multiplicity of Hamiltonian circuits in a graph. *Recent advances in graph theory*, pages 477–480, 1975.
- [143] S. Shelah. A compactness theorem for singular cardinals, free algebras, whitehead problem and transversals. *Israel Journal of Mathematics*, 21(4):319–349, 1975.

- [144] J. Steprāns. Strong q -sequences and variations on Martin's Axiom. *Canadian Journal of Mathematics*, 37(4):730–746, 1985.
- [145] J. Steprāns and W. Watson. Homeomorphisms of manifolds with prescribed behaviour on large dense sets. *Bulletin of the London Mathematical Society*, 19(4):305–310, 1987.
- [146] Y. Sternfeld. Mappings in dendrites and dimension. *Houston J. Math.*, 19(3):483–497, 1993.
- [147] R. Stong. Hamilton decompositions of Cartesian products of graphs. *Discrete Mathematics*, 90(2):169–190, 1991.
- [148] R. Thomas. A counter-example to 'Wagner's conjecture' for infinite graphs. In *Mathematical Proceedings of the Cambridge Philosophical Society*, volume 103, pages 55–57. Cambridge University Press, 1988.
- [149] R. Thomas. Well-quasi-ordering infinite graphs with forbidden finite planar minor. *Transactions of the American Mathematical Society*, 312(1):279–313, 1989.
- [150] A. G. Thomason. Hamiltonian cycles and uniquely edge colourable graphs. In *Annals of Discrete Mathematics*, volume 3, pages 259–268. Elsevier, 1978.
- [151] C. Thomassen. Reconstructing 1-coherent locally finite trees. *Commentarii Mathematici Helvetici*, 53(1):608–612, 1978.
- [152] C. Thomassen. Independent dominating sets and a second Hamiltonian cycle in regular graphs. *Journal of Combinatorial Theory, Series B*, 72(1):104–109, 1998.
- [153] C. Thomassen and A. Vella. Graph-like continua, augmenting arcs, and Menger's theorem. *Combinatorica*, 28(5):595–623, 2008.
- [154] L. B. Treybig. Mappings of finite oscillation at local separating points. In *Topology and order structures, Part 2 (Amsterdam, 1981)*, volume 169 of *Math. Centre Tracts*, pages 81–89. Math. Centrum, Amsterdam, 1983.
- [155] L. B. Treybig and L. E. Ward, Jr. The Hahn-Mazurkiewicz problem. In *Topology and order structures, Part 1 (Lubbock, Tex., 1980)*, volume 142 of *Math. Centre Tracts*, pages 95–105. Math. Centrum, Amsterdam, 1981.
- [156] W. T. Tutte. On hamiltonian circuits. *Journal of the London Mathematical Society*, 1(2):98–101, 1946.
- [157] E. Tymchatyn. Characterizations of continua in which connected subsets are arcwise connected. *Transactions of the American Mathematical Society*, 222:377–388, 1976.
- [158] J. Van Mill. *The infinite-dimensional topology of function spaces*, volume 64. Elsevier, 2002.
- [159] J. E. Vaughan. Small uncountable cardinals and topology. In *Open problems in topology*, pages 195–218. North-Holland, Amsterdam, 1990. With an appendix by S. Shelah.
- [160] B. Velickovic. Oca and automorphisms of $P(\omega)/\text{fin}$. *Topology and its Applications*, 49(1):1–13, 1993.
- [161] L. E. Ward, Jr. A generalization of the Hahn-Mazurkiewicz theorem. *Proc. Amer. Math. Soc.*, 58:369–374, 1976.
- [162] L. E. Ward, Jr. An irreducible Hahn-Mazurkiewicz theorem. *Houston J. Math.*, 3(2):285–290, 1977.
- [163] M. E. Watkins and D. M. Mesner. Cycles and connectivity in graphs. *Canad. J. Math.*, 19:1319–1328, 1967.
- [164] E. Westlund. Hamilton decompositions of certain 6-regular Cayley graphs on abelian groups with a cyclic subgroup of index two. *Discrete Mathematics*, 312(22):3228–3235, 2012.
- [165] G. T. Whyburn. Sets of local separating points of a continuum. *Bulletin of the American Mathematical Society*, 39(2):97–100, 1933.
- [166] G. T. Whyburn. *Analytic topology*. American Mathematical Society Colloquium Publications, Vol. XXVIII. American Mathematical Society, Providence, R.I., 1963.

- [167] D. Witte and J. Gallian. A survey: Hamiltonian cycles in Cayley graphs. *Discrete Mathematics*, 51(3):293–304, 1984.
- [168] D. R. Woodall. A note on a problem of Halin's. *Journal of Combinatorial Theory, Series B*, 21(2):132–134, 1976.
- [169] D. R. Woodall. Circuits containing specified edges. *J. Combinatorial Theory Ser. B*, 22(3):274–278, 1977.
- [170] J. Xu, Z.-H. Chen, H.-J. Lai, and M. Zhang. Spanning trails in essentially 4-edge-connected graphs. *Discrete Appl. Math.*, 162:306–313, 2014.