Tree-decompositions, Hamilton circles and Orientations of infinite Graphs

Dissertation

zur Erlangung des Doktorgrades
an der Fakultät für Mathematik, Informatik
und Naturwissenschaften
der Universität Hamburg

vorgelegt
im Fachbereich Mathematik
von
Marcel Koloschin

Hamburg
2023

Prüfungskommission

Vorsitz	:	Professor Thomas Schmidt	(Universität Hamburg)
Erstgutachter und Betreuer	:	PD Dr. Max Pitz	(Universität Hamburg)
Zweitgutachter und stellv. Vorsitz	:	Professor Reinhard Diestel	(Universität Hamburg)
Mitglied	:	Professor Jörg Teschner	(Universität Hamburg)
Mitglied und Schriftführer	:	PD Dr. Stefan Geschke	(Universität Hamburg)

Datum der Disputation	:	15. November 2023

Contents

1 Introduction

This dissertation is a collection of three different topics from infinite graph theory.

In the first topic we investigate the question whether the end space of a graph can be described in a certain way by a tree-decomposition. There are several ways how tree-decompositions can be related to ends of a graph, arguably the most natural one is the property of *displaying* a set of ends in the sense that the ends of the decomposition tree correspond bijectively to the ends of the graph or a prescribed subset of the ends.

A recent result by Carmesin [9] from 2014 states that it is always possible to display the set of undominated ends of a graph G. However, a full characterisation of the subsets of ends that can be displayed remained open. Carmesin's research in this topic was generally motivated by the wide field of separation systems.

In the last year, Pitz introduced the technique of enveloping a given subgraph, a powerful tool to find for any subgraph of a graph another subgraph with the same ends in the closure in $|G|$ but with finite adhesion. Those envelopes are perfect candidates for parts of a tree-decomposition and indeed allowed Pitz to obtain a shorter constructive proof of Carmesins result above [38]. This new technique was the starting point to think about whether we can find an answer for the general problem which sets of ends can be displayed. Our first result was that the graphs for which the whole end space can be displayed are exactly those with a normal spanning tree. By a result of Diestel [16], the graphs G with normal spanning trees are exactly those for which $|G|$ is completely metrizable.

In joint work with Thilo Krill and Max Pitz we further showed that the subsets of ends that can be displayed for any given graph G are exactly the G_δ sets of ends in $|G|$. In turn, the G_δ sets of ends are exactly those sets $\Psi \subseteq \Omega(G)$ for which $|G|_\Psi$ is completely metrizable. That way, we obtain a natural extension of the first result. Another approach to this field is to look for whether we can make the ends, which are not displayed, live in different parts of the decomposition tree. In the best case, we can find a bijection such that for each part and for each end of the decomposition tree, there is exactly one end of G living in that part or end. This approach leads to the property of *representing* the set $\Omega(G)$. This property is a strengthening of the concept of *distinguishing* ends by Carmesin (Carmesin did not require bijectivity). It turns out that the ends can always be represented whenever any subset of ends can be distributed along the parts of any tree-decomposition. This leads to our characterisation of graphs with a representable end space.

Our second topic will be Hamilton circles in powers of infinite graphs. The nth power G^n of a graph G is obtained from G by adding an edge between any two vertices for which its distance in G is at most n. As early as in 1960, Sekanina proved by induction that the third power G^3 of any connected finite graph G has a Hamilton cycle [33]. Georgakopoulos conjectured that by using topological circles in the Freudenthal compactification $|G^3|$ this theorem should extend to all countable connected graphs [24]. This conjecture was also reiterated by Reinhard Diestel [17] and Bojan Mohar [32].

In my master thesis, I disproved this conjecture and characterized the trees that have a Hamilton circle in their third power. Nevertheless is this not sufficient for a full characterization of all countable graphs which are Hamiltonian in the third power, since the end space of the third power of arbitrary graphs is way more complex than the same space for trees. Our main result in this field is a characterisation of the rayless graphs with a Hamilton circle in their third power. Unlike the majority of proofs about rayless graphs in infinite graph theory, our proof of this characterisation is not by induction on the rank of the graph, but instead uses an involved direct construction of the desired Hamilton circle.

Our second main result in this field is that the fourth and higher powers of all countable trees are Hamiltonian.

Our third and final topic is Nash Williams' orientation theorem from 1960 [34]. It states that every finite $2k$-edge-connected graph has a k-arc-connected orientation.

The question whether this statement holds for infinite graphs remains unsolved until today. Thomassen had asked in 1985 whether there is a function $f\colon \mathbb{N} \to \mathbb{N}$ such that any $f(k)$-edge-connected multigraph has a k-arc-connected orientation [43] and indeed in 2016, Thomassen achieved a marvellous breakthrough towards the orientation theorem by proving that every $8k$-edge-connected multigraph has a k-arc-connected orientation [44], giving $f(k) \leq 8k$.

We will refine this result by establishing an improved bound of $f(k) = 4k$, and further show the optimal result $f(k) = 2k$ for the class of locally finite graphs with countably many ends, from which at most one has odd degree. Especially does this include the class of one ended locally finite graphs. Also I hope that some of the techniques we use may be helpful for further research in this field.

2 Basic definitions and tools

In this thesis, if not otherwise stated $G = (V, E)$ is always any (possibly infinite) graph with vertex set V and edge set E. For graph theoretic terms we follow the terminology in the book *Graph Theory* by Reinhard Diestel [14]. Whenever we refer to commonly known established results in the field of graph theory, such as the *star-comb* lemma from the next subsection, they can usually be found in this book.

2.1 Facts about infinite graphs

There are three different structural ways how a graph can be infinite. A graph can have infinitely many components, it can contain a vertex of infinite degree, or it can contain a one way infinite path, called a *ray*. Since all topics of this dissertation are either only interesting for connected graphs or can be reduced to connected graphs by looking at each component separately, only the last two ones are interesting for our applications.

Remark 2.1. Every connected infinite graph does either contain an infinite star or a ray.

Graphs without vertices of infinite degree are called *locally finite*.

The following *star-comb* lemma is used frequently to find certain connected substructures in infinite graphs.

Given a set of vertices U, a *comb attached to U* consists of a ray R together with infinitely many disjoint R–U paths (possibly trivial). A *star attached to U* is a subdivided infinite star with all leaves in U. We call the paths from the center of the star to its leaves the *subdivided leaves* of the star.

Lemma 2.2 (Star-Comb lemma [14, Lemma 8.2.2])**.** *Let U be an infinite set of vertices in a connected graph G. Then G contains a star or a comb attached to U.*

In the special case of a locally finite graph, we will always find a comb and for a rayless graph, we will always find a star.

For two sets of vertices $A, B \subseteq V(G)$, we define the connectivity $\kappa_G(A, B)$ as the minimum number of vertices in G separating A from B. Note that κ may be any infinite ordinal. There are several Menger-type results for infinite graphs known. We will only use the simple result that $\kappa_G(A, B)$ is always equal to the maximal number of disjoint $A - B-$paths [14, Lemma 8.4.1].

For vertices $v, w \in V(G)$, we define $\kappa_G(v, w)$ as the maximal number of internally disjoint $v - w-$paths.

A graph G in which for every two vertices v, w holds $\kappa_G(v, w) < \infty$ is called *finitely separable*.

Further, we define $\lambda_G(v, w)$ as the maximal number of edge-disjoint $v - w-$paths.

Lemma 2.3. *Let $G = (V, E)$ be any connected graph with an infinite subdivided star S_w with center w and leaves $(v_i)_{i\in\mathbb{N}}$. If v is another vertex in G for which there is another infinite subdivided star S_v with center v and leaves in $(v_i)_{i\in\mathbb{N}}$, then $\kappa_G(v, w) = \infty$.*

Proof. For each i for which there is a $v - v_i$ path in S_v, there is also a $w - v_i-$path in S_w. The union of these paths contains a $v - w-$path P_i. Without loss of generality, we assume that there is one such path P_i for every $i \in \mathbb{N}$.

To choose internally disjoint paths recursively, it remains to show that given finitely many paths $P_{i_1}, P_{i_2}, P_{i_3}, ..., P_{i_n}$, we can find one more path $P_{i_{n+1}}$ which is internally disjoint from $P_{i_1}, P_{i_2}, P_{i_3}, ..., P_{i_n}$. To find $P_{i_{n+1}}$, let V' be the union of the inner vertex sets of $P_{i_1}, P_{i_2}, P_{i_3}, ..., P_{i_n}$. Now since the subdivided leaves of a star are disjoint apart from the center and $v, w \notin V'$, only finitely many of them can meet V'. Hence for both stars almost all subdivided leaves are disjoint from V'. This implies that we can find one subdivided leaf of S_v and one of S_w with the same endvertex $v_{i_{n+1}}$ such that both of them are disjoint from V'. It follows that also $P_{i_{n+1}}$ is disjoint from V' □

2.2 Ends and directions

Two rays in a graph are *equivalent* if no finite set of vertices separates them; the corresponding equivalence classes of rays are the *ends* of G. If ω is an end of G and $R \in \omega$, we call R an ω-ray. The set of ends of a graph G is denoted by $\Omega = \Omega(G)$.

The *degree* $\deg(\omega)$ of an end ω is the supremum of the sizes of collections of pairwise disjoint rays in ω; Halin showed that this supremum is always attained, see [14, Theorem 8.2.5]. Ends are called *thin* if they have finite degree, and *thick* otherwise.

We say that a vertex v *dominates* a ray R if there is a subdivided star with center v and leaves in R. Whenever two rays R and R' are equivalent, a vertex dominates R if and only if it dominates R'. Thus we can say that a vertex *dominates* an end ω if and only if it dominates one ray (and thus all rays) from ω. In this case, we say ω is *dominated*.

For every finite vertex-set $S \subseteq V$ and every $\omega \in \Omega$, there is a unique component of $G - S$ that contains a tail of every ω-ray. We denote this component by $C(S, \omega)$ and say that ω *lives in* $C(S, \omega)$. Further we define $\Omega(S, \omega) = \{\varphi \in \Omega \colon C(S, \varphi) = C(S, \omega)\}$ as the set of all ends that live in $C(S, \omega)$. We put $\hat{C}(S, \omega) = C(S, \omega) \cup \Omega(S, \omega)$.

If H is a subgraph of G, then rays equivalent in H remain equivalent in G; in other words, every end of H can be interpreted as a subset of an end of G, so the natural inclusion map $\iota\colon \Omega(H) \to \Omega(G)$ is well-defined. A subgraph $H \subseteq G$ is *end-faithful* if this inclusion map ι is a bijection from $\Omega(H)$ onto $\partial H \subseteq \Omega(G)$.[1]

A *direction* on G is a function d that assigns to every finite $S \subseteq V$ one of the components of $G - S$ so that $d(S) \supseteq d(S')$ whenever $S \subseteq S'$. For every end ω, the map $S \mapsto C(S, \omega)$ is easily seen to be a direction. Conversely, every direction is defined by an end in this way:

Theorem 2.4 (Diestel & Kühn [15])**.** *For every direction d on a graph G there is an end ω such that $d(S) = C(S, \omega)$ for every finite $S \subseteq V(G)$.*

2.3 Topologies on infinite graphs

The set of ends of G will be denoted by $\Omega(G)$. The space $|G|$ is defined as $V \dot\cup \Omega \dot\cup E'$, where E' is the disjoint union of continuum sized sets, one set (v, w) for each edge vw of G. Also we choose for each edge a fixed bijection between (v, w) and the real interval $(0, 1)$. For each end and each finite vertex set S, we define $C(S, \omega)$ as the unique component of $G - S$, in which each ray of ω has a tail and $\mathring{E}_\epsilon(S, \omega)$ as the set of all inner points of $S - C(S, \omega)$ edges at distance less than ϵ from $C(S, \omega)$. There are several topologies on $|G|$. The three most commonly used ones are the following:

Definition 2.5. The topology *TOP* on $|G|$ is the topology induced by the following basic open sets:

1.) For each edge vw, any inverse image of an open subset of $(0, 1)$ under our fixed bijection.

2.) For each vertex v, the union of half-open partial edges $[v, z)$, one for each edge e at v with an inner point z of e.

3.) For an end ω, a finite vertex set S, the set of all vertices, ends and inner edge points in $C(S, \omega)$ together with a union of half-open partial edges $(z, v]$, one from every $S - C(S, \omega)$ edge uv with $v \in C(S, \omega)$.

Definition 2.6. The topology *MTOP* on $|G|$ is the topology induced by the following basic open sets:

1.) For each edge vw, any inverse image of an open subset of $(0, 1)$ under our fixed bijection.

2.) For each vertex v and each $\epsilon > 0$, the set of all points p in topological edges (v, w), such

[1]In the literature, the term end-faithful subgraph is sometimes used only for subgraphs $H \subseteq G$ with $\partial H = \Omega(G)$.

that the images of v and p under our fixed bijection have distance smaller than ϵ in $(0,1)$.
3.) For an end ω, a finite vertex set S and each $\epsilon > 0$, the set of all vertices, ends and inner edge points in $C(S,\omega)$ together with the inner edge points of $S - C(S,\omega)$ edges with distance less than ϵ from its endpoint in $C(S,\omega)$.

Definition 2.7. The topology *VTOP* on $|G|$ is the topology induced by the following basic open sets:
1.) For each edge vw, any inverse image of an open subset of $(0,1)$ under our fixed bijection.
2.) For each vertex v and each $\epsilon > 0$, the set of all points p in topological edges (v,w), such that the images of v and p under our fixed bijection have distance smaller than ϵ in $(0,1)$.
3.) For an end ω, a finite vertex set S, the set of all vertices, ends and inner edge points in $C(S,\omega)$ together with all inner points of $S - C(S,\omega)$ edges.

Those topologies are originally defined by Reinhard Diestel and can be found in [16]. We state some of the basic properties without proving them:

The three topologies coincide for locally finite graphs. Also $|G|$ is compact in TOP or MTOP if and only if it is locally finite.

Theorem 2.8. [16]
The following statements are equivalent in VTOP for any graph $G = (V,E)$.

- *$|G|$ is compact.*
- *For any finite $S \subseteq V$ the graph $G - S$ has only finitely many components.*
- *Every closed set of vertices is finite.*

All three topologies induce the same end space $\Omega(G)$, thus in the following corollary the topology is not specified.

Corollary 2.9. [16]
The subspace $\Omega(G)$ of $|G|$ is compact if and only if for every finite $S \subseteq V(G)$ only finitely many components of $G - S$ contain a ray.

Theorem 2.10. [16]
Let G be a connected graph.

- *In MTOP, $|\mathrm{G}|$ is metrizable if and only if G has a normal spanning tree.*
- *In VTOP, $|\mathrm{G}|$ is metrizable if and only if none of its ends is dominated.*

- *In TOP, $|\mathrm{G}|$ is metrizable if and only if G is locally finite.*

Given a set of vertices $U \subseteq V(G)$, we write ∂U for its boundary, i.e. the set of ends in $\overline{U}$. It is well-known that $\omega \in \partial U$ if and only if there is a comb attached to U with spine in ω. Further this topological viewpoint allows us to define a (possibly infinite) path as a subspace of $|G|$ that is homeomorphic to the unit interval $[0,1]$ and a (possibly infinite) *circle* as a subspace of $|G|$ that is homeomorphic to the unit circle S^1.

2.4 Tree orders and normal trees

The *tree order* of a tree T with root r is a partial order on $V(T)$ which is defined by setting $u \leq v$ if u lies on the unique path rTv from r to v in T. Given $n \in \mathbb{N}$, the *nth level* $T_{[n]}$ of T is the set of vertices at distance n from r in T, and by $T_{[\leq n]}$ we denote the union over the first n levels. The *down-closure* of a vertex v is the set $\lceil v \rceil := \{\, u \colon u \leq v \,\}$; its *up-closure* is the set $\lfloor v \rfloor := \{\, w \colon v \leq w \,\}$. The down-closure of v is always a finite chain, the vertex set of the path rTv. A ray $R \subseteq T$ starting at the root is called a *normal ray* of T.

A rooted spanning tree T of a graph G is *normal* in G if the endvertices of every edge of G are comparable in the tree order of T. Normal spanning trees are always end-faithful [14, Lemma 8.2.3].

A rooted, not necessarily spanning, tree T contained in a graph G is *normal* in G if the endvertices of every T-path in G are comparable in the tree-order of T. Here, for a given subgraph $H \subseteq G$, a path P in G is said to be an *H-path* if P is non-trivial and meets H exactly in its endvertices. Clearly, if T is spanning, this reduces to the earlier condition, as in this case all T-paths are chords. We remark that for a normal tree $T \subseteq G$ the neighbourhood $N(D)$ of every component D of $G - T$ forms a chain in T. The following result can be found in [29].

Theorem 2.11. *Let G be a connected graph. For every open cover $\mathcal{O}$ of $\Omega(G)$, there is a rayless normal tree T in G such that for every component C of $G - T$ there is a set $O \in \mathcal{O}$ such that $\partial C \subseteq O$.*

We say a set of vertices U in a graph G has *finite adhesion*, if and only if every component of $G - U$ has a finite neighbourhood in U.

Lemma 2.12. *Let G be a connected graph and T a rayless normal tree in G. Then T has finite adhesion in G. Moreover, for every finite set $U \subseteq V(G)$ there is a rayless normal tree $T^* \supseteq T$ in G such that $U \subseteq V(T^*)$.*

Proof. For the proof that T has finite adhesion in G, let C be any component of $G - T$. Since T is normal, the neighbourhood of C is a chain in the tree order of T, and this chain is finite because T^* is rayless.

Next, let T^* be a rayless normal tree in G extending the tree T which contains maximally many vertices from U. We show that T^* contains all vertices from U. Suppose for a contradiction that there is a vertex $u \in U$ with $u \notin V(T^*)$ and let C be the component of $G - T^*$ containing u. We showed in the first paragraph of this proof that the neighbourhood $N(C)$ of C is a finite chain in the tree order of T^*. Let v be its maximal element and v' a neighbour of v in C. Then the union of T^* with the edge vv' and a v'–u path in C is again a rayless normal tree with T as a subgraph, contradicting the maximality of T. □

2.5 Tree-decompositions

A [rooted] *tree-decomposition* of a graph G is a pair $\mathcal{T} = (T, \mathcal{V})$ where T is a [rooted] tree and $\mathcal{V} = (V_t \colon t \in T)$ is a family of vertex sets of G called *parts* such that the following holds (see also [14, §12.3]):

(T1) for every vertex v of G there exists $t \in T$ such that $v \in V_t$;

(T2) for every edge e of G there exists $t \in T$ such that $e \in G[V_t]$; and

(T3) $V_{t_1} \cap V_{t_3} \subseteq V_{t_2}$ whenever t_2 lies on the t_1–t_3 path in T.

Let $e = xy$ be any edge of T and let T_x and T_y be the two components of $T - e$ with $x \in T_x$ and $y \in T_y$. Each edge $e = xy$ of T in a tree-decomposition gives rise to a separator $X_e := V_x \cap V_y$ called the separator *induced by* the edge e, which separates $A_x = \bigcup_{t \in T_x} V_t$ from $A_y = \bigcup_{t \in T_y} V_t$. The tree-decomposition has *finite adhesion* if all separators of G induced by the edges of T are finite.

2.6 Topological notions

A subspace Y of a topological space X is *discrete* if every singleton of Y is open in the subspace topology.

A G_δ*-set* of a topological space X is a countable intersection of open sets. An F_σ*-set* is a countable union of closed sets. Note that the complement of a G_δ-set is always a F_σ-set and vice versa.

Lemma 2.13. *Let G be a graph and $\Psi \subseteq \Omega(G)$. Then $V(G) \cup \Psi$ is F_σ in $|G|$ if and only if $G \cup \Psi$ is F_σ in $|G|$.*

Proof. The backwards direction follows from that fact that $V \cup \Omega$ is closed in $|G|$, so $V \cup \Psi$ is closed in $G \cup \Psi$, and closed subsets of F_σ-sets are themselves F_σ.

Conversely, assume $V \cup \Psi = \bigcup_{n\in\mathbb{N}} X_n$ is a countable union of closed sets X_n of vertices and ends in $|G|$. Without loss of generality, we have $X_n \subseteq X_{n+1}$. Let $V_n = X_n \cap V(G)$. Then $\bigcup_{n\in\mathbb{N}} V_n = V(G)$ and $G = \bigcup_{n\in\mathbb{N}} G[V_n]$. But then the induced subsets $G[X_n] := G[V_n] \cup X_n$ are also closed in $|G|$, and so $G \cup \Psi = \bigcup_{n\in\mathbb{N}} G[X_n]$ is F_σ in $|G|$, too. $\square$

A set of vertices U in a graph G is *dispersed* if it can be separated from any ray in G by a finite set of vertices. This is equivalent to the property of U being closed in $|G|$.

3 End spaces and tree-decompositions

3.1 Introduction

In this chapter we settle the question up to which complexity the topological spaces $|G|$ formed by an infinite graph G together with its ends can still be encoded by tree-decompositions of finite adhesion of the underlying graph G.

To state our results more precisely, recall that a *separation* of a graph G is an unordered pair $\{A, B\}$ of sets of vertices in G such that $A \cup B = V(G)$ and G has no edge between $A \setminus B$ and $B \setminus A$, which is equivalent to saying that its *separator* $A \cap B$ separates A from B. The cardinal $|A \cap B|$ is the *order* of the separation $\{A, B\}$ and the sets A, B are its *sides*.

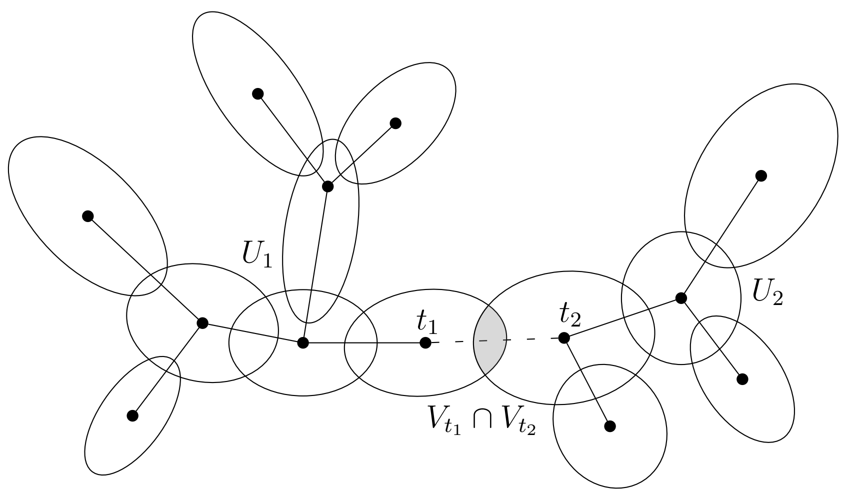

Figure 1: $V_{t_1} \cap V_{t_2}$ separates U_1 from U_2.

A longstanding quest in graph theory is to understand end spaces of infinite graphs that are not necessarily locally finite, cf. [11, 13, 16, 27, 29, 30, 39, 40].

Remember that the parts of a tree-decomposition mirror the separation properties of the tree: just like removing any edge $e = t_1t_2$ from T gives rise to two components T_1 and T_2 of $T - e$, so does removing $X_e := V_{t_1} \cap V_{t_2}$ from G separate any part of T_1 from any part of T_2, see Figure 1. More formally, writing $U_1 = \bigcup\{V_t : t \in T_1\}$ and $U_2 = \bigcup\{V_t : t \in T_2\}$, we require that $\{U_1, U_2\}$ is a separation of G with separator X_e. If all such separations are of finite order, we say the tree-decomposition has *finite adhesion*.

Now consider how the ends of a graph G interact with a tree-decomposition $\mathcal{T}$ of finite adhesion. As every edge e of T induces a finite order separation $\{A_e, B_e\}$ of G, any end of G has to choose one side of $T - e$, and we may visualize this decision by orienting e accordingly. Then for a fixed end, all the edges point either towards a unique node or towards a unique end of T, see Figure 2. In this way, each end of G *lives* in a part of $\mathcal{T}$ or *corresponds* to an end of T, and we may encode this correspondence by a map $f_\mathcal{T} \colon \Omega(G) \to V(T) \cup \Omega(T)$.

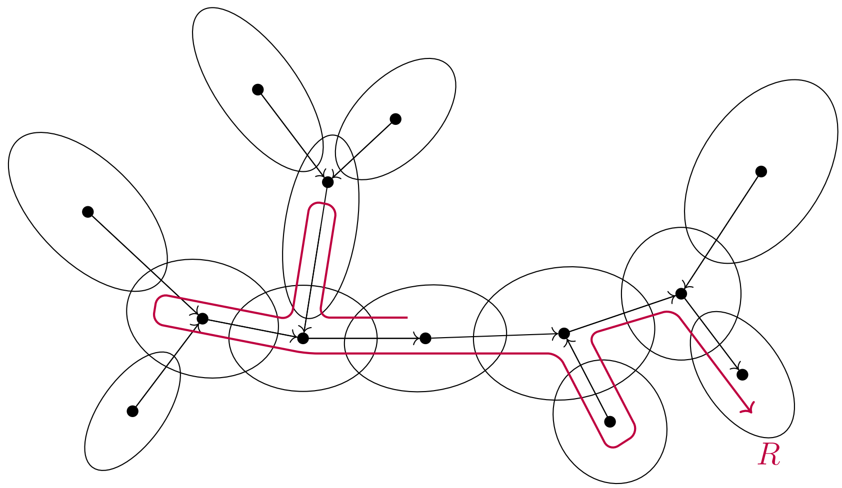

Figure 2: A ray R and its corresponding orientation of T

Tree-decompositions of finite adhesion have been used to study the structure of infinite graphs and their ends in e.g. [5, 7–11, 38, 42]. Of course, some tree-decompositions of finite adhesion carry more information about the ends than others. For one, information content may be measured in terms of injectivity of $f_\mathcal{T}$. Indeed, a tree-decomposition consisting of a single part contains zero information, whereas a tree-decomposition $\mathcal{T}$ of finite adhesion that *distinguishes* all the ends, i.e. where $f_\mathcal{T}$ is injective, contains more information about the end space – although it may still give false hints, as for example ends of T may not represent real ends of G. So even better would be a bijective $f_\mathcal{T}$, in which case we say that $\mathcal{T}$ *represents* the ends of G. On the other hand, while the trivial tree-decomposition into a single part always exists, some graphs G, such as the binary tree with one dominating vertex added to every rooted ray (cf. Section 3.9), are too complex to be distinguished or represented by a tree-decomposition of finite adhesion. Our first main result characterises precisely when these best-case scenarios occur; as a surprising by-product, we obtain that whenever a space $|G|$ can be distinguished by a tree-decomposition of finite adhesion, then it can also be represented. In fact, an even weaker condition suffices: As long as there is some tree-decomposition of finite adhesion into $\leq$1-*ended parts*, i.e. a tree-decomposition such that at most one end is mapped to any given part under $f_\mathcal{T}$, we also get a tree-decomposition representing $|G|$.

Let's call a set of vertices $U \subseteq V(G)$ *slender* if its closure $\overline{U} \subseteq |G|$ is scattered of finite Cantor-Bendixson rank; in other words, if successively taking the Cantor-Bendixson derivative of its closure $\overline{U} \subseteq |G|$ yields the empty set after finitely many iterations, cf. Section 2.6.

With this notion, our first main result reads as follows.

Theorem 3.1. *The following are equivalent for any connected graph G with at least one end:*

1. *There is a tree-decomposition of finite adhesion that represents $\Omega(G)$.*
2. *There is a tree-decomposition of finite adhesion that distinguishes $\Omega(G)$.*
3. *There is a tree-decomposition of finite adhesion into ≤ 1-ended parts.*
4. *$V(G)$ is a countable union of slender sets.*

It is clear that any assertion from (1) to (3) implies the next. The idea for (3) $\Rightarrow$ (4) is that for any fixed integer n, the union over all parts within distance n from the root is a slender set of vertices, and $V(G)$ clearly is a countable union of these sets. Thus, the main contribution behind Theorem 3.1 is the implication (4) $\Rightarrow$ (1), which employs recently developed techniques of *envelopes* from [30, 38] and *rayless normal trees* from [29]. The proof of Theorem 3.1 is given in Section 3.7.

A slightly different way to measure information captured by some tree-decomposition of finite adhesion is motivated by the observation that end spaces of trees are well-understood: They are precisely the completely ultra-metrizable spaces. This suggests preferring tree-decompositions $\mathcal{T}$ where $f_{\mathcal{T}}$ sends as many ends to $\Omega(T)$ as possible. In this case, there is hope to understand the subset $\Psi = f_{\mathcal{T}}^{-1}[\Omega(T)] \subseteq \Omega(G)$ called the *boundary* of the tree-decomposition, with the best case being that $\mathcal{T}$ [*homeomorphically*] *displays* its boundary, meaning that $f_{\mathcal{T}}$ restricts to a bijection [homeomorphism] between Ψ and $\Omega(T)$, cf. Figures 3 and 4.

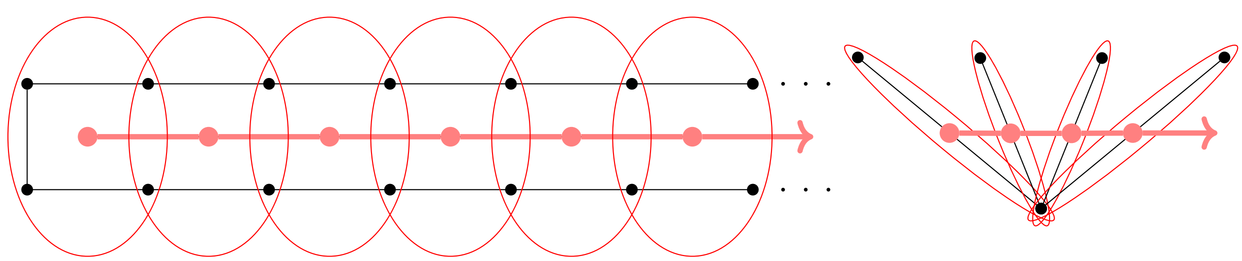

Figure 3: Examples of tree-decompositions (in red) of graphs (in black) failing to display their boundaries.

At first glance, however, it does not seem useful at all when $f_{\mathcal{T}}$ maps all ends of G into $\Omega(T)$ but the function is very much non-injective. However, this information is enough to

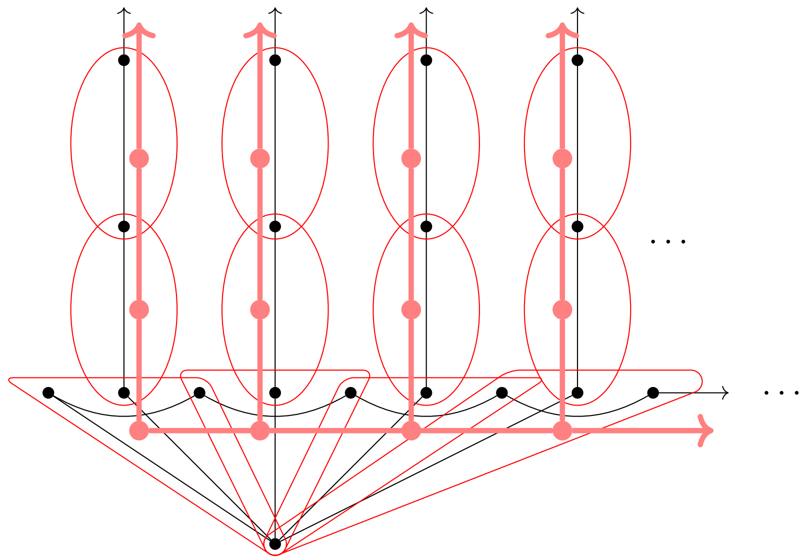

Figure 4: Example of a tree-decomposition (in red) that displays all ends of a countable star of rays (in black) but fails to display them homeomorphically.

guarantee a normal spanning tree, from which the space $|G|$ is easily understood. Indeed, given previous work in the field due to Jung and Diestel [12, 16, 26], it is not hard to verify that the following assertions are equivalent, see Theorem 3.17 for details:

- There is a tree-decomposition of finite adhesion that (homeomorphically) displays $\Omega(G)$.
- There is a tree-decomposition of finite adhesion with boundary $\Omega(G)$.
- $|G|$ is (completely) metrizable.
- $V(G)$ is a countable union of closed sets in $|G|$.
- G has a normal spanning tree.

Now our second main result provides a local version of the above equivalences, characterising precisely which subsets Ψ of $\Omega(G)$ can be (homeomorphically) displayed. Indeed, a striking, recent result by Carmesin [9] says that it is always possible to display the set of undominated ends of a graph G. In [8], Bürger and Kurkofka partially localized Carmesin's result by constructing tree-decompositions of finite adhesion (with additional desirable properties) that display the boundary ∂U of prescribed infinite sets of vertices $U \subseteq V(G)$ where none of the ends in ∂U are dominated. Carmesin also asked for a characterisation of those pairs of a graph G and a subset $\Psi \subseteq \Omega(G)$ for which G has a tree-decomposition displaying Ψ [9, p. 549].

This problem has also been reiterated in [7, Problem 3.22]. Theorem 3.2 below answers this question.

Another set of questions in infinite topological graph theory concerns so-called Ψ-graphs $|G|_\Psi$, i.e. subspaces of $|G|$ of the form $|G|_\Psi = G \cup \Psi \subseteq |G|$ for a set of ends $\Psi \subseteq \Omega(G)$. Ψ-graphs have been studied in connection with infinite matroids [5, 6, 19]: For example, the topological circles (copies of the unit circle S^1) in $|G|_\Psi$ form the cycles of an infinite matroid whenever Ψ belongs to the Borel σ-algebra of $\Omega(G)$ [5].

It turns out that the correct generalisation of the 3rd bullet above about metrizability of $|G|$ involves precisely the property of complete metrizability of Ψ-spaces.

Theorem 3.2. *For any connected graph G and a set Ψ of ends of G the following are equivalent:*

1. *There is a tree-decomposition of finite adhesion homeomorphically displaying Ψ.*
2. *There is a tree-decomposition of finite adhesion displaying Ψ.*
3. *There is a tree-decomposition of finite adhesion with boundary Ψ.*
4. *$|G|_\Psi$ is completely metrizable.*
5. *Ψ is G_δ in $|G|$.*

Note that from Theorem 3.2 one easily reobtains the above equivalences in the case $\Psi = \Omega$. Indeed, only item (5) needs to be commented on: For this, note that saying that $\Psi = \Omega$ is G_δ in $|G|$ means $\Psi = \Omega$ is a countable intersection of open sets, which turns out to be equivalent to $V(G)$ being a countable union of closed sets in $|G|$. Also note that Ψ being a G_δ means that Ψ is a fairly simple element of the Borel σ-algebra on $|G|$, and in fact, using Theorem 3.2 it is not hard to establish that $|G|_\Psi$ gives an infinite matroid in the special case from [5] where $\Psi \subseteq |G|$ is G_δ.

Carmesin's result that the undominated ends Ψ of any connected graph can always be displayed now follows easily from Theorem 3.2: Simply note that fixing any vertex v and considering the set $B_n(v)$ of all vertices in G within graph distance at most n from v, the set Ψ is the intersection of the countably many open sets $O_n = |G| \setminus \overline{B_n(v)}$ (for $n \in \mathbb{N}$) and hence G_δ, see Theorem 3.26.

Furthermore, Theorem 3.2 also provides tree-decompositions that (homeomorphically) display the undominated ends in the boundary ∂U of any fixed infinite set of vertices $U \subseteq V(G)$, strengthening the above mentioned result by Bürger and Kurkofka from [8]; see Theorem 3.27.

A number of natural questions remain on the topic which subsets of ends can be distinguished.

Problem 3.3. *Characterise which* $\Psi \subseteq \Omega(G)$ *can be distinguished.*

Given two distinct ends ω_1, ω_2 of a graph G write $n(\omega_1, \omega_2) \in \mathbb{N}$ for the minimal order of a separation in G that is oriented differently by ω_1 and ω_2. We say that a tree-decomposition $\mathcal{T}$ with decomposition tree T *efficiently distinguishes* a set of ends Ψ if $\mathcal{T}$ distinguishes Ψ with the additional property that for each $\psi_1 \neq \psi_2 \in \Psi$ there is an edge e on the path in T between $f_{\mathcal{T}}(\psi_1)$ and $f_{\mathcal{T}}(\psi_2)$ with $|X_e| = n(\omega_1, \omega_2)$.

Problem 3.4. *Characterise which* $\Psi \subseteq \Omega(G)$ *can be efficiently distinguished.*

An end ω of a graph is called *thin* if all families of disjoint ω-rays are finite, and *thick* otherwise. Our next problem extends a problem of Diestel [11], asking for which graphs there is a tree-decomposition of finite adhesion displaying precisely its thin ends. Carmesin [9] constructed a graph for which there is no such tree-decomposition, and we construct a different counterexample in Example 3.32 with help of our characterisation of displayable sets of ends from Theorem 3.2. We propose a different way in which a tree-decomposition of finite adhesion might distinguish the thin ends from the thick ends and ask which other bipartitions of $\Omega(G)$ can be distinguished in the same way:

Problem 3.5. *Characterise for which bipartitions* $\Omega(G) = \Omega_1 \sqcup \Omega_2$ *there is a tree-decompositions* $\mathcal{T}$ *of finite adhesion with* $f_{\mathcal{T}}(\Omega_1) \cap f_{\mathcal{T}}(\Omega_2) = \varnothing$.

We conclude with two problems concerning metrizability in end spaces.

Problem 3.6. *Characterise which subspaces* $\Psi \subseteq \Omega(G)$ *are metrizable or completely metrizable.*

Problem 3.7. *Characterise which spaces* $|G|_\Psi$ *are metrizable.*

For $\Psi = \Omega(G)$, an answer to Problem 3.6 is given in [29].

3.2 Basic definitions

Given a tree-decomposition $\mathcal{T} = (T, \mathcal{V})$ of finite adhesion of G, any end ω of G orients each edge $e = xy$ of T according to whether ω lives in a component of $G[A_x] - X_e$ or $G[A_y] - X_e$. This orientation of T points towards a node of T or to an end of T, and ω *lives* in that part for that node or *corresponds* to that end, respectively.

Let $f_{\mathcal{T}} \colon \Omega(G) \to V(T) \cup \Omega(T)$ be the function mapping every end of G to the node or end of T that it lives in or corresponds to, respectively. We say that $\mathcal{T}$ *distinguishes* the ends of G if $f_{\mathcal{T}}$ is injective, and it *represents* the ends of G if $f_{\mathcal{T}}$ is bijective.

We call $f_{\mathcal{T}}^{-1}[\Omega(T)]$ the *boundary* of $\mathcal{T}$, and $f_{\mathcal{T}}^{-1}[V(T)]$ the *interior* of $\mathcal{T}$. We say that $\mathcal{T}$ *displays* a subset $\Psi \subseteq \Omega(G)$ if Ψ is the boundary of $\mathcal{T}$ and $f_{\mathcal{T}} \restriction \Psi \to \Omega(T)$ is bijective, and it *homeomorphically displays* Ψ if $f_{\mathcal{T}} \restriction \Psi \to \Omega(T)$ is a homeomorphism. We say that $\mathcal{T}$ [*bijectively*] *distributes* a subset $\Xi \subseteq \Omega(G)$ if Ξ is the interior of $\mathcal{T}$ and $f_{\mathcal{T}} \restriction \Xi$ is injective [bijective]. Finally, we say that $\mathcal{T}$ *realises* [*represents*] a partition $\Omega(G) = \Xi \sqcup \Psi$ of the end space of G if $\mathcal{T}$ [bijectively] distributes Ξ and displays Ψ.

We conclude this section with a sufficient condition for tree-decompositions to (homeomorphically) display their boundary. We say a rooted tree-decomposition $(T, \mathcal{V})$ is *upwards connected* if for every edge $e \in E(T)$ with $x < y$ the induced subgraph $H_e := G[A_y \setminus A_x] = G[A_y] - X_e$ (with A_x, A_y and X_e as above) is non-empty and connected (or equivalently, H_e is a component of $G - X_e$).

Lemma 3.8. *Every upwards connected rooted tree-decomposition $\mathcal{T} = (T, \mathcal{V})$ of finite adhesion of a graph G homeomorphically displays its boundary.*

Proof. Let Ψ be the boundary of $\mathcal{T}$. We show that $f := f_{\mathcal{T}} \restriction \Psi : \Psi \to \Omega(T)$ is a homeomorphism.

For the proof that f is injective, let $\psi_1 \neq \psi_2 \in \Psi$ and let R_i be the $f(\psi_i)$-ray in T starting in the root of T for $i = 1, 2$. There is a finite vertex set $S \subseteq V(G)$ such that ψ_1 and ψ_2 live in different components of $G - S$. By (T1) there is a finite subtree T' of T containing the root of T such that $S \subseteq \bigcup_{t \in T'} V_t =: G'$. We denote the unique T'–$(T \setminus T')$ edge in R_i by e_i for $i = 1, 2$. Then ψ_i lives in H_{e_i} (as defined above) which is a component of $G - G'$ since $\mathcal{T}$ is upwards connected. Since ψ_1 and ψ_2 live in different components of $G - S$, they also live in different components of $G - G'$. It follows that $H_{e_1} \neq H_{e_2}$. Therefore $e_1 \neq e_2$, $R_1 \neq R_2$, and thus $f(\psi_1) \neq f(\psi_2)$.

Next, for the proof that f is onto, for each end ω of T we find an end $\psi \in \Psi$ such that $f_{\mathcal{T}}(\psi) = \omega$. Let $R = re_0v_1e_1v_2e_2\ldots$ be the ω-ray in T starting in the root of T. We have

$\bigcap_{i\in\mathbb{N}} H_{e_i} = \varnothing$ because each H_{e_i} contains only vertices from parts V_t such the distance of t to the root of T is greater than i. In particular, for every finite subset S of $V(G)$ there is a minimal integer i such that $S \cap H_{e_i} = \varnothing$. Since H_{e_i} is connected and non-empty, there is a unique component $d(S)$ of $G - S$ with $H_{e_i} \subseteq d(S)$. The function d defines a direction on G because the components $H_{e_0} \supseteq H_{e_1} \supseteq \dots$ are nested and non-empty. By Theorem 2.4, there is an end ψ of G such that $C(S, \psi) = d(S)$ for every finite subset $S \subseteq V(G)$. In particular, we have $d(X_{e_i}) = H_{e_i}$ for the separator X_{e_i} corresponding to e_i, and hence ψ lives in H_{e_i} for all $i \in \mathbb{N}$. Consequently, ψ lies in the boundary of $\mathcal{T}$ and $f(\psi) = \omega$.

We now argue that f is continuous (this part of the argument works for any tree-decomposition displaying Ψ and doesn't yet require upwards connectedness). Indeed, let $\psi \in \Psi$ and $f(\psi) = \omega \in \Omega(T)$. For continuity, consider an arbitrary basic open neighbourhood $\Omega_T(T', \omega)$ of $\omega \in \Omega(T)$. Since T is a tree, there is a unique $C(T', \omega)$–T' edge $e = tt'$. Then $X_e = V_t \cap V_{t'}$ is finite since $\mathcal{T}$ had finite adhesion. Now $C_G(X_e, \psi)$ lies completely on one side of the separation (A_e, B_e), and so all ends in $C_G(X_e, \psi)$ orient e towards ω, showing that $f[\Omega_G(X_e, \psi)] \cap \Psi \subseteq \Omega_T(T', \omega)$ as desired.

Finally, we show that f^{-1} is continuous (this part fails without upwards connectedness, cf. Figure 4).Let $f^{-1}(\omega) = \psi \in \Psi$ as before and consider a basic open neighbourhood $\Omega_G(S, \psi) \cap \Psi$ of $\psi \in \Psi$. Let T' be a finite subtree of T which contains the root of T and such that $S \subseteq \bigcup_{t \in V(T')} V_t$. Since ψ orients e towards ω and H_e is connected, it follows that $H_e = C(X_e, \psi) \subseteq C(S, \psi)$. Thus, all ends that orient e towards ω live in $C(S, \psi)$, giving $f^{-1}[\Omega_T(T', \omega)] \subseteq \Omega(S, \psi) \cap \Psi$ as desired. $\square$

3.3 Tree-decompositions displaying all ends

In this section we answer the question which graphs have a tree-decomposition displaying all ends. It turns out that those are exactly the graphs with a normal spanning tree. A characterisation of those graphs by forbidden minors can be found in [37].

Theorem 3.9. *The following are equivalent for any connected graph G:*

1. *There is an upwards connected tree-decomposition of finite adhesion with connected parts that homeomorphically displays $\Omega(G)$.*
2. *There is a tree-decomposition of finite adhesion displaying $\Omega(G)$.*
3. *There is a tree-decomposition of finite adhesion with boundary $\Omega(G)$.*
4. *$|G|$ is (completely) metrizable.*
5. *$\Omega(G)$ is G_δ in $|G|$.*
6. *G has a normal spanning tree.*

Proof. The equivalence (5) $\Leftrightarrow$ (6) is a well-known result by Jung characterising the existence of normal spanning trees [26]. In Jung's language, a connected graph has a normal spanning tree if and only if $V(G)$ is a countable union of dispersed sets; since dispersed sets are precisely the sets of vertices which are closed in $|G|$, this is equivalent to $V = V(G)$ being F_σ in $|G|$. By Lemma 2.13, this is equivalent to G being F_σ in $|G|$, which by taking complements is the same as $\Omega(G)$ being G_δ in $|G|$.

The equivalence (4) $\Leftrightarrow$ (6) is due to Diestel [16].[2] The implications (1) $\Rightarrow$ (2) and (2) $\Rightarrow$ (3) are trivial.

For (3) $\Rightarrow$ (5) suppose we have a tree-decomposition $(T, \mathcal{V})$ with root r of finite adhesion with boundary $\Omega(G)$. We claim that $G[D_n]$ is closed, where

$$D_n := \bigcup_{t \in T^{\leq n}} V_t.$$

Indeed, for any end ω of G there is a unique ray $R = t_0 t_1 t_2 \dots$ starting at the root $t_0 = r$ corresponding to this end. Then $V_{t_n} \cap V_{t_{n+1}}$ is a finite separator that separates D_n from the tails of all ω-rays. Hence no end lives in the closure of D_n, so $G[D_n]$ is closed. It follows from

[2]For (6) $\Rightarrow$ (4), Diestel only verifies that his metric is topologically compatible; but it not hard to see that his metric is in fact complete. See also Theorem 3.17.

property (T1) and (T2) of a tree-decomposition that $G = \bigcup_{n\in\mathbb{N}} G[D_n]$ is F_σ, so by taking complements in $|G|$, we see that $\Omega(G)$ is G_δ in $|G|$.

Lastly, we show (6) $\Rightarrow$ (1). Something similar has been done in [12]. Assume that G has a normal spanning tree T with root r. For every vertex t of T, we define $V_t := \lceil t \rceil$ and show that $\mathcal{T} := (T, (V_t)_{t\in T})$ is a tree-decomposition of G of finite adhesion that homeomorphically displays all its ends. Since T is normal, the end vertices of any edge vw of G are comparable in the tree order. If say $v < w$, then e belongs to the part V_w per definition, giving (T2). Further, if a vertex lies in two parts V_v and V_w, it lies in $\lceil v \rceil \cap \lceil w \rceil$ and hence in all V_t for vertices t on the unique v–w path in T. Thus we get property (T3), so we have a tree-decomposition. It is clear that all parts are connected and since all parts are finite, also all adhesion sets are finite. Finally, $\mathcal{T}$ is clearly upwards connected. Therefore it follows from Lemma 3.8 that $\mathcal{T}$ homeomorphically displays its boundary, which contains all ends of G since all parts are finite. $\square$

3.4 Envelopes

Let G be a connected graph. An *envelope* for a set of vertices $U \subseteq V(G)$ is a set of vertices $U^* \supseteq U$ of finite adhesion (i.e. such that every component of $G - U^*$ has only finitely many neighbours in U^*) with $\partial U^* = \partial U$. In [30, Theorem 3.2] it is proven that every set of vertices in a connected graph admits a connected envelope.

In the following, however, we need a stronger notion of an envelope that works for a set $X \subseteq V(G) \cup \Omega(G)$ of vertices and ends (and in particular, for a set X consisting of ends only): An *envelope* for such a set $X \subseteq V(G) \cup \Omega(G)$ is a set of vertices $X^* \supseteq X \cap V(G)$ of finite adhesion such that $\partial X^* = \overline{X} \cap \Omega(G)$, where the closure $\overline{X}$ of X is taken in $|G|$.

Theorem 3.10. *Any set consisting of vertices and ends in a graph G admits an envelope.*

Proof. Let $X \subseteq V(G) \cup \Omega(G)$ be a given set of vertices and ends in a graph G, and write $V(X) := X \cap V(G)$. Let $\mathcal{R}$ be an inclusionwise maximal set of pairwise disjoint rays of ends in $\overline{X}$. Put

$$X' := V(X) \cup \bigcup_{R \in \mathcal{R}} V(R)$$

and let $\mathcal{S}$ be the set of all centres of (infinite) stars attached to X'. We will show that

$$X^* := X' \cup \mathcal{S}$$

is an envelope for X. The verification relies on the following two claims:

Claim 3.11. *If S is a finite set of vertices and C is a component of $G - S$ such that $X' \cap C$ is finite, then $X^* \cap C = X' \cap C$.*

Only $X^* \cap C \subseteq X' \cap C$ requires proof. For this consider some $v \in \mathcal{S}$. By definition, v is the centre of an infinite star attached to X'. Since S is finite and X' meets C finitely, it follows that $v \notin C$. Hence, $C \cap S = \varnothing$ and so $X^* \cap C = X' \cap C$ as claimed.

Claim 3.12. *If S is a finite set of vertices and C is a component of $G - S$ such that $\overline{C} \cap X = \varnothing$, then $X^* \cap C$ is finite.*

To see the claim, consider some finite set of vertices S, and assume that C is a component of $G - S$ such that $\overline{C}$ avoids X. First, we show that $\overline{C} \cap \overline{X} = \varnothing$. For this, observe that the set $\overline{C} \cup \mathring{E}_{1/2}(S, C)$ is open and disjoint from X and so it is disjoint from $\overline{X}$. In particular, $\overline{C}$ is disjoint from $\overline{X}$. Hence every ray $R' \in \mathcal{R}$ meets C finitely. Furthermore, every ray from $\mathcal{R}$ which meets C also meets S, and since the rays in $\mathcal{R}$ are pairwise disjoint, at most $|S|$ rays from $\mathcal{R}$ meet C. So $\bigcup_{R\in\mathcal{R}} V(R)$ meets C finitely, and hence so does X'. By Claim 3.11, also $X^* \cap C = X' \cap C$ is finite. This establishes the claim.

To see $\partial X^* = \overline{X} \cap \Omega(G)$, we show both inclusions separately. For $\supseteq$ consider any end $\varepsilon \notin \partial X^*$. Then $C(S,\varepsilon) \cap X^* = \varnothing$ for some finite set of vertices S. Consider a ray R in ε that is completely contained in $C(S,\varepsilon)$. Then R is disjoint from any ray in $\mathcal{R}$. By maximality of $\mathcal{R}$, this means that $\varepsilon \notin \overline{X}$.

For $\subseteq$ consider any end $\varepsilon \notin \overline{X}$. Then there is a finite set of vertices S such that $\hat{C}(S,\varepsilon)$ avoids X. By Claim 3.12, also X^* intersects $C(S,\varepsilon)$ finitely, witnessing $\varepsilon \notin \partial X^*$.

To see that X^* has finite adhesion, suppose for a contradiction that there is a component C of $G - X^*$ with infinite neighbourhood. Then by a routine application of the Star-Comb Lemma 2.2, we either find a star or a comb attached to X^* whose centre v or spine R is contained in C. The ray case results in an immediate contradiction as follows: If ε denotes the end with $R \in \varepsilon$, then the comb attached to X^* with spine R witnesses that $\varepsilon \in \partial X^*$. Since $\partial X^* = \overline{X} \cap \Omega(G)$ by the earlier observation, we get $R \in \varepsilon \in \overline{X}$. But then the existence of R contradicts the maximality of $\mathcal{R}$.

In the star case, note that for all finite sets of vertices S disjoint from v, the component C of $G - S$ containing v meets X^* infinitely. Then C also meets X' infinitely by Claim 3.11. But then it is straightforward to inductively construct a star with centre v attached to X', violating the maximality of $\mathcal{S}$. The completes the proof that X^* is an envelope for X. $\square$

Note that the envelopes constructed in Theorem 3.10 are in general neither connected nor end-faithful. But we can easily obtain both properties with the following construction.

For a given subgraph $H \subseteq G$ of finite adhesion, we define a *torso-extension* $H' \supseteq H$ as follows: First, we make H induced. Then for each component C of $G-H$, let $T_C \subseteq G[C \cup N(C)]$ be a finite tree such that all vertices from $N(C)$ are leaves of T_C. We add all these T_C to H to obtain H'.

Lemma 3.13. *Let G be connected. Whenever $H \subseteq G$ is a subgraph of finite adhesion, then every torso-extension H' is an end-faithful connected subgraph of G of finite adhesion with $\partial H' = \partial H$.*

Proof. Since inside of each component of $G - H$ we only add a finite subgraph to H, also H' has finite adhesion.

By construction, every vertex of $H' \setminus H$ is connected via a finite path in H' to a vertex of H. Hence for connectivity of H' it remains to show that there is a path in H' between every two vertices $v, w \in H$.

Since G is connected, there is a v–w path P in G. We consider P as a sequence of edges between vertices of H and segments inside of components C of $G - H$ together with their end-vertices in $N(C)$. After replacing each of those segments in a component C by a path in T_C between the same end-vertices, we obtain a finite v–w walk P' contained in H'. So H' is connected.

To see that $\partial H' = \partial H$, only $\subseteq$ requires proof. If $\omega \notin \partial H$, then ω lives in a unique component C of $G - H$. Since $H' \cap C$ is finite it follows that ω also lives in a unique component C' of $G - H'$ with $C' \subseteq C$ and hence $\omega \notin \partial H'$ by finite adhesion of H'.

We now argue that H' contains an ω-ray for every end ω in $\partial H' = \partial H$. Suppose without loss of generality that $H \neq \varnothing$ and fix any ω-ray $R = r_0 r_1 r_2 \ldots$ in G with $r_0 \in V(H)$. By finite adhesion of H, the ray R contains infinitely many vertices of H. We will construct a ray $R' \subseteq H'$ that meets R infinitely as follows: If $R \subseteq H'$, there is nothing to do. Otherwise, let r_{n_0} be the first vertex on R outside of H', and consider the component $C_0 \ni r_{n_0}$ of $G - H$. Let r_{k_0} be the last vertex of R in C_0. Replace $r_{n_0 - 1} R r_{k_0 + 1}$ by an r_{n_0-1}–r_{k_0+1} path P_0 in $T_{C_0} \subseteq H'$ and call the resulting ray R_1. Note that $R_1 \cap H \subseteq R \cap H$. Now we iterate the same step for R_1 to find a new ray R_2 and so on. This yields a sequence of rays $R_1, R_2, R_3, \ldots$ with $R_n \cap H \subseteq R \cap H$ and which agree on larger and larger initial segments contained in H'. The union of these segments is a ray $R' \subseteq H'$ with $R' \cap H \subseteq R \cap H$, so R' is an ω-ray in H' as desired.

To see that H' is end-faithful, it remains to show that any two rays R_1 and R_2 in H' that are equivalent in G are also equivalent in H'. By assumption there is a collection $\mathcal{P}$ of infinitely many disjoint R_1–R_2 paths in G. We will find infinitely many such paths in H'. Let P be an R_1–R_2 path in G with endvertices $r_1 \in R_1$ and $r_2 \in R_2$. As in the second paragraph, we find a r_1–r_2 walk P' in H'. Consider the finitely many components of $G - H$ that meet P' and delete from $\mathcal{P}$ all paths that meet one of these components – by finite adhesion, $\mathcal{P}$ remains infinite. So we can find another R_1–R_2 path in H' disjoint to the first one. Iterating this construction, we find infinitely many disjoint R_1–R_2 paths in H', showing that R_1 and R_2 in are also equivalent in H'. □

A result for torsos of parts in tree-sets similar to Lemma 3.13 is proven in [21, Section 2.6].

Corollary 3.14. *Any set consisting of vertices and ends in a connected graph G has a connected, end-faithful envelope.*

Whenever we refer to *the envelope* of X inside a connected graph G, we assume that we fixed one possible end-faithful connected choice and call it $\mathcal{E}_G(X)$.

3.5 From topology to tree-decompositions

In this section we employ the envelope technique in order to construct a tree-decomposition of finite adhesion adapted to some prescribed topological information. Roughly, given an infinite graph G and an increasing sequence of closed subsets $X_0 \subseteq X_1 \subseteq X_2 \subseteq \dots$ in $|G|$ such that $V(G) \subseteq \bigcup_{n\in\mathbb{N}} X_n$, we construct a tree-decomposition $\mathcal{T} = (T, \mathcal{V})$ of finite adhesion such that precisely the ends of X_n live in parts indexed by the first n levels of T, and all other ends get displayed.[3]

However, we also want a device that ensures that all ends of some prescribed subcollection Δ of ends in $\bigcup_{n\in\mathbb{N}} X_n$ live in pairwise distinct parts of $\mathcal{T}$. It turns out that this can be achieved provided that each $\Delta_n := \Delta \cap (X_n \setminus X_{n-1})$ is a discrete set.

Lemma 3.15. *Let G be a graph and $\Xi \subseteq \Omega(G)$. Suppose that there is a sequence $X_0 \subseteq X_1 \subseteq X_2 \subseteq \dots$ of subsets of $V(G) \cup \Xi$ that are closed in $|G|$ with $V(G) \cup \Xi = \bigcup_{n\in\mathbb{N}} X_n$. Denote $\Xi_n := X_n \cap \Omega(G)$ and let Δ_n be a discrete subset of $\Xi_n \setminus \bigcup_{i<n} \Xi_i$ for all $n \in \mathbb{N}$. Then there exists a sequence of induced subgraphs of finite adhesion $G_0 \subseteq G_1 \subseteq G_2 \subseteq \dots$ of G such that the following holds for all $n \in \mathbb{N}$:*

(i) *For every component C of $G - G_n$, the set $(C \cap G_{n+1}) \cup N(C)$ is connected in G;*

(ii) $\partial G_{3n+2} \subseteq \Xi_n \setminus \Delta_n$;

(iii) *for every component C of $G - G_{3n+2}$, there is at most one end from Δ_n contained in ∂C;*

(iv) $X_n \cap V(G) \subseteq V(G_{3n+3})$;

(v) $\partial G_{3n+3} = \Xi_n$.

Proof. Set $G_0 := \varnothing$. We will inductively define subgraphs $G_0, G_1, \dots$ of G all of finite adhesion so that *(i)* – *(v)* are satisfied.

Every step of the construction follows the same general pattern: To construct G_{n+1} from G_n consider the current set $\mathcal{C}_n$ of components of $G - G_n$. For every $D \in \mathcal{C}_n$ we consider the subgraph $\tilde{D} := G[D \cup N(D)]$ of G. Each time we will define a set of vertices $V_D \subseteq V(\tilde{D})$ of finite adhesion in $\tilde{D}$ containing $N(D)$. Then also $G_{n+1} := G_n \cup \bigcup_{D\in\mathcal{C}_n} V_D$ has finite adhesion in G since any component C of $G - G_{n+1}$ is also a component of $\tilde{D} - V_D$ for some $D \in \mathcal{C}_n$. Furthermore, we will make sure that V_D is connected so that *(i)* is satisfied.

[3] In the actual proof, we arrange for technical reasons that the ends of X_n live precisely in parts indexed by the first $3n + 3$ levels of T.

Next, we make two observations concerning the end space of $\tilde{D}$, which both follow from the fact that $N(D)$ is finite: Firstly, we have $\partial D = \partial \tilde{D}$ in G, and secondly, the inclusion map ι as mentioned in Section 2.2 is a homeomorphism from $\Omega(\tilde{D})$ to $\partial D \subseteq |G|$. Via this homeomorphism, we will in the following identify the spaces $\Omega(\tilde{D})$ and $\partial D \subseteq |G|$.

Now for the actual construction of the sequence $G_0, G_1, \dots,$ we proceed in steps of three. Suppose that G_{3n} has already been defined. We demonstrate how to recursively construct

$$G_{3n} \rightsquigarrow G_{3n+1} \rightsquigarrow G_{3n+2} \rightsquigarrow G_{3n+3} = G_{3(n+1)}$$

in order to satisfy $(i) - (v)$ for the three indices $3n+1$, $3n+2$ and $3n+3$.

1. Step $3n \rightsquigarrow 3n+1$.
Let D be any component from $\mathcal{C}_{3n}$. Since Δ_n is discrete in $|G|$, also $\Delta_n \cap \partial D$ is discrete in ∂D. Thus there is a set $\mathcal{O}_D = \{O_\omega : \omega \in \Delta_n \cap \partial D\}$ of open subsets of ∂D with $O_\omega \cap \Delta_n = \{\omega\}$ for all $\omega \in \Delta_n \cap \partial D$. Applying Corollary 3.14 we consider the envelope

$$V_D := \mathcal{E}_{\tilde{D}}(((\Xi_n \cap \partial D) \setminus \bigcup \mathcal{O}_D) \cup N(D))$$

which is a connected vertex set of finite adhesion in $\tilde{D}$ (cf. Figure 5).

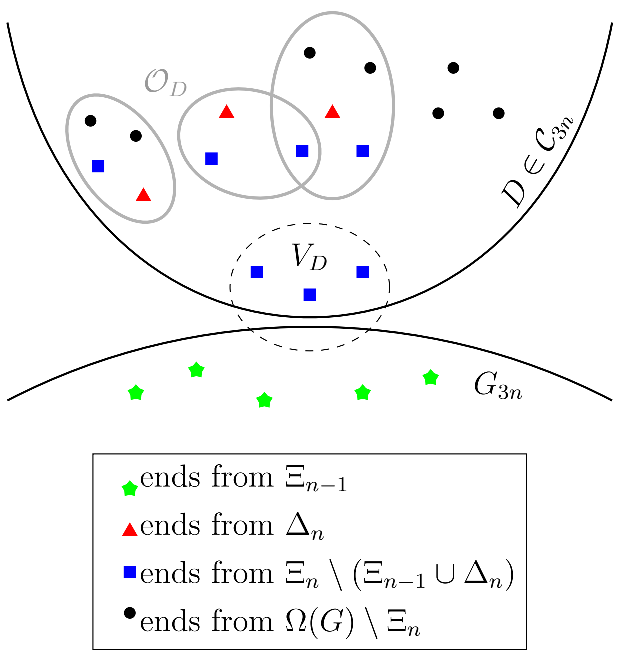

Figure 5: Construction step $3n \rightsquigarrow 3n+1$.

We now determine which ends are contained in ∂V_D. Since both X_n and $\Omega(G)$ are closed in $|G|$, also $\Xi_n = X_n \cap \Omega(G)$ is closed in $|G|$. Hence $(\Xi_n \cap \partial D) \setminus \bigcup \mathcal{O}_D$ is closed in the subspace

∂D of $|G|$. Since $N(D)$ is finite and therefore does not have any ends in its closure, it follows from the definition of an envelope and our identification of $\Omega(\tilde{D})$ with ∂D that

$$\partial V_D = \overline{(((\Xi_n \cap \partial D) \setminus \bigcup \mathcal{O}_D) \cup N(D))} \cap \partial D = \overline{(\Xi_n \cap \partial D) \setminus \bigcup \mathcal{O}_D} = (\Xi_n \cap \partial D) \setminus \bigcup \mathcal{O}_D.$$

Hence by (v) for G_{3n}, the graph $G_{3n+1} = G_{3n} \cup \bigcup_{D \in \mathcal{C}_{3n}} V_D$ satisfies

(vi) $\partial G_{3n+1} \subseteq \Xi_n \setminus \Delta_n$.

Next, we show that

(vii) for every component C of $G - G_{3n+1}$, there is an open cover $\mathcal{O}$ of ∂C such that each set from $\mathcal{O}$ contains at most one end from Δ_n.

Let C be a component of $G - G_{3n+1}$ and D' the component of $G - G_{3n}$ with $C \subseteq D'$. We show that *(vii)* is fulfilled with

$$\mathcal{O} := \{O \cap \partial C : O \in \mathcal{O}_{D'}\} \cup \{\partial C \setminus \Xi_n\}.$$

Clearly, all sets in $\mathcal{O}$ are open in ∂C and contain at most one end from Δ_n. For the proof that $\partial C \subseteq \bigcup \mathcal{O}$, we observe that C and $V_{D'}$ are disjoint and the neighbourhood of C is finite. Therefore, ∂C and $\partial V_{D'}$ are disjoint. Since $\partial V_{D'} = (\Xi_n \cap \partial D') \setminus \bigcup \mathcal{O}_{D'}$, we have $\partial C \cap \Xi_n \subseteq \bigcup \mathcal{O}_{D'}$ and therefore $\partial C \subseteq \bigcup \mathcal{O}$.

2. Step $3n+1 \rightsquigarrow 3n+2$.
Let D be any component from $\mathcal{C}_{3n+1}$. By *(vii)* there exists an open cover $\mathcal{O}$ of ∂D such that each set from $\mathcal{O}$ contains at most one end from Δ_n (cf. Figure 6). Then by Theorem 2.11, there is a rayless normal tree T in $\tilde{D}$ such that for every component C' of $\tilde{D} - T$ there is a set $O \in \mathcal{O}$ with $\partial C' \subseteq O$. By Lemma 2.12 there exists a rayless normal tree T^* in $\tilde{D}$ such that $V(T) \cup N(D) \subseteq V(T^*)$. We define $V_D := V(T^*)$. Then every component C of $G - V_D$ is contained in a component C' of $G - T$ and thus there is a set $O \in \mathcal{O}$ with $\partial C \subseteq O$. Then by *(vii)*, C contains at most one end from Δ_n. Hence $G_{3n+2} = G_{3n+1} \cup \bigcup_{D \in \mathcal{C}_{3n+1}} V_D$ satisfies *(iii)*. Furthermore, T^* has finite adhesion in $\tilde{D}$ by Lemma 2.12. Finally, normal trees are end-faithful by [14, Lemma 8.2.3], so from the fact that T^* is rayless it follows that $\partial T^* = \varnothing$. Therefore $\partial G_{3n+2} = \partial G_{3n+1}$ and *(ii)* is a consequence of *(vi)*.

3. Step $3n+2 \rightsquigarrow 3n+3$.
Again let D be any component from $\mathcal{C}_{3n+2}$ (the components of $G - G_{3n+2}$). We define

$$V_D := \mathcal{E}_{\tilde{D}}((X_n \cap \overline{D}) \cup N(D)).$$

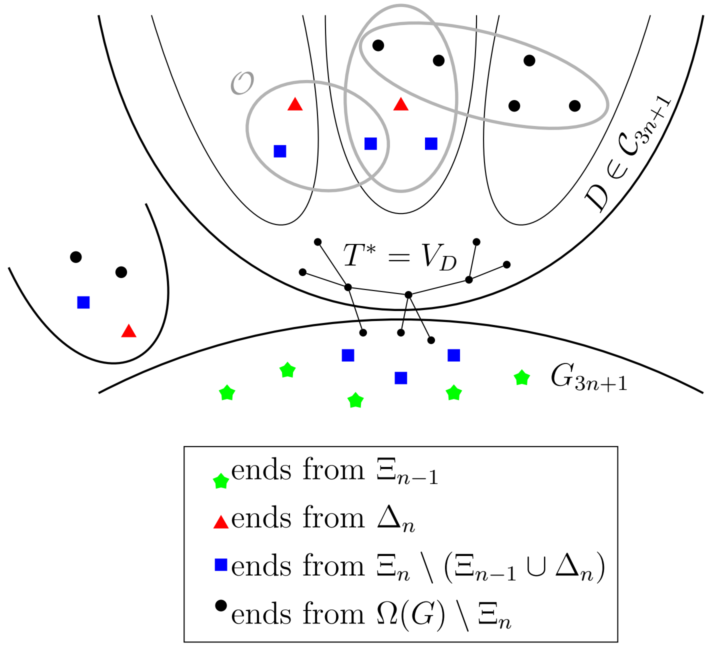

Figure 6: Construction step $3n+1 \rightsquigarrow 3n+2$.

Then it follows from the definition of an envelope that

$$X_n \cap V(\tilde{D}) \subseteq ((X_n \cap \overline{D}) \cup N(D)) \cap V(\tilde{D}) \subseteq V_D.$$

Therefore $G_{3n+3} = G_{3n+2} \cup \bigcup_{D \in \mathcal{C}_{3n+2}} V_D$ satisfies *(iv)*. Furthermore, since $N(D)$ is finite and X_n is closed, we have

$$\partial V_D = \overline{((X_n \cap \overline{D}) \cup N(D))} \cap \partial D = X_n \cap \partial D = \Xi_n \cap \partial D.$$

Then together with *(ii)* we obtain $\partial G_{3n+3} = \Xi_n$ which proves *(v)*. □

Theorem 3.16. *Let G be a connected graph and $\Xi \subseteq \Omega(G)$. Suppose that there is a sequence $X_0 \subseteq X_1 \subseteq X_2 \subseteq \dots$ of subsets of $V(G) \cup \Xi$ that are closed in $|G|$ with $V(G) \cup \Xi = \bigcup_{n \in \mathbb{N}} X_n$. Denote $\Xi_n := X_n \cap \Omega(G)$ and let Δ_n be a discrete subset of $\Xi_n \setminus \bigcup_{i<n} \Xi_i$ for all $n \in \mathbb{N}$. Then there is an upwards connected tree-decomposition $\mathcal{T} = (T, \mathcal{V})$ of finite adhesion with connected parts which homeomorphically displays $\Omega(G) \setminus \Xi$ such that the boundary of every part contains at most one end from $\bigcup_{n \in \mathbb{N}} \Delta_n$.*

Proof. Let $G_0 \subseteq G_1 \subseteq G_2 \subseteq \dots$ be the sequence from Lemma 3.15 with properties *(i)* – *(v)* and suppose without loss of generality that $G_0 = \varnothing$. This sequence gives rise to a tree-decomposition $\mathcal{T} = (T, \mathcal{V})$ of finite adhesion and into connected parts as follows: Write $\mathcal{C}_n$ for the set of components of $G - G_n$. We define a tree order $\leq_T$ on $T := \bigsqcup_{n\in\mathbb{N}} \mathcal{C}_n$ as follows: For all $C_n \in \mathcal{C}_n$ and $C_m \in \mathcal{C}_m$, let $C_n \leq_T C_m$ if and only if $C_n \supseteq C_m$ and $n \leq m$; this will be our decomposition tree. Note that $G_0 = \varnothing$ ensures T has a root whose associated part is G. The part corresponding to a node $C \in \mathcal{C}_n$ of T will be $N(C) \cup (C \cap G_{n+1})$ (which is precisely the set V_C from the proof of Lemma 3.15). Then it is readily checked that all properties (T1) – (T3) of a tree-decomposition are satisfied, in particular (T1) holds by *(iv)*. All parts of $\mathcal{T}$ are connected by *(i)*.

It is clear from the construction that $\mathcal{T}$ is upwards connected. Furthermore, by *(v)* the interior of $\mathcal{T}$ is Ξ and hence its boundary is $\Omega(G) \setminus \Xi$. Therefore $\mathcal{T}$ homeomorphically displays $\Omega(G) \setminus \Xi$ by Lemma 3.8.

It is left to show that in every part of $\mathcal{T}$ there lives at most one end from $\bigcup_{n\in\mathbb{N}} \Delta_n$. For any $n \in \mathbb{N}$, we have $\Delta_n \subseteq \partial G_{3n+3} \setminus \partial G_{3n+2}$ by *(ii)* and *(v)*.

Since this inclusion holds for all $n \in \mathbb{N}$, it follows that $\partial G_{3n+3} \setminus \partial G_{3n+2}$ does not contain ends from $\Delta_{n'}$ for any $n' \neq n$. Furthermore, by *(iii)* every component in $\mathcal{C}_{3n+2}$ contains at most one end from Δ_n in its boundary. Hence all ends from Δ_n are contained in the boundaries of parts of the form $N(C) \cup (C \cap G_{3n+3})$ for $C \in \mathcal{C}_{3n+2}$, and in the boundary of every such part there is no end from $\Delta_{n'}$ for any $n' \neq n$ and at most one end from Δ_n. This finishes the proof. $\square$

3.6 Tree-decompositions displaying sets of ends

In this section we will prove our characterisation announced in Theorem 3.2 of *displayable* subsets of $\Omega(G)$, i.e. subsets which can be (homeomorphically) displayed by a tree-decomposition of finite adhesion.

Theorem 3.17. *For any connected graph G and any set Ψ of ends of G the following are equivalent:*

1. *There is an upwards connected tree-decomposition of finite adhesion with connected parts that homeomorphically displays Ψ.*
2. *There is a tree-decomposition of finite adhesion displaying Ψ.*
3. *There is a tree-decomposition of finite adhesion with boundary Ψ.*
4. *$|G|_\Psi$ is completely metrizable.*
5. *Ψ is G_δ in $|G|$.*

Proof. We demonstrate the following sequence of implications:

$$\begin{tikzcd}
 & (2) \arrow[r, Rightarrow] & (3) \arrow[dr, Rightarrow] & & \\
(1) \arrow[ur, Rightarrow] \arrow[dr, Rightarrow] & & & (5) \arrow[r, Rightarrow] & (1) \\
 & (4) \arrow[urr, Rightarrow] & & &
\end{tikzcd}$$

The implications (1) $\Rightarrow$ (2) $\Rightarrow$ (3) are trivial.

(1) $\Rightarrow$ (4): Let $(T, \mathcal{V})$ be a tree-decomposition of finite adhesion of G homeomorphically displaying Ψ with a fixed root r of T. We begin by defining a complete metric d_T on $V(T) \cup \Omega(T)$. Assign to every $e \in E(T)$ a number $\ell(e)$: If $e \in E(T)$ is a T^n–T^{n+1} edge (i.e. an edge between level n and level $n+1$ of T), we set $\ell(e) = 1/2^n$. If P is a (possibly infinite) path in T, we say that the finite number $\sum_{e \in E(P)} \ell(e)$ is the *length* of P. Now we define $d_T(x, y)$ for all $x, y \in V(T) \cup \Omega(T)$: If x and y are both vertices, let $d_T(x, y)$ be the length of the unique x–y path in T. If x is a vertex and y is an end, then let $d_T(x, y)$ be the length of the unique ray from y which starts in x. Similarly, if both x and y are ends, let $d_T(x, y)$ be the length of the unique double ray in T between x and y. It is straight-forward to check that d_T defines a complete metric on $V(T) \cup \Omega(T)$.

We now use d_T to define a metric d on $G \cup \Psi$. For every vertex $v \in V(G)$, let v_T be the least vertex of T with respect to the tree order such that v is contained in the part V_{v_T} (this is well-defined according to (T3)). Additionally, for every end $\omega \in \Psi$, let ω_T be the end of T which ω corresponds to. For all $x, y \in V \cup \Psi$, we define

$$d(x,y) = \begin{cases} 0 & \text{if } x = y, \\ 1/2^n & \text{if } x \neq y \in V(G) \text{ and } x_T = y_T \text{ lies in the } n\text{th level of } T, \\ d_T(x_T, y_T) & \text{if } x_T \neq y_T. \end{cases}$$

Next, we prove that d is a metric on $V(G) \cup \Psi$. It is clear that $d(x,x) = 0$ and $d(x,y) > 0$ for all $x \neq y$ and that d is symmetric. We show that triangle inequality holds: Let x, y, z be pairwise distinct elements of $V(G) \cup \Psi$. We need to show that

$$d(x,z) \leq d(x,y) + d(y,z). \tag{$*$}$$

Clearly, $(*)$ holds if $x_T = y_T = z_T$. If $x_T = y_T \neq z_T$, then $d(x,z) = d(y,z)$ and hence $(*)$ follows. A similar argument works if $y_T = z_T$. Next, suppose that $x_T = z_T \neq y_T$ and let n be the level of x_T in T. Then $d(x,z) = 1/2^n$ and since $\ell(e) \geq 1/2^n$ for every edge e of T with endvertex x_T also $d(x,y) \geq 1/2^n$, which proves $(*)$. Finally, if x_T, y_T and z_T are pairwise distinct, then $(*)$ follows from the triangle inequality for d_T. This finishes the proof of $(*)$.

For the proof that d is complete, let $(x_n)_{n\in\mathbb{N}}$ be a Cauchy-sequence in $V(G) \cup \Psi$. Hence $((x_n)_T)_{n\in\mathbb{N}}$ is a Cauchy-sequence in T because $d_T(v_T, w_T) \leq d(v,w)$ for all $v, w \in V(G) \cup \Psi$. If $((x_n)_T)_{n\in\mathbb{N}}$ is eventually constant, then $(x_n)_{n\in\mathbb{N}}$ is eventually contained in V_t for some $t \in V(T)$. If t lies in the nth level of T, then $d(v,w) \geq 1/2^n$ for all $v \neq w \in V_t$. Hence also $(x_n)_{n\in\mathbb{N}}$ is eventually constant. Otherwise, if $((x_n)_T)_{n\in\mathbb{N}}$ is not eventually constant, then $((x_n)_T)_{n\in\mathbb{N}}$ converges to an end ω of T and thus $(x_n)_{n\in\mathbb{N}}$ converges to the end of G which corresponds to ω. Finally, we extend d to a complete metric on $G \cup \Psi$ by relating every edge vw of G linearly to a real closed interval of length $d(v,w)$. We omit the details.

It is left to show that the metric d induces the subspace topology on $G \cup \Psi$ inherited from $|G|$. We need to show for any given $x \in G \cup \Psi$ that

(†) every MTop-basic open neighbourhood of x in $G \cup \Psi$ contains an open ε-ball around x with respect to d, and vice versa.

This is clear if x is an inner point of an edge. Next, let $x \in V(G)$ be a vertex and n the level of x_T in T. Then (†) is true because every edge of G which has x as an endvertex has length at least $1/2^n$ and at most 1.

Now suppose that $x \in \Psi$ and let $\hat{C}_\varepsilon(S, x)$ be a basic open neighbourhood of x in $|G|$ for some $\varepsilon \leq 1$. Let n be the maximum level of T containing a vertex s_T for some $s \in S$. We show that the open ball B in $|G|$ with respect to d with radius $\varepsilon/2^n$ and centre x is a subset of $\hat{C}_\varepsilon(S,x)$. First, consider the open ball B' in T with respect to the metric d_T with radius $\varepsilon/2^n$ and centre x', where x' is the end of T which x corresponds to. Let e be the edge of T which is contained in the normal x'-ray in T and connects a node u_n form the nth level of T to a node u_{n+1} from the $n+1$st level. Then B' is completely contained in the closure of the component D of $T-e$ with $u_{n+1} \in V(D)$ since

$$d_T(u_{n+1}, x') = \sum_{i \geq n+1} 1/2^i = 1/2^n \geq \varepsilon/2^n.$$

In particular, every vertex in B' lies in the $n+1$st level of T or above. Next, it follows from the definition of the metric d that every vertex in B is contained in a part V_t with $t \in B' \subseteq \overline{D}$, but no vertex of B can be contained in a part V_t such that the level of t in T is at most n. Therefore all vertices in B and similarly also all ends in B are contained in $\overline{H_e}$, where H_e is the subgraph of G from the definition of upwards connectedness. Since H_e is disjoint from S, connected by upwards connectedness of $\mathcal{T}$, and x orients e towards x', we have $\overline{H_e} \subseteq \hat{C}(S,\omega)$. Hence all vertices and ends in B and all edges with both endvertices in B are contained in $\hat{C}_\varepsilon(S,x)$; it is left to show the same for points of edges in B with only one endvertex in B. Every such edge f, however, has its other endvertex in V_{u_n} by (T3), and as u_n lies in the nth level of T, the length of f with respect to d is at least $1/2^n$. Recall that any point p on f in B has distance less than $\varepsilon/2^n$ to x and therefore also to the end vertex of f in B. Thus p is contained in $\hat{C}_\varepsilon(S,x)$, as desired.

Conversely, let B be an open ε-ball around x with respect to d of radius $0 < \varepsilon \leq 1$. Let $\omega \in \Omega(T)$ be the end of T corresponding to x and R the rooted ω-ray in T. Choose $n \in \mathbb{N}$ such that $1/2^n < \varepsilon$ and let $t^i \in V(T)$ be the node in $R \cap T^i$ for $i \in \{n+2, n+3\}$. Then define S as the separator induced by the edge $t^{n+2}t^{n+3}$ of T in G. Now $C := \hat{C}_{1/2^{n+1}}(S,x)$ is a subset of B: Let y be any point in C; we have to show that $d(y,x) < \varepsilon$. First suppose that $y \in C(S,x)$ and let w be a vertex from the part $V_{t^{n+3}}$. For any point $z \in C(S,x)$ we have

$$d(w,z) \leq \sum_{i \geq n+3} 1/2^i = 1/2^{n+2}.$$

Hence

$$d(y,x) \leq d(y,w) + d(w,x) \leq 1/2^{n+2} + 1/2^{n+2} = 1/2^{n+1} < \varepsilon.$$

Next, suppose that y is an inner point of an S–$C(S,x)$ edge with endvertex v in $C(S,x)$. We have seen above that $d(v,x) \leq 1/2^{n+1}$. Hence it follows from the choice of C that

$$d(y,x) \leq d(y,v) + d(v,x) \leq 1/2^{n+1} + 1/2^{n+1} < \varepsilon$$

which proves $C \subseteq B$.

(4) $\Rightarrow$ (5): Assume that $|G|_\Psi$ is completely metrizable. We claim that

- Ψ is G_δ in $|G|_\Psi$, and
- $|G|_\Psi$ is G_δ in $|G|$.

This implies (5) as being G_δ is transitive.

Since closed subsets of metrizable spaces are always G_δ [22, Corollary 4.1.12], we get that Ψ is G_δ in $|G|_\Psi$. Next, by a well-known result of Čech [22, Theorem 4.3.26] all completely metrizable spaces, and so in particular $|G|_\Psi$, are Čech-complete, and by [22, Exercise 3.9.A], all Čech-complete spaces are G_δ in their closures. Thus we conclude that $|G|_\Psi$ is G_δ in its closure $|G|$.

(3) $\Rightarrow$ (5): Let $(T, \mathcal{V})$ be a tree-decomposition of finite adhesion of G with boundary Ψ. Fix a root r of T and denote by E_n the set of all edges between the nth and $n+1$st level of T. For every edge $e \in E_n$, let (A_e, B_e) be the respective separation of G such that $V_r \subseteq A_e$ and let $S_e = A_e \cap B_e$ be the corresponding finite adhesion set. Note that A_e contains every part V_t with $t \in T^{\leq n}$. We denote

$$\mathcal{C}_e := \bigcup \{\hat{C}_{1/2}(S_e, \omega) : \omega \in \partial B_e\}.$$

Then $O_n := \bigcup_{e \in E_n} \mathcal{C}_e$ is an open set in $|G|$ because it is a union of open sets. We show that $\Psi = \bigcap_{n \in \mathbb{N}} O_n$. Clearly, $\Psi \subseteq \bigcap_{n \in \mathbb{N}} O_n$. For the converse inclusion, let $\omega \in \bigcap_{n \in \mathbb{N}} O_n$. We show that ω does not live in any part of $(T, \mathcal{V})$ and therefore lies in the boundary of $(T, \mathcal{V})$. Indeed, if $\omega \in \partial V_t$ for $t \in T^n$, then ω is not contained in O_{n+1}, a contradiction.

(5) $\Rightarrow$ (1): Let $\Psi \subseteq \Omega(G)$ be a G_δ set in $|G|$. Hence $G \cup \Xi$ where $\Xi := \Omega(G) \setminus \Psi$ is an F_σ set in $|G|$ and by Lemma 2.13, also $V(G) \cup \Xi$ is an F_σ set in $|G|$. This means that $V(G) \cup \Xi = \bigcup_{n \in \mathbb{N}} X_n$ is a countable union of sets X_n which are closed in $|G|$, we may assume that $X_0 \subseteq X_1 \subseteq \cdots$. By applying Theorem 3.16 (with $\Delta_n = \varnothing$) there is an upwards connected tree-decomposition of finite adhesion into connected parts that homeomorphically displays $\Psi = \Omega(G) \setminus \bigcup_{n \in \mathbb{N}} X_n$. $\square$

Corollary 3.18. *Displayable sets of ends are completely metrizable.*

Proof. The implication (2) $\Rightarrow$ (4) in Theorem 3.17 says that for every displayable set of ends $\Psi \subseteq \Omega(G)$ in a graph G we have that $|G|_\Psi$ is completely metrizable. Since $\Psi \subseteq |G|_\Psi$ is closed, and closed subspaces of completely metrizable spaces are again completely metrizable, it follows that Ψ is completely metrizable. $\square$

Corollary 3.19. *Let G be a graph with a displayable set of ends $\Psi \subseteq \Omega(G)$ and let Φ be a subset of Ψ. Then Φ is (homeomorphically) displayable if and only if Φ is a G_δ set in Ψ.*

Proof. Immediate from (2) $\Leftrightarrow$ (5) in Theorem 3.17 and transitivity of the G_δ-property. $\square$

Corollary 3.20. *Let G be a graph with a normal spanning tree. Then a subset $\Phi \subseteq \Omega(G)$ is (homeomorphically) displayable if and only if Φ is a G_δ set in $\Omega(G)$.*

Proof. Follows from (6) $\Rightarrow$ (5) in Theorem 3.9 together with the previous corollary for $\Psi = \Omega(G)$. $\square$

3.7 Tree-decompositions distributing sets of ends

In this section we characterise which subsets of ends can be distributed by a tree-decomposition of finite adhesion. Recall that a topological space $X \subseteq Z$ has a *σ-discrete expansion in Z* if it can be written as a disjoint union $X = \bigsqcup_{n\in\mathbb{N}} X_n$ such that all X_n are discrete and all $Y_n := \bigcup_{i\leq n} X_i$ are closed in Z.

Theorem 3.21. *Let G be a connected graph and $\Xi \subseteq \Omega(G)$ a subset of ends of G. Then the following are equivalent:*

(i) There is a tree-decomposition of finite adhesion distributing Ξ.

(ii) $V(G)$ is a countable union of slender vertex sets U_n such that $\bigcup_{n\in\mathbb{N}} \partial U_n = \Xi$.

(iii) $V(G) \cup \Xi$ has a σ-discrete expansion in $|G|$.

(iv) There is an upwards connected tree-decomposition of finite adhesion with connected parts realising $(\Xi, \Xi^\complement)$.

Proof. We will show a cyclic chain of implications. For $(i) \Rightarrow (ii)$, suppose we have a tree-decomposition $(T, \mathcal{V})$ with root r of finite adhesion that distributes Ξ.

We define

$$U_n = \bigcup_{t\in T^{\leq n}} V_t.$$

By property (T1) of a tree-decomposition, it is clear that $V(G) \subseteq \bigcup_{n\in\mathbb{N}} U_n$. Since Ξ is the interior of $(T, \mathcal{V})$, we also have $\Xi = \bigcup_{n\in\mathbb{N}} \partial U_n$ as desired.

Furthermore, each U_n is slender: Clearly, all vertices are isolated in $|G|$. Additionally, $\partial U_n \setminus \partial U_{n-1}$ consists of at most one end for each part V_t for $t \in T^n$ and hence all ends in $\partial U_n \setminus \partial U_{n-1}$ are isolated points of U_n. Therefore, each $\overline{U_n}$ has Cantor-Bendixson rank at most $n + 1$ by induction.

For $(ii) \Rightarrow (iii)$, suppose $V(G)$ is a countable union of slender vertex sets U_n such that $\bigcup_{n\in\mathbb{N}} \partial U_n = \Xi$. Without loss of generality, the sequence of the U_n is increasing. Write $X_n = \overline{U_n}$ and let $Y_0 = X_0$ and $Y_{n+1} = X_{n+1} \setminus X_n$. By assumption, each Y_n has finite Cantor-Bendixson rank say k_n. Recall that $Y_n^{(0)} := Y_n$ and $Y_n^{(i+1)}$ denotes the derived space of $Y_n^{(i)}$ for all $i \in \mathbb{N}$. Since Y_n has rank k_n, we have $Y_n^{(k_n)} = \varnothing$. Let $Z_{n,i} := Y_n^{(i)} \setminus Y_0^{(i+1)}$ be the subset of Y_n consisting of all elements that get deleted when forming $Y_n^{(i+1)}$ for $0 \leq i \leq k_n - 1$. We claim that

$$Z_{0,k_0-1}, Z_{0,k_0-2}, \dots, Z_{0,0}, Z_{1,k_1-1}, Z_{1,k_1-2}, \dots, Z_{1,0}, Z_{2,k_2-1}, Z_{2,k_2-2}, \dots$$

is the desired σ-discrete expansion of $V(G) \cup \Xi$.

First of all, since $V(G) \cup \Xi = \bigcup_{n\in\mathbb{N}} Y_n$ and this union is disjoint, the above sequence has union $V(G) \cup \Xi$. By the definition of rank, it is also clear that all sets in the sequence are discrete. It remains to show that the union over finite initial segments is closed. Clearly, each such union is of the form

$$Y = X_n \cup Z_{n+1,k_{n+1}-1} \cup \cdots \cup Z_{n+1,i} \subseteq X_{n+1}$$

for some $i < k_{n+1}$, and this set is closed in $|G|$ as X_{n+1} is closed in $|G|$ and Y is closed in X_{n+1} by the definition of the Cantor-Bendixson rank.

For $(iii) \Rightarrow (iv)$, let $(X'_n)_{n\in\mathbb{N}}$ be a σ-discrete expansion for $V(G) \cup \Xi$. Then we apply Theorem 3.16 for the closed sets $X_n := \bigcup_{i\le n} X'_i$ and the discrete sets $\Delta_n := X'_n \cap \Omega(G)$ to obtain an upwards connected tree-decomposition of G of finite adhesion into connected parts displaying $\Xi^{\complement}$ such that all ends from $\Xi = \bigcup_{n\in\mathbb{N}} \Delta_n$, and hence all ends from the interior of $\mathcal{T}$ live in pairwise distinct parts. In other words, this tree-decomposition realises $(\Xi^{\complement}, \Xi)$.

Next, it is clear that (iv) implies (i), which completes the proof. □

We have now all results in place to prove our main result Theorem 3.1 from this paper, the following theorem contains even more equivalent properties:

Theorem 3.22. *The following are equivalent for any connected graph G with at least one end:*

1. *There is an upwards connected tree-decomposition of finite adhesion that represents $\Omega(G)$ such that all parts induce connected subgraphs.*
2. *There is a tree-decomposition of finite adhesion that represents all ends in $\Omega(G)$.*
3. *There is a tree-decomposition of finite adhesion that distinguishes all ends in $\Omega(G)$.*
4. *There is a tree-decomposition of finite adhesion into ≤ 1-ended parts.*
5. *Some subset $\Xi \subseteq \Omega(G)$ of ends can be distributed.*
6. *$V(G)$ is a countable union of slender sets.*

Proof. The implications $(1) \Rightarrow (2) \Rightarrow (3) \Rightarrow (4) \Rightarrow (5)$ are trivial. The implication $(5) \Rightarrow (6)$ follows from $(i) \Rightarrow (ii)$ in Theorem 3.21. Finally, for $(6) \Rightarrow (1)$ note that due to $(ii) \Rightarrow (iv)$ in Theorem 3.21, we immediately get from (6) that there is an upwards connected tree-decomposition of finite adhesion into connected parts that realises $(\Xi, \Xi^{\complement})$. But then it follows from the subsequent Lemma 3.23 that there also is such a tree-decomposition $\mathcal{T}'$ that represents some partition (Ξ', Ψ') of $\Omega(G)$ with $\Xi \subseteq \Xi'$, and so $\mathcal{T}'$ represents all ends in $\Omega(G)$ as desired. $\square$

Lemma 3.23. *If a connected graph G with at least one end admits a tree-decomposition $\mathcal{T}$ of finite adhesion that realises some partition (Ξ, Ψ) of $\Omega(G)$, then there also is such a tree-decomposition $\mathcal{T}'$ that represents some partition (Ξ', Ψ') of $\Omega(G)$ with $\Xi \subseteq \Xi'$.*

Moreover, whenever $\mathcal{T}$ has connected parts or is upwards connected, we can obtain the same for $\mathcal{T}'$.

Proof. Suppose we are given a tree-decomposition $(T, \mathcal{V})$ of finite adhesion realising some partition (Ξ, Ψ) of $\Omega(G)$. We will perform two rounds of contractions on T to make sure that we represent some partition (Ξ', Ψ') of $\Omega(G)$ with $\Xi \subseteq \Xi'$.

First, pick a maximal family $\mathcal{R}$ of disjoint rays in T such that no end of G lives in a part corresponding to one of the nodes of a ray in $\mathcal{R}$. Then consider a new tree-decomposition $(\dot{T}, \dot{\mathcal{V}})$ where $\dot{T}$ is obtained from T by contracting each ray in $\mathcal{R}$. For every $R \in \mathcal{R}$ we define a corresponding part $\dot{V}_R = \bigcup_{t \in R} V_t$. Since the set of separators of $(\dot{T}, \dot{\mathcal{V}})$ is a subset of the set of separators of $(T, \mathcal{V})$, it follows that also $(\dot{T}, \dot{\mathcal{V}})$ has finite adhesion. And since by assumption on $(T, \mathcal{V})$ there corresponds precisely one end of G to any ray $R \in \mathcal{R}$, it follows that $(\dot{T}, \dot{\mathcal{V}})$ realises (Ξ', Ψ') where Ξ' is the union of Ξ together with all ends of G that correspond to a ray in $\mathcal{R}$, and Ψ' is its complement.

Next, note that by maximality of $\mathcal{R}$, every ray of $\dot{T}$ contains infinitely many nodes whose corresponding parts in $\dot{\mathcal{V}}$ contain an end of G. Therefore, if we pick any partition $\mathcal{P}$ of $V(\dot{T})$ into subtrees such that each subtree P contains a unique node for which there is an end ω_P of G living in the corresponding part of $\dot{\mathcal{V}}$, then all $P \in \mathcal{P}$ are necessarily rayless.

Now consider a new tree-decomposition $(T', \mathcal{V}')$ where T' is obtained from $\dot{T}$ by contracting each subtree in $\mathcal{P}$. Naturally, $V(T') = \mathcal{P}$, and for each $P \in V(T')$ we define $V'_P = \bigcup_{t \in P} \dot{V}_t$. Since $\mathcal{T}'$ arises from $\mathcal{T}$ by contracting subtrees, it is clear that $\mathcal{T}'$ has finite adhesion, connected parts, or is upwards connected if the same is true for $\mathcal{T}$. Lastly, $(T', \mathcal{V}')$ now represents the partition (Ξ', Ψ'), as in each part V'_P there lives precisely the single end ω_P from Ξ', and since all P were rayless and $(\dot{T}, \dot{\mathcal{V}})$ displays Ψ', also $(T', \mathcal{V}')$ displays Ψ'. $\square$

Corollary 3.24. *If a connected graph G with at least one end admits a rayless tree-decomposition $\mathcal{T}$ of finite adhesion that distributes $\Omega(G)$, then there also is such a tree-decomposition that bijectively distributes $\Omega(G)$. Moreover, whenever $\mathcal{T}$ has connected parts or is upwards connected, we can obtain the same for $\mathcal{T}'$.*

3.8 Tree-decompositions distributing all ends

In the previous section we stated a topological characterisation for the sets of ends that can be distributed. If we are interested in distributing all ends of G, we can obtain a combinatorial characterisation in terms of the underlying graph.

The following is a convenient description of the Cantor-Bendixson rank of the space $V \cup \Omega(G) \subseteq |G|$ due to Jung [26, §3]: The *rank* $r(x)$ of a vertex or an end x in a graph $G = (V, E)$ is defined as follows: all vertices have rank 0. An end ω has rank 1, if there is a finite set $S \subseteq V$, such that $\hat{C}(S, \omega)$ contains no other end. For an ordinal α, we say an end ω has rank α, if it has not already been assigned a smaller rank and if there is a finite set $S \subseteq V$ such that all ends in $\hat{C}(S, \omega)$ have been assigned a rank, and all these ranks are strictly smaller than α.

For a graph G in which every end has a rank (i.e. for graphs where $V \cup \Omega(G)$ is scattered), we define the *end-rank* $r(G)$ as the supremum of the ranks of all points in $V \cup \Omega(G)$.[4]

Theorem 3.25. *The following are equivalent for any connected graph G:*

(i) There is an upwards connected rayless tree-decomposition of finite adhesion with connected parts distributing $\Omega(G)$.

(ii) There is a tree-decomposition of finite adhesion distributing $\Omega(G)$.

(iii) $V \cup \Omega(G)$ has a σ-discrete expansion.

(iv) G contains no end-faithful subdivision of the full binary tree T_2.

(v) Every end of G has a rank, i.e. $\Omega(G)$ is scattered.

Moreover, if $\Omega(G) \neq \varnothing$, we may add

(vi) There is an upwards connected rayless tree-decomposition of finite adhesion with connected parts bijectively distributing $\Omega(G)$.

Proof. $(i) \Leftrightarrow (ii) \Leftrightarrow (iii)$ is a special case of Theorem 3.21.
For the implication $(iii) \Rightarrow (iv)$ note that any subspace of $V \cup \Omega(G)$ inherits the property of having a σ-discrete expansion. However, the end space of a binary tree does not have a σ-discrete expansion: Indeed, any discrete set in a compact metric space is just countable; but the end space of a binary tree is uncountable, so not a countable union of countable sets.

[4] We remark that in this formulation, $r(G)$ and the Cantor-Bendixson rank of $V \cup \Omega(G)$ may differ by ± 1.

The equivalence $(iv) \Leftrightarrow (v)$ is the content of Jung's [26, Satz 4].

We prove $(v) \Rightarrow (iii)$ by transfinite induction on the end-rank α of G. In the base case $r(G) = 0$, i.e. when $\Omega(G) = \emptyset$, we may take the trivial expansion consisting just of the vertex set.

Now let $\alpha > 0$, and suppose that all graphs of rank $< \alpha$ admit a σ-discrete expansion. First, let $\Phi \subseteq \Omega(G)$ consist of all ends of rank α. Clearly, Φ is a closed discrete subset of $\Omega(G)$. By Corollary 3.14, there is a connected envelope U for Φ, i.e. U is a connected set of vertices in G of finite adhesion such that $\partial U = \Phi$. Write $\mathcal{P}$ for the collection of components of $G - U$, and note that for each $P \in \mathcal{P}$, all ends living in P have rank $< \alpha$.

Now for each component $P \in \mathcal{P}$ individually, consider a collection $\mathcal{C}_P = \{C_P(S_\omega, \omega) : \omega \in \Omega(P)\}$ such that each set $C_P(S_\omega, \omega)$ witnesses the rank of ω inside the graph P. By Theorem 2.11, there is a rayless normal tree N_P in P such that every component D of $P - N_P$ is included in an element of $\mathcal{C}_P$ and hence satisfies $r(D) < \alpha$. Note that $U' = U \cup \bigcup_{P \in \mathcal{P}} N_P$ also is an envelope for Φ, but now, writing $\mathcal{P}'$ for the collection of components of $G - U'$, we have $r(D) < \alpha$ for every $D \in \mathcal{P}'$. By induction assumption, each $D \in \mathcal{P}'$ admits a σ-discrete expansion

$$V(D) \cup \Omega(D) = \bigcup_{n \geq 1} X_{D,n}.$$

Then $X_0 := \overline{U'} = U' \cup \Phi$ together with

$$X_n := \bigcup_{D \in \mathcal{P}'} X_{D,n}$$

for $n \geq 1$ gives the desired σ-discrete expansion of $V \cup \Omega(G)$. Indeed, to see that $X_0 \cup X_1 \cup \cdots \cup X_n$ is closed for every $n \in \mathbb{N}$, note that every end ω of G outside of this set lives in some component D for $D \in \mathcal{P}'$. Let $S \subseteq V(D)$ be finite such that $\hat{C}_D(S, \omega)$ is a basic open set inside $V(D) \cup \Omega(D)$ separating ω from the closed set $X_{D,1} \cup \cdots \cup X_{D,n}$. But then $\hat{C}_G(S \cup N(D), \omega)$ is a basic open neighbourhood of ω in $V \cup \Omega(G)$ witnessing that ω does not belong to the closure of $X_0 \cup X_1 \cup \cdots \cup X_n$. This completes the induction step and the proof of $(v) \Rightarrow (iii)$.

Finally, the moreover part $(i) \Leftrightarrow (vi)$ is immediate from Corollary 3.24. $\square$

Using different methods, Polat showed that $\Omega(G)$ has a σ-discrete expansion if and only if every end of G has a rank [40, Theorem 8.11].

3.9 Applications

3.9.1 Tree-decompositions displaying special subsets of ends

Through our main characterisation, we can now give a short proof of the main result from Carmesin's [9].

Theorem 3.26. *Every connected graph G has a tree-decomposition of finite adhesion with connected parts that displays precisely the undominated ends of G.*

Proof. Let Ξ be the set of all ends of G which are dominated. By Theorem 3.17 it suffices to show that $\Omega(G) \setminus \Xi$ is a G_δ set in $|G|$, and by Lemma 2.13 it is equivalent to show that $V(G) \cup \Xi$ is an F_σ set in $|G|$. Choose an arbitrary vertex $u \in V(G)$ and for all $n \in \mathbb{N}$ write X_n for the set of all vertices of G with distance at most n to u. We show that $V(G) \cup \Xi = \bigcup_{n\in\mathbb{N}} \overline{X_n}$. We have $V(G) = \bigcup_{n\in\mathbb{N}} X_n$ because G is connected. It is left to show that the ends in $\bigcup_{n\in\mathbb{N}} \overline{X_n}$ are precisely the dominated ends of G.

Consider any end $\omega \in \Omega(G)$ and let R be an ω-ray in G. First, suppose that ω is dominated and let v be the centre of an infinite subdivided star S with leaves in R. Furthermore, suppose that $v \in X_n$. Then $S - v$ is a comb attached to $N(v) \subseteq X_{n+1}$ and therefore ω is contained in $\overline{X_{n+1}}$.

Now assume for a contradiction that some $\overline{X_n}$ contains an undominated end ω, and choose n minimal with that property. Then there is a comb C attached to X_n with spine $R \in \omega$. By minimality of n, there is an infinite set $\mathcal{T}$ of teeth of C which lie in $X_n \setminus X_{n-1}$. The neighbourhood of $\mathcal{T}$ in X_{n-1} is finite, again by minimality of n. Since every vertex in $X_n \setminus X_{n-1}$ has a neighbour in X_{n-1}, there is vertex $v \in X_{n-1}$ with infinitely many neighbours in $\mathcal{T}$. Hence ω is dominated by v, a contradiction. $\square$

The following generalises a corresponding result from [8, Theorem 2].

Theorem 3.27. *For every infinite set of vertices U in a connected graph G, there is a tree-decomposition of G of finite adhesion that displays precisely the undominated ends of ∂U.*

Proof. Without loss of generality, we may assume that U has finite adhesion (Theorem 3.10).

Consider the contraction minor $H \preceq G$ obtained from G by contracting each component C of $G - U$ to a single vertex v_C (of finite degree).

Claim 3.28. *The inclusion $U \hookrightarrow H$ induces a bijection $\partial U \to \Omega(H)$ that preserves the property of being dominated.*

This claim is proven just like Lemma 3.13.

Claim 3.29. *The contractions resulting in H induce a natural continuous surjection $f\colon |G| \to |H|$.*

To see that f is continuous, consider some end $\omega \in |G|$. If $\omega \notin \partial U$, then $f(\omega) = v_C$ for some component C, and f is continuous at ω. If $\omega \in \partial U$, then $f(\omega) = \omega' \in \Omega(H)$ by Claim 3.28. Let $C_H(X', \omega')$ be an arbitrary basic open neighbourhood around ω' in H. Let $X \subseteq U$ be the finite set of vertices where we replace every vertex of the form v_C in X' by $N(C)$. It remains to verify that

$$f[C_G(X, \omega)] \subseteq C_H(X', \omega').$$

But this is clear: for every $v \in C_G(X, \omega)$, any $v - \omega$-ray R avoiding X is mapped to a locally finite connected subgraph in H avoiding X' which includes an $f(v) - \omega'$-ray R'.

Now we apply Theorem 3.26 inside H to see that there is a tree-decomposition of finite adhesion displaying the undominated ends Ψ of H. Hence Ψ is G_δ in $|H|$ by Theorem 3.17, say $\Psi = \bigcap_{n \in \mathbb{N}} O_n$ with O_n open in $|H|$. But then by Claim 3.29,

$$f^{-1}(\Psi) = f^{-1}(\bigcap_{n \in \mathbb{N}} O_n) = \bigcap_{n \in \mathbb{N}} f^{-1}(O_n)$$

is G_δ in $|G|$. Thus $f^{-1}(\Psi)$ can be displayed by a tree-decomposition of finite adhesion of G, again by Theorem 3.17. This completes the proof as $f^{-1}(\Psi)$ is the set of all undominated ends in ∂U by Claim 3.28. $\square$

3.9.2 Counterexamples

Consider the full infinite binary tree T_2, and let $X \subseteq \Omega(T_2)$ be any set of ends. A *binary tree with tops* X is the graph with vertex set $T_2 \sqcup X$, all edges of T_2, and such that the neighbourhood of $x \in X$ consists of infinitely many nodes on its corresponding normal ray in T_2.

We reobtain Carmesin's observation that a T_2 with uncountably many tops does not admit a tree-decomposition of finite adhesion displaying all its ends, but now with significantly shorter proof.

Example 3.30. *No binary tree with uncountably many tops admits a tree-decomposition of finite adhesion displaying all its ends.*

Proof. These graphs do not have normal spanning trees by [18, Proposition 3.3], and so the result follows from Theorem 3.9. $\square$

With only a little more work, we can prove the following stronger result by Carmesin [9, p.7].

Example 3.31. *No binary tree with uncountably many tops admits a tree-decomposition of finite adhesion distinguishing all its ends.*

Proof. Let G be a binary tree with uncountably many tops. Suppose for a contradiction that $V(G)$ is a countable union of slender sets. Then one of the slender sets U contains uncountably many of the tops. Write $\mathcal{R}$ for the set of all normal rays of T_2 which have a corresponding top in U. We call a vertex v of T_2 *good*, if it lies in uncountably many rays from $\mathcal{R}$. It is clear that the root of T_2 is good. We now show that for each good vertex v, there are two incomparable good vertices above v in the tree-order:

Suppose not for a contradiction. It is clear that at least one upper neighbour in T_2 of each good vertex is good. This implies that there is a ray R of good vertices above v. Since per assumption all good vertices above v are comparable, no other vertex above v outside the ray R is good. But this ray has only countable many neighbours in T_2. As no such neighbour above v is good, every neighbour of R above v lies on only countably many rays from $\mathcal{R}$. But then also v lies on only countably many rays from $\mathcal{R}$, which is a contradiction since v is good.

From this claim follows that there is a subdivided binary tree inside G such that each branch vertex is good.

It follows that ∂U itself contains the end space of a subdivided binary tree. But the end space of a binary tree is not scattered, a contradiction. It follows from Theorem 3.22 that $\Omega(G)$ cannot be distinguished. $\square$

We conclude this section with a new example of a graph G witnessing that the thin ends of G cannot always be displayed, that is based on topological considerations only (a different example is given by Carmesin in [9, Example 3.3]). More precisely, since displayable subsets of ends are always completely metrizable by Corollary 3.18, it suffices to construct a graph where the thin ends are not completely metrizable.

As a warm-up, consider the binary tree T, and call a normal ray of T *rational* if its corresponding $0-1$-sequence becomes eventually constant, and *irrational* otherwise. Let $\Sigma \subseteq \Omega(T_2)$ be the subspace of rational ends. By Sierpinski's characterisation [41], every countable metric space without isolated points – so in particular Σ – is homeomorphic to the rational numbers $\mathbb{Q}$. Thus, Σ is not completely metrizable, and hence not displayable.

We now modify T such that all irrational ends become thick, and all rational ends remain thin. A binary tree with *fat* tops Z is a graph with vertex set $T \sqcup Z$, all edges of T_2, and such that the neighbourhood of $z \in Z$ consists of infinitely many nodes on some normal ray R_z of T_2. Thus, the difference between a tree with tops and tree with fat tops is that a normal ray may now have more than one top vertex.

Example 3.32. *There is a binary tree with uncountably many fat tops such that its thin ends cannot be displayed.*

Construction. Starting from the binary tree T, let $\{R_i \colon i \in \mathbb{N}\}$ be an enumeration of the rational rays in T. We now add infinitely many top-vertices above each irrational ray, and connect them to their rays such that

1. each top-vertex z dominates its corresponding irrational ray R_z, and

2. for each rational ray R_i, at most i vertices on R_i have top-vertices as neighbours.

Once the construction has finished, it is clear that the resulting graph G is as desired. The end space $\Omega(G) = \Omega(T)$ remains unchanged. From (2) it is easy to see that every rational end $\omega_i \ni R_i$ has end-degree 1 in G (and hence is thin), since the corresponding ray R_i has a tail of vertices of degree 3 whose edges are cut edges. All irrational ends are dominated by their infinitely many top-vertices and thus become thick.

By Sierpinski's characterisation [41], the set of rational / thin ends is homeomorphic to the rational numbers $\mathbb{Q}$, so not completely metrizable, and hence not displayable by Corollary 3.18.

It remains to describe how to connect a top-vertex z to its irrational ray R_z. For each top-vertex z and every $j \in \mathbb{N}$, let r_z^j be the $\leq_T$-minimal vertex in $R_z \setminus (R_0 \cup \ldots \cup R_j)$. Now let the neighbours of z be exactly the vertices in $\{r_z^0, r_z^1, r_z^2, \ldots\}$. Since $r_z^0 \leq r_z^1 \leq r_z^2 \leq \cdots$ is cofinal in R_z, the top-vertex z dominates R_z, establishing property (1).

Next, consider the ith rational ray R_i. Again, for $j \leq i$ let r_i^j be the $\leq_T$-minimal vertex in $R_i \setminus (R_0 \cup \ldots \cup R_j)$ (if it exists). Then it is clear that $r_i^0, r_i^1, \ldots, r_i^{i-1}$ are the only vertices on R_i adjacent to top-vertices, giving (2). $\square$

4 Hamilton circles in powers of infinite graphs

4.1 Introduction

For a given graph $G = (V, E)$, we obtain its nth power G^n by adding an edge between any two vertices for which its distance in G is at most n. As early as in 1960, Sekanina proved that the third power G^3 of any connected finite graph G has a Hamilton cycle [33]. The original proof was by induction, showing that for a fixed root r, we can find a Hamilton cycle for which one of the two edges at r lies in G itself. For our application we will control a bit more how the Hamilton cycle lies in the graph, thus we will state a constructive proof in Chapter 4.2.3 of a slightly stronger result.

A general approach to extend finite theorems to locally finite graphs is by working in the topological space $|G|$ and using a compactness argument. Agelos Georgakopoulos proved the theorem about Hamilton circles in the third power for locally finite graphs, by using the compactness principle:

Theorem 4.1. [24] *If G is a connected locally finite graph, then G^3 has a Hamilton circle.*

A circle in this context is a topological circle in $|G^3|$ defined as in the end of section 2.3. Even though Georgakopoulos conjectured that this is also true for all countable connected graphs, there are several counterexamples. To understand those counterexamples, we have to think about the fact that the endspace $\Omega(G^3)$ may differ structurally from $\Omega(G)$ itself.

We define V_∞ as the set of vertices of infinite degree in G. For each vertex $v \in V_\infty$, its G-neighborhood becomes an infinite clique already in G^2, which contains new rays that are not necessarily equivalent to any ray of G. In G^3 or even higher powers of G, some of those cliques will belong to the same end.

The starting point for this chapter are the results of my master thesis, in which I characterized the trees with a Hamilton circle in their third power. We will introduce this result by giving the key counterexamples and further motivate why the problem is even more complex for arbitrary countable graphs.

Lemma 4.2. *If T is a countable tree, then the set of ends in T^3 can be written as a disjoint union $\Omega_1 \cup \Omega_2$, with the following properties:*

1. *There exists a canonical injection from Ω_1 to the set of ends of T, in which each end in Ω_1 is mapped to a subset of itself and the image of this injection is the set of ends of T not containing a ray of vertices in V_∞. We call the ends in Ω_1* preserved ends.

2. *Ω_2 consists of one end for each component K of $T[V_\infty]$, where every ray of that end meets the union of that component with its T-neighborhood infinitely often. We call such an end a* new end *in K.*

In the following counterexamples, T is rayless and $T[V_\infty]$ has only one component, hence T^3 is one-ended and a Hamilton circle would be a spanning double-ray together with the new end. To explain those examples, we will use some terms without precisely defining them. However, we will state the ideas of them based on our examples. The precise definitions and full characterisation of all countable trees T, for which T^3 has a Hamilton circle, can be found in my master thesis.

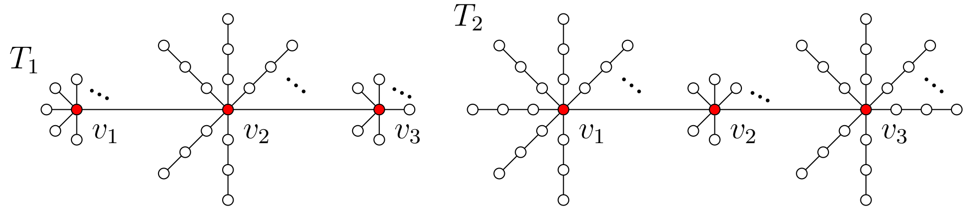

To understand why no spanning double-ray exists, let us assume that we had a spanning double-ray R. Consider for example the left subdivided star S in $T := T_2$. Each of its leaves has 3 neighbors in T^3: The other two vertices on its subdivided edge and the center v_1 of the star. But v_1 has only two neighbors in R, implying that for all but at most 2 of the leaves, its neighbors in R are the two subdividing vertices:

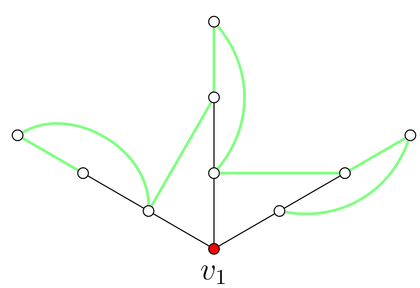

The T-neighbor of such a leaf also has another neighbor in R. With a similar argument, we obtain that for almost all subdivided leaves, that this neighbor is a T-neighbor of v_1. We say that such components of $G - S$ are *not well-coverable*. We obtain a ray in R, consisting of the vertices of infinitely many components of $G - v$. We call such a ray *captured*. If there are two disjoint captured rays B_1, B_2 in R, then $R - (B_1 \cup B_2)$ must be finite, but still contains infinitely many leaves of one (or both) non-subdivided stars, a contradiction.

Similarly we find such a captured ray for each of the subdivided stars. In T_2 we obtain one such captured ray for each of the subdivided stars. In T_1, if one captured ray B does not cover almost all of the components of $T_1 - v_2$, we can find a second one in $T_1^3 - B$ and obtain the same contradiction. If not, then $S := v_2 \cup N_{T_1}(v_2)$ is a separator of T_1^3 and $S|_{R-B}$ is a finite separator of $R - B$. But $R - B$ contains infinitely many paths between $N_T(v_1)$ and $N_T(v_3)$, each of them containing a vertex in $S|_{R_2}$, a contradiction.

Every other counterexample for trees fails in a similar way to be hamiltonian. A component of $T[V_\infty]$ which prohibits the existence of a Hamilton circle in one of those two ways is called splitting. A Hamilton circle of T^3 exists exactly when there is no such splitting component. The following theorem is shown in my master thesis:

Theorem 4.3. *If T is a tree, then T^3 has a Hamilton circle if and only if no component of $T[V_\infty]$ is splitting.*

On first glance, one may think that the characterization of the countable connected graphs with a Hamilton circle in their third power might be somehow directly related to the characterisation for trees, for example by finding Hamilton circles in the third power of a certain spanning tree T. But in general $T \subseteq G$ does not imply $|T^3| \subseteq |G^3|$. The graph G^3 can have fewer ends than T^3. Even the existence of a Hamilton circle for each spanning tree is not enough to find one in G^3:

Lemma 4.4. *There is a rayless graph G such that G^3 has no Hamilton circle, but every spanning tree of G has a Hamilton circle in its third power.*

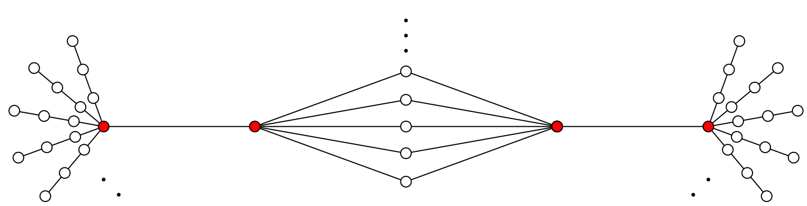

Conversely, there are graphs with a Hamilton circle in their third power, for which there is none for every spanning tree:

Lemma 4.5. *There is a rayless graph G such that G^3 has a Hamilton circle, but no spanning tree of G has a Hamilton circle in its third power.*

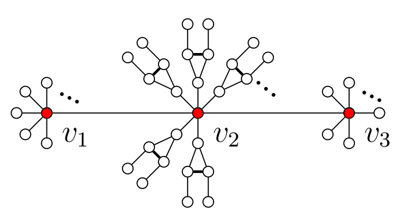

In the first example G^3 is one ended but has two vertices inducing captured rays and still infinitely many vertices left in the middle, thus with a similar argument as for the first of our initial examples, one can prove that G^3 has no Hamilton circle. But for each spanning tree T we keep only one from the infinitely many paths in the middle and thus obtain two different new ends in T^3 for components of T_∞ of size 2. Because now each of the components has only one vertex forcing a captured arc, both of them are not splitting anymore, so we can construct a Hamilton circle in T^3.

In the second example, the finite components of $G - v_2$ are well-coverable in a similar way as in our definition for trees, so we do not obtain a captured arc, but when we choose a spanning-tree T, we turn each of the components in one or two not well-coverable ones, so we obtain no Hamilton circle in T^3.

The precise notions and a proof why those are indeed counterexamples is included in our characterisation of Section 4.2.

Those two examples motivate the extension of our question from trees to arbitrary countable graphs. In the next section we present a characterization of all rayless graphs with a Hamilton circle in their third power. It turns out that our characterization of the graphs with Hamilton circle in their third power can be formulated in a similar way with another definition of *splitting* components. However, the proofs contain a lot more interesting and challenging technical details.

Theorem 4.6. *For a countable rayless graph G, G^3 has a Hamilton circle if and only if no class of $V_\infty/\sim$ is splitting.*

The last section of this chapter is about higher powers of trees. It seems natural that every countable tree has a Hamilton circle in the fourth or higher power and indeed, this is the case as we will prove in this section. But due to the fact that the endspace may change for different powers, the proof is not straightforward. In particular, it is worth mentioning that to build a Hamilton circle of T^n we might be forced to use edges of $E(T^n) \setminus E(T^{n-1})$. Because of that, we cannot apply induction on n or any similar argument. Instead, we will do a direct construction of a Hamilton circle of T^n for each $n \geq 4$.

Theorem 4.7. *For a countable tree T and any $n \geq 4$, T^n has a Hamilton circle.*

4.2 Hamilton circles in the third power of rayless graphs

4.2.1 The end space of the third power of G

The goal of this section is to characterize the ends of G^3 for any rayless connected graph $G = (V, E)$. Since G itself does not have any ends, all of its rays arise when building the third power.

We define $V_\infty \subseteq V$ as the set of vertices of V with infinite degree. For every $v \in V_\infty$, its neighborhood in G becomes an infinite clique in G^2, and hence also in G^3. To understand the endspace of G^3, it is crucial to find out for which vertices the new rays in G^3 arising through those cliques belong to the same end. We will describe this concept by the following equivalence relation:

Definition 4.8. For a graph $G = (V, E)$, we define the graphs $G_\infty^\sim = (V_\infty, E^\sim)$ on V_∞ with edges between each two vertices $v, w \in V_\infty$ for which either $vw \in E(G)$ or $\kappa_G(v, w) = \infty$.

For $v, w \in V_\infty$, we say $v \sim w$ whenever they are in the same component of $G_\infty^\sim$. Two vertices which are equivalent in this relation are called *weakly equivalent*.

The weak equivalence classes can be further partitioned into *strong equivalence classes*:

Definition 4.9. For a graph $G = (V, E)$, we define the graph $G_\infty^\approx = (V_\infty, E^\approx)$ on V_∞ with edges between each two vertices $v, w \in V_\infty$ for which $\kappa_G(v, w) = \infty$.

For $v, w \in V_\infty$, we say $v \approx w$ whenever they are in the same component of $G_\infty^\approx$. Two vertices which are equivalent in this relation are called *strongly equivalent*.

Lemma 4.10. *If $G = (V, E)$ is a rayless graph, then $G_\infty^\sim$ and $G_\infty^\approx$ are also rayless.*

Proof. Suppose there was a ray $R = v_1, v_2, v_3, ...$ in $G_\infty^\sim$. We construct a ray in G recursively. Define $P_0 = v_1$, and construct a sequence of paths $P_1 \subset P_2 \subset P_3...$ satisfying the following:

(i) P_i begins in v_1 and ends in a vertex v_k of R

(ii) P_i contains no v_i for an $i > k$.

Given P_n ending in v_k, we define P_{n+1} as follows: If there is an edge $v_k v_{k+1}$ in G, then we obtain P_{n+1} by adding this edge to P_n. This edge is not already in P_n, because of (ii). It is clear that P_{n+1} satisfies (i) and (ii).

If there is no edge $v_k v_{k+1}$ in G, then there are infinitely many internally disjoint $v_k - v_{k+1}-$paths in G. Let P be one of them which meets P_n only in v_k. Also let v_l be the first vertex of R (after v_k) on this path. Then we obtain P_{n+1} from P_n by adding $v_k P v_l$. Again it is clear that (i) and (ii) hold for P_{n+1}.

The union $\bigcup_{i\in\mathbb{N}} P_i$ is a ray in G, a contradiction. The graph $G_\infty^{\approx}$ is a subgraph of $G_\infty^{\sim}$ and hence also rayless. $\square$

For the rest of this section, let $G = (V, E)$ be a fixed rayless graph. We will see that its third power G^3 has exactly one end for each weak class of V_∞.

Lemma 4.11. *Let K be a component of $G_\infty^{\approx}$ or $G_\infty^{\sim}$. For a vertex $w \in V_\infty$, for which there is an infinite subdivided star in G with center w and infinitely many leaves $(v_i)_{i\in\mathbb{N}}$ in K follows that also $w \in K$.*

Proof. Since K is infinite and $G_\infty^{\approx}$, $G_\infty^{\sim}$ are rayless (4.10), according to the star-comb Lemma 2.2, there is a subdivided infinite star S in $G_\infty^{\approx}$ or $G_\infty^{\sim}$ with a center v and leaves in $v_1, v_2, v_3, \ldots$.

Let $v_{i_1}, v_{i_2}, v_{i_3}, \ldots$ be the leaves of S. We will show that we also find an infinite star in G, by recursively defining infinitely many internally disjoint paths starting at v and ending in different vertices from $v_{i_1}, v_{i_2}, v_{i_3}, \ldots$: For two vertices, which are weakly, but not strongly equivalent, we call the G-edge between them *important*.

Assume that we have already constructed finitely many (possibly none) such paths. Let $\bigcup P$ be the union of these paths. Because $\bigcup P$ is finite, it meets only finitely many important edges. Let v_{i_k} be one leaf of S such that the $v - v_{i_k}-$path S_k in S contains no vertex from $\bigcup P$, and also for each two vertices on S_k, which are not strongly equivalent, does $\bigcup P$ not contain the associated important edge. We choose for every edge between two strongly equivalent vertices on S_k one path in G between them, which does not meet $\bigcup P$ and for the other edges, we choose the associated important edge.

The union of all the chosen paths and edges is a finite connected graph, disjoint from $\bigcup P$. This graph contains the desired $v - v_{i_k}-$path in G.

After we have done this step countably many times, we obtain the infinitely many internally disjoint paths starting at v and ending in different vertices from $v_{i_1}, v_{i_2}, v_{i_3}, \ldots$. Now these two subdivided stars with centers v and w satisfy the properties of Lemma 2.3. It follows that $\kappa_G(v, w) = \infty$ and hence $w \in K$. $\square$

Lemma 4.12. *Let $[v] \neq [w]$ be two strong or weak classes of V_∞. Then there are at most finitely many internally disjoint $[v] - [w]-$paths in G.*

Proof. Assume for a contradiction that there are infinitely many internally disjoint $[v]-[w]-$paths in G. Let G' be the graph G after identifying all vertices of $[w]$ (Note that $G[[w]]$ is not necessary connected).

Then we can apply Lemma 4.11 to G' and obtain a vertex $v' \in [v]$ with infinitely many internally disjoint paths to the contraction vertex w in G'. In G itself those paths are between v' and $[w]$.

If infinitely many of them are ending in the same vertex $w' \in [w]$ then $v' \approx w'$. If not, then they end in infinitely many different vertices in $[w]$ so we find a subdivided infinite star in G with center v' and leaves in $[w]$.

We apply Lemma 4.11 a second time to show that $v' \in [w]$ and thus $[v] = [w]$, a contradiction. $\square$

Lemma 4.13. *Let v, w be in V_∞. Then $v \approx w$ if and only if there are infinitely many internally edge-disjoint paths between them.*

Proof. For the forward implication, consider a $v - w-$path in $G_\infty^\approx$. Choosing for each edge of this path one of the infinitely many paths in G between its endvertices, we find a finite subgraph containing a $v - w-$path. By choosing those paths recursively such that none of it meets any edge of the constructed graphs before, we obtain infinitely many edge-disjoint $v - w-$paths.

For the backward direction, assume that there are infinitely many internally edge-disjoint paths $P_1, P_2, \ldots$ between v and w. We call those paths *original*. We will recursively construct a $v - w-$path $v = w_0, w_1, \ldots, w_n = w$ in $G_\infty^\approx$.

For each original path let v_i be the first vertex after v on P_i. Since from each v_i there is a path to w in G, the union of these paths is a connected subgraph G_1 of $G - v$. Because G is rayless, we obtain a star with a center w_1 in G_1 with leaves in $\{v_1, v_2, \ldots\}$. (2.2). Now it follows from 4.11 that $vw_1 \in E(G_\infty^\approx)$.

If any $w_i = w$, we are done. If not, we define G_{i+1} as the subgraph of G_i as the union of the subpaths between w_i and w from the infinitely many original paths that meet w_i. Now we can obtain w_{i+1} inside of G_{i+1} in the same way as we obtained w_1 in G_1. Since each G_{i+1} is disjoint from the paths between $w_0, w_1, \ldots, w_i$ apart from w_i, it follows that $w_{i+1} \notin \{w_0, w_1, \ldots, w_i\}$.

We can always continue the sequence, unless $w_{i+1} = w$ at some point, but if this construction does not end, we would obtain a ray in $G_\infty^\approx$ and hence also in G (4.10), a contradiction. It follows that $v = w_0 \approx w_1 \approx \ldots \approx w_n = w$. $\square$

The assumption for G to be rayless is indeed essential for the last lemma to hold. The following *Farey graph* is a counterexample of a not rayless graph.

Example 4.14. *Let H_0 be the graph K_2 with 2 vertices v, w and one edge vw. Now let H_1 be the triangle C_3 with vertices v, w, x . Now to obtain H_{i+1} for $i > 1$, add for each edge in $E(H_i) \setminus E(H_{i-1})$ one parallel edge to it and subdivide it once. Let H be the limit of this sequence. Now in H every vertex has infinite degree, but for each two vertices there are at most 3 internally disjoint paths between them, thus no two vertices are strongly equivalent. By subdividing every edge of H, we can even obtain a graph in which every weak class is a singleton.*

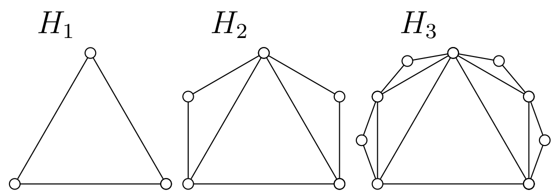

Substantial research on the Farey graph was done by Jan Kurkofka. One of his results was that the Farey graph is the unique graph up to minor-equivalence that is infinitely edge-connected but such that every two vertices can be finitely separated [28].

Lemma 4.15. *Let v, w be in V_∞. Then $v \sim w$ if and only if there is no finite $v-w-$separator with finite neighborhood in G.*

Proof. Suppose for a contradiction that $v \sim w$ and there is a finite $v - w-$separator S with finite neighborhood in G. There is a $v - w-$path in $G_\infty^\sim$ that witnesses $v \sim w$. We will show that for each edge v_1v_2 of P, there is a path between v_1 and v_2 in G that avoids S. The union of these paths contains a $v - w-$path in G avoiding S, a contradiction.

In the case that there is an edge between v_1v_2 in G, each $v_1 - v_2-$separator contains v_1 or v_2 and hence would have infinite neighborhood. Hence the edge v_1v_2 avoids S. If there are infinitely many $v_1 - v_2-$paths in G, there is no finite $v_1 - v_2-$separator in G at all. Hence, there is a $v_1 - v_2-$path avoiding S.

Suppose now that $v \nsim w$. We will construct a finite $v - w-$separator with finite neighborhood: By Lemma 4.12, there is a finite maximal set $\{P_1, ..., P_n\}$ of internally disjoint $[v]_\sim - [w]_\sim-$paths in G.

Now let S_n be the set of vertices of finite degree on those paths. If S_n is a $v - w-$separator, then we are done. If not, there is a $[v]_\sim - [w]_\sim-$path P_{n+1} avoiding S_n. Define S_{n+1} in the same way as S_n for the set of paths $P_1, ..., P_{n+1}$. Iterate this as long as possible. If we do not get a finite separator (with finite neighborhood) after finitely many steps, we obtain infinitely many different paths $P_1, P_2, ...$ in G between $[v]_\sim$ and $[w]_\sim$ from which each of them meets $P_1, ..., P_n$. This implies that the union P of these paths has at most n components. Further, since $v \nsim w$, every $v-w-$path contains at least one vertex of finite degree, so each path we construct does indeed contain at least one vertex that does not lie on any of the earlier paths. This implies that P is infinite. It follows that P has at least one infinite component and since G is rayless, there is an infinite subdivided star in P. It follows from 4.11 that its center is also in $[v]_\sim$, a contradiction. □

To understand how new rays in G^3 arise and which pairs of them are equivalent, we introduce the concept of manifestations. This will not only lead to our main result of this subsection that we obtain exactly one new end for each weak class but further gives us an idea of how those ends look like.

Definition 4.16. For an element v of V_∞, we define a *simple manifestation* of v in G^3 as a ray in G^3 in the G-neighborhood of v.

Lemma 4.17. *For each edge v_1v_2 in $G_\infty^\sim$, any simple manifestations of v_1 and v_2 are equivalent in G^3*

Proof. Let R_1 be a simple manifestation of v_1 and R_2 be a simple manifestation of v_2. If there is an edge v_1v_2 in G, then every vertex of R_1 is adjacent to every vertex of R_2 and the rays are clearly equivalent. If $v_1 \approx v_2$, then let $P_1, P_2, P_3, ...$ be infinitely many disjoint paths between $N_G(v_1)$ and $N_G(v_2)$ (In this context, we allow paths to have only one vertex, if it is in $N_G(v_1) \cap N_G(v_2)$.)

If infinitely many of these paths are already between R_1 and R_2, then we are done. If not, then we assume without loss of generality that none of them are. Adding one vertex of R_1 and one of R_2 to each path in a disjoint way gives us infinitely many disjoint $R_1 - R_2-$paths in G^3. □

Lemma 4.18. *Two simple manifestations of vertices v_1 and v_2 are in the same end of G^3, if and only if $v_1 \sim v_2$.*

Proof. The backward direction follows by induction on the distance of v_1, v_2 in $G_\infty^\sim$ from Lemma 4.17. Suppose now for a contradiction that simple manifestations R_1 and R_2 of v_1 and v_2 are in the same end of G^3, but $v_1 \nsim v_2$. It follows from Lemma 4.15 that there is a finite $v-w-$separator S with finite neighborhood in G; clearly S cannot contain v or w. We show that $S \cup N_G(S)$ separates the rays R_1 and R_2 in G^3: If not, there is a $R_1 - R_2-$path P avoiding $S \cup N_G(S)$ in G^3. We replace each edge xy of P by a $x-y-$path P_{xy} of length at most 3 in G. Because x and y does not lie in the neighborhood of S, they both have distance at least 2 from S and the path P_{xy} cannot meet S. Hence after replacing all edges in P like this, we find a $R_1 - R_2-$path $P' := v_1, v_2, ..., v_n$ avoiding S in G. By adding the edge vv_1, if v is not already in P' and $v_n w$, if w is not already in P', we obtain a $v - w-$path in G avoiding S, a contradiction.

□

Definition 4.19. For an element $v \in V_\infty$, a *manifestation* of v is a ray R in G^3 such that there is an infinite subdivided star in G with center v and leaves in R.

Lemma 4.20. *For each element $v \in V_\infty$, any two manifestations of v are equivalent.*

Proof. Due to Lemma 4.18 it suffices to show that each manifestation R of v is equivalent to any simple manifestation R' of v: Consider a subdivided star with center v and leaves in R. The union of R, this subdivided star and the clique in G^3 consisting of $N_G(v)$ is a subgraph of G^3 containing R and R', but no finite set separating them. Hence R and R' are also equivalent in G^3. □

Corollary 4.21. *Two manifestations of elements $v_1, v_2 \in V_\infty$ are equivalent if and only if $v_1 \sim v_2$.*

Proof. The backward direction follows from 4.20 and 4.18.
If two manifestations R_1, R_2 of elements $v_1, v_2 \in V_\infty$ are equivalent, then we also find with 4.20 two simple manifestations R_1', R_2' of v_1, v_2 such that $R_1 \sim R_1'$ and $R_2 \sim R_2'$. It follows that $R_1' \sim R_2'$, which implies with the forward direction of 4.18 that $v_1 \sim v_2$. □

Proposition 4.22. *For any rayless graph $G = (V, E)$, there is a canonical bijection between $V_\infty/\sim$ and $\Omega(G^3)$, in which every equivalence class $[v]_\sim$ of $V_\infty/\sim$, is mapped to the end G^3 consisting of the rays that meet $[v]_\sim \cup N_G([v]_\sim)$ in infinitely many vertices. The rays in one such end are exactly the manifestations of all vertices from the class.*

Proof. Let R be any ray of G^3. Applying the star-comb Lemma 2.2 in G on $V(R)$, we can show that R is indeed a manifestation of an element $v \in V_\infty$. The Lemma 4.21 finishes our proof. □

We call the image of a weak class $[v]_\sim$ under the bijection from the previous lemma the end *induced* by $[v]_\sim$.

4.2.2 Counterexamples of graphs with no Hamilton circle in their third power

As already mentioned in the introduction, the existence of a Hamilton circle in the third power of a rayless graph is dependent on each of the new ends or, in terms of 4.22, on each weak class. In this section we will precisely define the concept of captured arcs and splitting classes. Further, we conclude the section with a proof that every rayless graph with a splitting class fails to be Hamiltonian in its third power.

Again, for this whole section let $G = (V, E)$ be a countable connected rayless graph.

Lemma 4.23. *Let $[v]$ be any strong or weak class. Then $[v]$ is either infinite or $G - [v]$ has infinitely many components.*

Proof. Assume that $[v]$ is finite. Since all vertices in $[v]$ have infinite degree, it follows that $N_G([v])$ is infinite. Suppose for a contradiction that $G - [v]$ has only finitely many components. This implies that there is a component K of $G - [v]$ containing infinitely many elements of $N_G([v])$. From the star comb Lemma 2.2 and the fact that G is rayless follows that there is a subdivided star in K with a center w and teeth in $N_G([v])$. It follows from 4.12 that $w \in [v]$, a contradiction. □

Definition 4.24. Let $[v]$ be any strong or weak class. A component H of $G - [v]$ is called *not well-coverable*, if

(i) $|V(H)| > 2$,

(ii) $|V(H) \cap N_G([v])| = 1$, say $V(H) \cap N([v]) = \{r\}$,

(iii) $d_H(r)$ is finite and $H - r$ has exactly $d_H(r)$ components, and

(iv) each component of $H - r$ has at least 2 vertices.

Otherwise, H is called *well-coverable*. Vertices in *well-coverable* components of $G - [v]$ are called $[v]$-*good* or simply *good*.

We call the neighbors of $[v]$ in a component K of $G - [v]$ the $[v]$-*roots* or simply *roots* of K.

Lemma 4.25. *A not well-coverable component of $G - [v]_\approx$ for a strong class $[v]_\approx$ cannot meet any other strong class inside $[v]_\sim$.*

Proof. If a component K of $G - [v]_\approx$ meets another strong class inside $[v]_\sim$, this means that it has one vertex which is weakly, but not strongly equivalent to a vertex of $[v]_\approx$. Thus K has a root in V_∞ and is well-coverable. □

Lemma 4.26. *For a subgraph $H \subseteq G$, whenever $N_G(G-H) \cap H$ is finite and contains only vertices of finite degree, then every end of G^3 lives either in H or in $G-H$, so $|H^3|$ can be well-defined as a subspace of $|G^3|$.*

Proof. Follows from Lemma 4.15 and Proposition 4.22. □

Consider a Hamilton circle C of $|G^3|$ and a subgraph $H \subseteq G$ as in the previous lemma. Inside of H is C not necessarily connected. Possibly are the vertices and ends of $|H^3|$ covered by multiple segments of C. Each of these segments is either an arc or a singleton vertex. For simplicity we think of such singletons also as arcs (for which both endvertices are the same) and call a cover of $V(H) \cup \Omega(H^3)$ of disjoint arcs in $|H^3|$ an *arc cover* of H^3.

Lemma 4.27. *Let $H \subsetneq G$ a subgraph, for which $N_G(G-H) \cap H$ is finite. If an arc cover of H^3 is induced by a Hamilton circle C of G^3, all its arcs are between vertices. Furthermore are the endvertices of the arcs of G-distance at most 3 from $G-H$ in G.*

Proof. Each endpoint of an arc from our cover is a point through which C leaves H. The subgraphs H^3 and $(G-H)^3$ have no ends in common by lemma 4.26, hence C cannot leave H through an end. Further for every edge between H and $G-H$ in G^3 is the distance between its two endvertices at most 3 in G, which implies the second part of the lemma. □

If the G-distance to $G-H$ from an endvertex of an arc in $|H^3|$ is lower, there are usually more neighbors outside of H. To make the construction of a Hamilton circle possible, we try to make this distance sufficiently low. This is the crucial difference between well-coverable and not well-coverable components.

We distinguish two types of arc covers of a component H of $G-[v]$ according to whether there is a root of H as a singleton arc or not. In the first case, the best thing possible for a one-rooted component would be if $H-r$ is covered by arcs with endvertices in $N_H(r)$. But this is not possible for not well-coverable components:

Lemma 4.28. *Let $[v]$ be any strong or weak class and H be a not well-coverable component of $G-[v]$ with root r. Then every arc cover of H^3 in which the root r is a singleton has at least one arc with an endpoint outside of $\{r\} \cup N_H(r)$.*

Proof. Let X be the set of singletons in our arc cover apart from r. We may assume that $X \subseteq N_H(r)$, otherwise we are done. Let $\{A_1, A_2, ...\}$ be the remaining finite or countable set of disjoint arcs covering the vertices and ends of $|H^3| - r - X$.

Suppose for a contradiction that each of the arcs has two endvertices in $N_H(r)$. Per definition of not well-coverable, the root r has finite degree (see Definition 4.24 (iii)), say $N_H(r) = \{v_1, v_2, ..., v_l\}$, consequently there are only finitely many arcs $\{A_1, A_2, ..., A_k\}$ and $k \leq (l - |X|)/2$.

Also per definition does $H - r$ have exactly l components, each of them of size at least 2. This implies that each of these components contains exactly one element of $N_H(r)$ and at least one element of distance at least 2 from r (see properties (iii) and (iv)).

Since r has finite degree, it is a finite separator with finite neighborhood. Now Lemma 4.15 implies that each weak class lives in at most one component of $H - r$. Also as in the proof of Lemma 4.15, we can show that $\{r\} \cup N_G(r)$ is a separator in H^3, which leaves at least l components of $H^3 - r - N_H(r)$. We choose l of them and call them $K_1, K_2, ...K_l$. The closures of these components have no ends in common. It follows that for each K_i there are at least two edges between K_i and $N_H(r)$ in the union of the A_i. All in all there are at least $2l$ such edges, but since there is at least one arc A_1, at least two of the vertices $v_1, v_2, ..., v_l$ are endvertices of this arc and hence have degree 1 in A_1. Each other vertex has degree at most 2 in the union of the A_i, because the arcs are disjoint. Hence, there are at most $2l - 2$ edges between $N_H(r)$ and $K_1 \cup K_2 \cup ... \cup K_l$, a contradiction. □

The following definition of *captured arcs* is the fundamental concept of this section. To motivate this definition and to get an idea of how we obtain a captured arc in the proof of 4.30, let us consider the case that a class $[v]_\approx \in [v]_\sim / \approx$ is finite and has only finitely many $[v]_\approx$-good neighbors.

Let us evaluate the question, what the possible neighbors of the arc from the previous lemma in a Hamilton circle of G^3 are. In the optimal case the arc in $|H^3| - r$ is between a vertex in $N_H(r)$ and one vertex of H-distance 2 from r.

The endvertex which is further away from r has as G^3-neighbors outside of $|H^3|$ only some vertices from $[v]_\sim$, namely the $G-$neighbors of r in $[v]_\sim$. Those are per assumption only finitely many. So, by the lack of enough such outside neighbors, almost all of these components cannot contain such an arc and hence their roots cannot be a singleton in the arc cover. That way the best case left for almost all not well-coverable components is an arc between the root r and one of its H-neighbors. The neighbors outside of H in G^3 of the vertices in $N_H(r)$ are vertices of $[v]_\sim$ and roots from some other components of $G - [v]_\sim$. If almost all of those components are not well-coverable, then almost all times the adjacent arc in C is in another not well-coverable component. We will see that we can recursively obtain an arc that we call a *captured arc* through infinitely many of those components. Formally, we define a captured arc as follows:

Definition 4.29. We call an arc $A \subseteq |G^3|$ *captured* by $[v]_\approx$, if it satisfies the following:

(A1) All vertices of A lie in not well-coverable components of $G - [v]_\approx$,

(A2) A contains the root of infinitely many not well-coverable components of $G - [v]_\approx$ and

(A3) A starts in a vertex and ends in the end of G^3 induced by $[v]_\sim$ (see Proposition 4.22).

If an arc is captured by strong class $[v]_\approx$, we call it a $[v]_\approx$-*captured arc* or just *captured arc*.

Lemma 4.30. *Let C be a Hamilton circle of G^3 and $v \in V_\infty$. If a class $[v]_\approx \in [v]_\sim / \approx$ is finite and has only finitely many $[v]_\approx$-good neighbors, then C contains a $[v]_\approx$-captured arc.*

Proof. In this proof, we call components of $G - [v]_\approx$ *bad*, if they are not well-coverable. We know from Lemma 4.23 that $G - [v]_\approx$ has infinitely many components. Because $[v]_\approx$ has only finitely many good neighbors, there are also only finitely many well-coverable components of $G - [v]_\approx$, which we call $K_1, ..., K_j$. Let $K_{j+1}, K_{j+2}, ...$ be the bad components of $G - [v]_\approx$. The root of a bad component K_i we call r_i.

We will show that for almost all of the components K_i of $G - [v]_\approx$ the following assertion holds:

(*) K_i is bad and the arc A_i containing r_i in the arc cover of K_i is not a singleton arc. Furthermore is one C-neighbor of the second endvertex w_i of A_i the root r_j of another bad component of $G - [v]_\approx$.

To see this claim, we will show that every possible way in which a component can fail to satisfy (*) can happen for only finitely many components:

- Only finitely many of the components are not bad.
- For each bad component K_i for which its arc cover contains r as a singleton, due to Lemma 4.28, at least one of the arcs A'_i partitioning $G^3[K_i - r_i]$ has an endvertex v_i with G-distance 3 from $[v]_\approx$. All G^3-neighbors of v_i, which are not in K_i are in $[v]_\approx$. In C does every vertex have degree 2, hence this can happen only for at most $2|[v]_\approx|$ bad components of $G - [v]_\approx$. ($[v]_\approx$ is finite per assumption) For the remaining components does the arc A_i have two endvertices r_i and w_i.

- Since bad components have only one root, each w_i has G-distance at least 2 from $[v]_\approx$. Hence $N_{G^3}(w_i) \setminus K_i \subseteq [v]_\approx \cup N_G([v]_\approx)$. It follows again that at most $2|[v]_\approx|$ of them have a vertex of $[v]_\approx$ as C-neighbor outside K_i and every other such vertex has a root of a component K_j as C-neighbor, from which almost all are bad.

We define an auxiliary graph Z with vertex set $\{K_{j+1}, K_{j+2}, ...\}$ and an edge $K_x K_y$, whenever the two components K_x and K_y satisfy (*) and there is a C-edge between the bad arcs A_x and A_y. Per definition has Z maximal degree 2 or less.

If Z contains a finite circle, the arcs A_i for the elements of this circle together with the C-edges between them induce a circle $C' \subsetneq C$, a contradiction. It follows that each component of Z is either a singleton, a finite path, a ray or a double-ray. Each component of Z which is a singleton or a finite path contains one element K_i which does not satisfy (*). Since this can only happen finitely often, it follows that Z has only finitely many such components, hence Z contains a ray.

The arcs according to (*) inside the components on this ray together with the C-edges between them induce an arc A in C. It is clear that A satisfies $(A1)$ and $(A2)$. Whenever we delete any initial segment of A, the remaining arc still contains infinitely many roots of not well-coverable components of $G - [v]_\approx$. It follows that the end induced by $[v]_\sim$ is in the closure of A but not as an inner point. Hence A ends in the end induced by $[v]_\sim$ and satisfies $(A3)$. □

Corollary 4.31. *Let A be a Hamilton arc of G^3 and $v \in V_\infty$. If a class $[v]_\approx \in [v]_\sim / \approx$ is finite and has only finitely many good neighbors, then C contains a captured arc in the not well-coverable components of $G - [v]_\approx$.*

Proof. The proof is analogous to the proof of 4.30 apart from the fact that there might be two fewer components satisfying the claim (*) (one for each endpoint of the Hamilton arc). □

Remark 4.32. In the context from 4.30, according to 4.25 only the well-coverable components of $G - [v]_\approx$ can meet another strong class of $[v]_\sim / \approx$, hence almost all bad components of $G - [v]_\approx$ are also components of $G - [v]_\sim$. So our captured arcs lie especially in the not well-coverable components of $G - [v]_\sim$, whose neighborhood is in $[v]_\approx$.

Definition 4.33. A strong class $[v]_\approx$ of $V_\infty / \approx$ is called *bad* if it is finite and has only finitely many good neighbors. A weak class $[v]_\sim$ of $V_\infty / \sim$ is called *splitting*, if $|[v]_\sim / \approx| \geq 3$ and $[v]_\sim / \approx$ contains either two or more bad classes or one bad class $[v]_\approx$ that separates $[v]_\sim$ in $G_\infty^\sim$.

Theorem 4.34. *If $V_\infty / \sim$ has a splitting class $[v]_\sim$, then G^3 has no Hamilton circle.*

Proof. Suppose for a contradiction that G^3 has a Hamilton circle C.

Let ω be the end induced by $[v]_\sim$ (see Proposition 4.22). If there are two bad classes $[k_1]_\approx$ and $[k_2]_\approx$ in $[v]_\sim / \approx$ with captured arcs R_1 and R_2 ending in ω (Lemma 4.30), the captured arcs have only finitely many vertices in common (else this would imply $k_1 \approx k_2$, a contradiction), so after deleting an initial segment we may assume that they are disjoint. After deleting R_1 and R_2 from C, it remains an arc $A := C \setminus (R_1 \cup R_2)$ between the initial vertices x_1 and x_2 from R_1 and R_2.

Because $|[v]_\sim / \approx| \geq 3$, there is at least one third strong class $[k_3]_\approx$. Per definition all vertices from a captured arc R_i lie in the not well-coverable components of $G - [k_i]_\approx$. Now Lemma 4.25 implies that none of these components meets $[k_3]_\approx$. Hence there are also infinitely many neighbors of $[k_3]_\approx$, which are not covered by R_1 and R_2. These neighbors must be contained in A. Since A is closed and ω lives in the boundary of $N_G([k_3]_\approx)$ it follows that ω is a boundary point of A, obtaining a third arc ending in ω, a contradiction.

If there is only one bad class $[k]_\approx$ in $[v]_\sim / \approx$, then it separates at least two vertices w and w' in $G^\sim_\infty$. If C contains two disjoint captured arcs induced by $[k]_\approx$, we obtain the same contradiction as before with the G-neighbors of $[w]_\approx$. If not, then there is only one captured arc R. We may assume that R contains all but finitely many components of $G - [k]_\approx$, else $G - R$ has a Hamilton arc $C - R$ and $[k]_\approx$ has also only finitely many good neighbors in $G - R$, so we would find another captured arc (Corollary 4.31). Let $\{r_1, ..., r_k\}$ be the set of all roots of the finitely many components from $G - [k]_\approx$ not covered by R .

Since $[k]_\approx$ separates w and w' in $G^\sim_\infty$, according to Lemma 4.15 there is a finite $w - w'-$separator S in $G - [k]_\approx$ with finite neighborhood. It follows that each $w - w'-$path in G^3 avoiding S contains an edge for which its endvertices are separated by $[k]_\approx$ in G. Thus each such path contains a vertex of $N_G([k]_\approx)$, so $Z := [k]_\approx \cup N_G([k]_\approx) \cup S \cup N_G(S)$ is a separator of w and w' in G^3. Then $Z - R = [k]_\approx \cup \{r_1, ..., r_k\} \cup S \cup N_G(S)$ is a finite separator of w and w' in $G^3 - R$. But after deleting R from C, it remains an arc A' between the initial vertex x of R and the end ω. This arc contains only finitely many arcs between $N_G(w)$ and $N_G(w')$ (each such arc uses one vertex from $Z - R$ or the end ω), and hence lies eventually in one of these components. But ω is in the closure of both components, so A' must have ω already as an inner point, a contradiction. □

4.2.3 Covering finite subgraphs

After we characterized counterexamples by splitting classes, it remains for us to show the converse that every countable rayless graph without a splitting class has indeed a Hamilton circle in its third power. We will do this by constructing Hamilton circles in several steps. A good way to start with our construction of Hamilton circles is to think about covering the vertex set of finite subgraphs. However we will later apply the results of this section not exactly on subgraphs of G itself, but subgraphs of contraction minors.

From [33] we already know that we can find for each finite graph G a Hamilton circle of G^3 containing any chosen prescribed edge of G. Since our application of this finite statement will be a recursive construction of infinite Hamilton circles, we need more control about how edges of G^3 lie in this circle. Consider a finite component of $G - [v]$ for any week or strong class $[v]$. Remember that its roots are defined as the vertices in the neighborhood of $[v]$.We may need to make sure that our contracted vertices are incident with edges of G^2 that can be replaced in a later step. Further we may need for each finite subgraph an edge of G incident with one of its roots.

This motivates the following definitions:

Definition 4.35. We say, a subspace X of G^3 *respects* a set of vertices V' of G, if there is an injection $i : V' \to E(G) \cap E(X)$ such that $v \in i(v)$ for all $v \in V'$. For such a given function, we say $i(v)$ *respects* v. Also whenever we use the phrase e *respects* v without defining a function i, we assume that we fixed one with $i(v) = e$.
We say, a subspace X of G^3 *pays attention to* a set of vertices V' of G, if there is an injection $i : V' \to E(G^2) \cap E(X)$ such that $v \in i(v)$ for all $v \in V'$. For such a given function, we say $i(v)$ *pays attention to* v. Also whenever we use the phrase e *pays attention to* v without defining a function i, we assume that we fixed one with $i(v) = e$.

Further, whenever we cover the vertices of a finite subgraph by a path, it is important to know which are the endvertices of this path or more precise, by how far they are away in G from a root of such a subgraph. We will see in this section that the best such distance possible depends on whether a finite subgraph is well-coverable or not. With this thought in mind, we are ready to prove several detailed ways of covering finite subgraphs, which can be utilized as a toolbox for constructing Hamilton circles and arcs. For the sake of simplicity, we will do some constructions for trees and then apply them to arbitrary finite connected graphs.

Lemma 4.36. *For any finite rooted tree (T, r) with at least 3 vertices and an independent set V' of vertices of T, there is a Hamilton circle C of T^3 which respects r and pays attention* *to V'. Moreover can we construct the circle such that both edges incident with r are in T^2.*

Proof. Let $T = (V, E)$ be a finite tree with root r and at least 3 vertices. We define an equivalence relation on the set of leaves of T, in which two leaves are equivalent whenever they have the same lower neighbor in T. We call the equivalence classes under this relation *bundles*. Define $L_0 := \{r\}$ and for each $i \in \mathbb{N}$ let L_i be the ith distance class of r. We say that a Hamilton circle C of T^3 covers a bundle with lower neighbor v *smoothly*, if there is an enumeration $\{l_0, l_1, ..., l_k\}$ of the bundle such that C contains the segment $vl_0l_1...l_k$ or its inverse. In that case, we can pay attention to the whole bundle with the following function:

$$i(l_m) = \begin{cases} vl_0 & \text{if } m = 0, \\ l_{m-1}l_m & \text{if } 0 < m \leq k \end{cases}$$

We apply induction after the height h of T and show that we find a Hamilton circle C of T^3 with the following slightly stronger properties:

(i) One of the edges in C incident with r is in T and the other is in T^2.

(ii) C pays attention to $(V' \cup L_h) \setminus \{l \in L_{h-1} | l \text{ is no leaf }\}$.

(iii) Each bundle of T is covered smoothly by C.

To see how (i)-(iii) imply the lemma, note that whenever a vertex v is in $V' \cap L_{h-1}$, none of its upper neighbors is in V' because V' is an independent set of vertices. So whenever this vertex is not payed attention to due to (ii), it has by property (iii) an upper neighbor l_0 with the edge vl_0 paying attention to it. We can use this edge to pay attention to v instead of l_0 and hence pay attention to V'.

For $h = 1$, let $v_0, v_1, ..., v_n$ be the neighbors of r. Let C be the Hamilton circle $rv_0v_1v_2...v_nr$. This circle clearly satisfies (iii) and the function

$$i(v_k) = \begin{cases} rv_0 & \text{if } k = 0, \\ v_{k-1}v_k & \text{if } k > 0 \end{cases}$$

witnesses that all vertices are payed attention to. Property (i) is satisfied due to the edge rv_0 if r is in V'.

Let now a tree T of height h be given together with an independent set of vertices V'. We define $T' := T - L_h$. Per inductive assumption there is a Hamilton circle C_{h-1} of T'^3 satisfying the properties (i)-(iii) for $V' \setminus L_h$ and $h - 1$.

Now for each vertex v in L_{h-2} that has upper upper neighbors, we modify C_{h-1} as follows: We name the upper neighbors of v in T as $l_0, l_1, ..., l_k$ and for each l_i we name the bundle of its upper neighbors as $l_{i,0}, l_{i,1}, ...l_{i,k_i}$. We will modify C_{h-1} to cover $\lfloor v \rfloor$ as follows:

Since we covered each bundle smoothly, we may assume without loss of generality that C_{h-1} contains the segment $vl_0l_1...l_k$.
We replace this segment by $vl_{0,0}l_{0,1}, ...l_{0,k_0}l_0l_{1,0}l_{1,1}, ...l_{1,k_1}l_1...l_{k,0}l_{k,1}, ...l_{k,k_k}l_k$, leaving out the bundle of leaves whenever a vertex l_i has no upper neighbors. The property (i) is clearly satisfied for the new circle. That (iii) is satisfied is easy to check since each bundle $l_{i_0}, l_{i_1}, ...l_{i_{k_i}}$ is covered by a path according to the definition of smoothly. Since we need no longer pay attention to vertices in the $h-1th$ level it is clear that all those new edges covering the bundles smoothly can be used to satisfy the part of (ii) that we shall pay attention to L_h. But the initial circle C_{h-1} did not pay attention to the vertices of V' in the $h-2th$ level. It remains to show that C_h does this. For a vertex $v \in (V' \cap L_{h-2})$ we can use the edge $vl_{0,0}$ to pay attention to it, since $l_{0,0}$ is payed attention to with $l_{0,0}l_{0,1}$ or $l_{0,0}l_0$ if its bundle is a singleton. □

Corollary 4.37. *For any finite rooted tree (T, r) with at least 3 vertices and an independent set V' of vertices of T, there is a Hamilton path of T^3 between r and and one T-neighbor of r which pays attention to $\{r\} \cup V'$*

Proof. Delete the edge respecting r from a Hamilton circle as in 4.36 □

Note that those statements about finite trees can also be applied for any finite graph, by using a spanning tree.

Corollary 4.38. *For any finite graph $G = (V, E)$ with at least 3 vertices, any vertex $r \in V$ and an independent set V' of vertices of G, there is a Hamilton circle C of G^3 which respects r and pays attention to V'.*

Corollary 4.39. *For any finite graph $G = (V, E)$ with at least 3 vertices, any vertex $r \in V$ and an independent set V' of vertices of G, there is a Hamilton path of G^3 between r and and one G-neighbor of r which pays attention to $\{r\} \cup V'$*

Let us remember how we obtained captured arcs around a strong class $[v]/\approx$ in our counterexamples. Whenever we cover the vertices of a not well-coverable component as in 4.38, we leave it through a vertex which is a G-neighbor of the root r. Whenever the next vertex is a root of another not well-coverable component, we end again in a vertex of G-distance 2 from $[v]/\approx$. Apart from some finite exceptions, the only way to leave the components around $[v]/\approx$ and reach another strong class is through well-coverable components.

Remark 4.40. A well-coverable component of $G - [v]$ for a strong or weak class $[v]$ has been defined as the negation of a not well-coverable component as defined in Definition 4.24. Thus a well-coverable component K has either more then one root, or of it has only one root r, it satisfies one of the following:

(i) $|V(K)| \leq 2$,

(ii) $d_K(r) = \infty$

(iii) $K - r$ has less then $d_K(r)$ components, or

(iv) one component of $K - r$ has only one vertex

To cover a well-coverable component, we will either use a path between two roots of one such component (see (C")) in the following lemma) or we will use only a single root. In the second case, we need to cover the rest of the well-coverable component with an arc between two neighbors of this root (see (C') in the following lemma).

As already mentioned before, paying attention to independent sets of vertices will play an essential role in our construction of Hamilton circles for infinite rayless graphs. However, when leaving out the root of a well-coverable component of $G - V_\infty$ and covering the rest, there will be one exception, in which we cannot pay attention to an arbitrary independent set of vertices. For this exception, we will define the property of being *sovereign* for independent vertex sets in well-coverable components as follows:

When K is well-coverable with only one root r and $K - r$ has exactly $d_K(r)$ components, it follows from the Definition 4.24 that there is at least one of those components of $K - r$ with only one vertex. (see the last property of Remark 4.40) In this case, we call an independent set not containing r *sovereign*, if there is one such component of size 1 which is also not part of that set. In every other case, we use *sovereign* as a synonym for independent.

Lemma 4.41. *Let $[v]$ be any strong or weak class in G and K be a finite component of $G - [v]$. Then K is well-coverable if and only if one of the following holds:*

(C') For every root r of K and any sovereign set $V' \subseteq V(K - r)$, there is a Hamilton path of $K^3 - r$ with endvertices of G-distance at most 2 from V_∞, which pays attention to V' or

(L) K has at most 2 vertices.

If K has at least two roots r and r', then also:

(C") For any sovereign set $V' \subseteq V(K-r)$, there exists a Hamilton path of K^3 with endvertices r and r' which pays attention to V'.

Proof. For the backward implication assume that K is not well-coverable. From Definition 4.24 we obtain that K has at least 3 vertices. The negation of (C') follows from Lemma 4.28.

Assume now that K is a finite well-coverable component of $G-[v]$ and (L) does not hold. Let $V' \subseteq V(K)$ be a sovereign set. In the case that K has more than one root r, let r' be another one of them. We will show (C') and then (C"):

To show (C'), we will construct a Hamilton-path of $K^3 - r$ with endvertices of G-distance at most 2 from $[v]$: Because K is finite, we can assume without loss of generality that K is a tree. Also does K contain a $r' - r-$path $P' = x_0x_1x_2...x_nx_{n+1}$ (with $x_0 = r'$ and $x_{n+1} = r$). For each vertex x_i on this path, let X_i be the subgraph of K induced by x_i and all components of $K - x_i$, which are disjoint from P'. For each X_i with $i \leq n$, which has size more than 1, we will replace the edge x_ix_{i+1} of P' by a path covering $V(X_i)$ and paying attention to all vertices of $X_i \cap V'$. If $X_i = \{x_i, w_i\}$, then define $P_i := x_iw_ix_{i+1}$. Clearly the edge x_iw_i pays attention to w_i.

If $|X_i| > 2$, then from 4.38 we obtain a Hamilton circle C_i of X_i^3, which respects x_i with an edge x_iw_i and pays attention to $X_i \cap V'$. In this case, we define $P_i := C_i - x_iw_i + w_ix_{i+1}$ and $P' := x_0P_0w_0w_1P_1x_1w_2P_2x_2w_3...x_{n-1}w_nP_nx_n$. The path P' covers all vertices from the component of $K - r$, which contains r' and pays attention to all its vertices in V'. If $K - r$ has no other component, then we are done. If not, let X' be the union of r and the remaining components. According to Lemma 4.39, there is a Hamilton path H of X'^3 between r and one G-neighbor w of r, which pays attention to $\{r\} \cup V' \cap X'$ such that the H-neighbor w' of r has distance at most 2 from r in X'. Hence there is an edge x_nw' in K and we can define $P := x_0P'x_nw'Hw$. This path covers all vertices of $K - r$, pays attention to all leaves and its endvertices have G-distance at most 2 from $[v]$.
To obtain the Hamilton path as in (C"), add the edge wr to the constructed path P.

In the case that K has only one root r, it follows from Definition 4.24 that there is one component X of $G - r$ which contains at least two neighbors of r (see (iii) in Remark 4.40) or has only one vertex (see (iv) in Remark 4.40).

Let $K_1, K_2, ..., K_n$ be the other components of $G - r$. Again, we apply Lemma 4.39 to find for each K_i^3 a Hamilton-path Q_i with first endvertex of K-distance 1 and second endvertex of K-distance 2 from r, which pays attention to $K_i \cap V'$. Consider the union of the Q_i

together with all edges between the second endvertex of each Q_i and the first endvertex of Q_{i+1} (whenever Q_{i+1} exists). Let v be one G-neighbor of r in X and add the edge from the second endvertex of Q_n to v to the required path. If X has more than one vertex, it has another G-neighbor v' of r. In this case, we can apply (C") to find a path covering $V(X^3)$ from v to v', which pays attention to $X \cap V'$. We add this path and finish our construction. $\square$

4.2.4 Rayless graphs with a single weak class

In this whole section, let $G = (V, E)$ be a rayless graph for which $|V_\infty / \sim| = 1$ and for which its one weak class is not splitting. To understand this section better, it might be helpful to keep in mind that this graph will arise later as a contraction minor from some original graph with possibly multiple weak classes.

We call a strong class $[v]_\approx$ *good*, if it is not bad. This is exactly the case, whenever it is infinite or if it has infinitely many good neighbors. Let $\mathcal{C}$ be the set of components of $G - [v]_\sim = G - V_\infty$. We call its elements *flaps*.

Note that each flap is finite, since it contains only vertices of finite degree and because G is rayless per assumption. We will construct a Hamilton circle and a Hamilton arc of G^3. Since G^3 is one ended according to 4.22, it remains to build a spanning double ray in the first case and two disjoint rays spanning V in the second case. To do so, we apply the results of the last section to the elements of $\mathcal{C}$. For our main result later, the results of this section will be applied to certain contraction minors. We will now carefully define a specific setup to obtain some additional properties.

Some flaps that might be not well-coverable in the original graph might become well-coverable after certain contractions. We will *pretend* as if they were not well-coverable. Also if there is some bad class $[v]_\approx$, we will not use the property of being well-coverable for flaps adjacent to that class either. To be able to keep track of that, let us assume that we fixed a colouring of $V_\infty \cup N_G(V_\infty)$ into the colors red and green satisfying the following:

(C1) Every vertex of $[v]_\sim$ is green. Every other green vertex is in a well-coverable flap,

(C2) Every well-coverable flap contains at most one green root,

(C3) If $|[v]_\sim / \approx| \geq 3$, then every class $[v]_\approx$ apart from at most one not separating $G_\infty^\approx$ does either contain infinitely many green vertices or has infinitely many green vertices in its neighborhood and

(C4) Every vertex $v \in [v]_\sim$ has either infinitely many green neighbors in $G - [v]_\sim$ or none.

We will see in the next section in the proof of our main result 4.55 that we can always obtain such a coloring. For now, we just assume that the coloring is given.

Note that the third property (C3) is equivalent to the statement that V_∞ remains not splitting if we change all well-coverable flaps without a green root into not well-coverable flaps.

Every root of a not well-coverable flap is red. Together with the first property, we obtain that a strong class $[v']_\approx$ is good if and only if that class together with its G-neighborhood has infinitely many green vertices.

For each flap with only red roots, we cover it in a way according to 4.39 (so, we do not care, whether it is well-coverable or not). Components with a green root are always well-coverable and we might cover them with a path omitting the green root as in 4.41.

We define a *module* as the vertex set of a whole flap or as the vertex set of a flap without a green vertex. We sometimes use the word module also for the subgraph of G induced by the vertex set of a module. Furthermore, we will show that we can again choose one vertex that our Hamilton circle or arc respects and one independent set of vertices to which we pay attention for each flap. As mentioned before, there is one exception for how we can choose those independent sets: Whenever we want to cover a one-rooted well-coverable flap and omitting the root as in 4.41, we can only pay attention to a sovereign set as defined in the last section. We need this kind of cover for those flaps with a green root.

For all other flaps we do not need the information whether they are well-coverable or not and cover them in a way according to 4.39. So we assume for the whole section that we fixed for every flap an independent set of vertices inside that flap which is sovereign for the flaps containing one green root and arbitrarily chosen for all other flaps. We call a standard subspace of $|G^3|$ *sovereign* if it pays attention to the chosen independent sets of each module which it covers all vertices.

Our strategy to construct a Hamilton circle is as follows: We will split G into two parts and construct a spanning ray of the third power of each part. We will make sure that in each part, we can move around to cover modules and green vertices in arbitrary order without leaving any vertices uncovered. To do so, we define the relation of being connectable on V_∞. This relation basically states that for an equivalence class $V' \subseteq V_\infty$, after we already used finitely many vertices to construct an initial segment of a ray, we can still extend any finite path constructed so far by connecting an arbitrary vertex of V' or a flap adjacent to V'. Using this property, we will be able to build a spanning ray or double ray recursively for each such class V'.

Definition 4.42. A sovereign path P in G^3 from a vertex v' to a vertex w' is called *structured*, if it satisfies the following:

(S1) The endvertices v' and w' of P both have G-distance at most 2 from V_∞,

(S2) $V(P)$ is the union of v', modules and green vertices, (S3) if $v' \notin V_\infty$, then P meets the flap of v' only in v' and

(S4) if $w' \notin V_\infty$, then P contains a module containing w'.

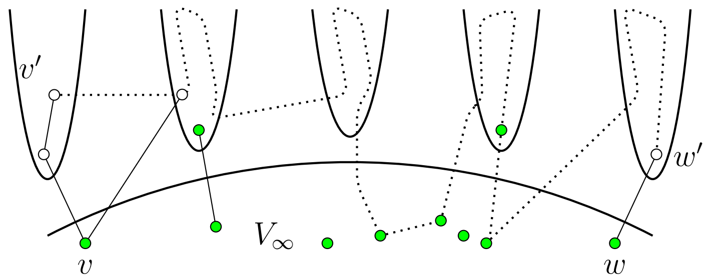

Figure 7: A structured path from v' to w'

Our Hamilton circle and arc will be the union of infinitely many such structured paths. To keep track of the already used parts of the graph, we define a *blocker* as any finite vertex set which is the union of a set of green vertices and vertex sets of whole modules. Note that each module has finitely many vertices, so we do not have to worry about blockers becoming infinite by adding finitely many modules.

Definition 4.43. We say that a vertex $v \in V_\infty$ is *connectable* to $w \in V_\infty$, if given any blocker B, any vertex $v' \notin B$ with G-distance at most 2 from v and any *target* avoiding B, which is either w itself or a flap $K \in \mathcal{C}$ with w as G-neighbor, there is a structured path P satisfying the following:

(S5) P begins in v',

(S6) P ends in w, if w is the target and otherwise in a vertex $w' \in K$ with G-distance at most 2 from w,

(S7) P avoids B and

(S8) if $v = v'$, then the first edge of P lies in G.

In the case $v = v' = w = w'$, we define the singleton vertex v also as a structured path.

Lemma 4.44. *'Being connectable to' is a reflexive and transitive relation on V_∞.*

Proof. We show first that every element $v \in V_\infty$ is connectable to itself: If our target is v itself, then the desired structured path is either an edge from v' to v or a single vertex, if $v = v'$. As structured path from v' to a given flap $K \in \mathcal{C}$, we can take a cover of K as in 4.39 together with an edge in G^3 from v' to a root of K, which is a G-neighbor of v. That way is (S8) satisfied. The properties (S1)-(S7) are straightforward to check.
Let now v be connectable to w and w be connectable to x. We will show that v is connectable to x: We assume that we are given any blocker B, any vertex v' avoiding B with G-distance at most 2 from v and any target that is either x or a flap K avoiding B with x as G-neighbor. Because v is connectable to w, there is a structured path P from v' to w satisfying (S5)-(S8). Because of (S2) this path consists of v', modules and green vertices. We obtain another blocker B_+ by adding these green vertices and modules together with the flap containing v' to B. Since modules are finite, B_+ remains finite as well. Because w is connectable to x, there is another structured path P satisfying (S5)-(S8), beginning in w and ending in a vertex $x' \in V(K)$ with G-distance at most 2 from x or in x itself, if x is the target. For the union of these paths are the properties (S1)-(S8) immediately clear per definition, hence we conclude that v is connectable to x. □

Definition 4.45. We say v and w *are* connectable, whenever v is connectable to w and w is connectable to v.

The idea behind the concept of being connectable is that we obtain equivalence classes in which we can move back and forth any number of times always using only finitely many modules and green vertices and hence still have infinitely many left to continue the construction recursively. With the setup defined, this construction will be fairly easy if we have only one single equivalence class of the relation of being connectable. Also for two classes, we will not have a problem, because we can construct two rays to our new end, one in each class.

The following two lemmas give sufficient conditions for vertices from V_∞ to be connectable. First, we show that we are always connectable inside strong classes. And second we show that we can always move out of good strong classes. In particular, when two adjacent classes are both good, then they lie in the same equivalence class of the relation of being connectable.

Lemma 4.46. *For every edge $vw \in E(G^{\approx}_\infty)$, its endvertices v and w are connectable.*

Proof. Assume that we are given any blocker B, any vertex v' avoiding B with G-distance at most 2 from v and any target which is either a flap K avoiding B with w as G-neighbor or w itself.

Per definition of $\approx$ there are infinitely many $v-w-$paths in G. Let P be one of them, such that no vertex of P lies in B and no vertex of P lies in a flap that contains a vertex of B. This is possible, since B is finite. For every vertex of finite degree on P, we add the whole flap containing it and obtain a graph P_+ that also avoids B. Now P_+ contains a sequence in $V_\infty \cup \mathcal{C}$ of pairwise different elements, from which each element is adjacent in G to the next one.

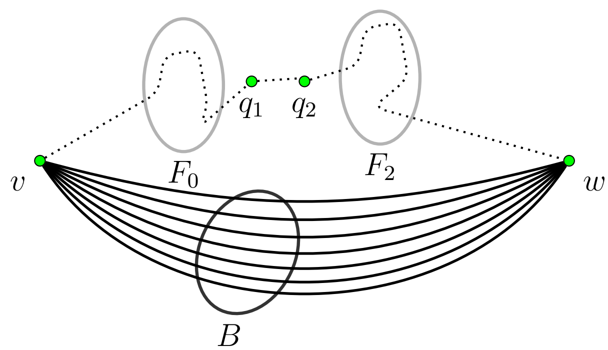

Figure 8: A structured path between two vertices v, w for which $vw \in E(G_\infty^{\approx})$. The Blocker B is finite, so we can find a $v-w-$path avoiding B.

Let $v = q_0, ..., q_l = w$ be the elements of V_∞ of this sequence in order. Between each two of them is either an edge in G or a flap.

We will now construct a structured path witnessing that v and w are connectable along this sequence. At first, we consider the case that $v = v'$ and the target is w.

We will construct a path between each two successive elements q_j and q_{j+1}: If there is no flap between q_j and q_{j+1}, we just use the edge q_jq_{j+1}.

Now assume that there is a flap F_j between q_j and q_{j+1}. Per assumption, both of q_j and q_{j+1} have one G-neighbor in F_j. If q_j and q_{j+1} have no such G-neighbor in F_j in common, then F_j has at least two roots and thus is well-coverable per Definition 4.24, so we can find a sovereign Hamilton-path P_{F_j} of F_j^3 between a G-neighbor of q_j and a G-neighbor of q_{j+1} (4.41). Otherwise, if q_j and q_{j+1} have one G-neighbor r_j in F_j in common, then we can find a sovereign Hamilton-path P_{F_j} of F_j^3 between r_j and one G-neighbor of r_j (4.39).

In both cases P_{F_j} is a sovereign path between a G-neighbor r_{F_j} of q_j and a vertex w_{F_j} which has distance at most 2 of q_{j+1} in G. So we add the edges $q_jr_{F_j}$ and $w_{F_j}q_{j+1}$ to obtain the desired path.

We show that the union of these paths satisfies (S1)-(S8): The properties (S1) and (S3)-(S7) are immediately clear, and (S2) holds, since each vertex in V_∞ is green and the vertex set of each flap is a module. Further is the first edge vq_2 or vr_{F_1}. In both cases is this an edge from G itself.

Now if $v' \neq v$, we can switch the first vertex to v'. Since v' has G-distance of at most 2 from v is the new first edge still in G^3.

If the target is a flap K instead of the vertex w, we can find a sovereign Hamilton-path P of K^3 between a G-neighbor r of w and one vertex w' with G-distance at most 2 from w. In the case that the target was w, the last edge on our path was in G^2. Thus we can replace w by r and add P, so we end up in w' to satisfy (S4) and (S6). The remaining properties are clear per construction. $\square$

Corollary 4.47. *Any two strongly equivalent vertices $v, w \in V_\infty$ are connectable.*

Proof. Combine Lemma 4.46 and Lemma 4.44. $\square$

Lemma 4.48. *For an edge $vw \in G_\infty^\sim$, for which $[v]_\approx$ is good, it follows that v is connectable to w.*

Proof. If $v \approx w$, then the statement is already shown in 4.47. Thus, we only have to consider the case that there is an edge between v and w in G.

Assume that we are given any blocker B, any vertex v' avoiding B with G-distance at most 2 from v and any target which is either a flap K avoiding B with w as G-neighbor or w itself.

We will construct the structured path in several steps. In each step, we might continue constructing from both sides, every time reducing the quest of constructing the remaining path to the construction of a new path between the two new endvertices. The properties (S1),(S3),(S4),(S5), (S6) and (S8) are satisfied at the beginning of our construction. Thus, in later steps we only have to check on (S2) and (S7).

Every time we construct something, we will avoid B to satisfy (S7). For simplicity, we assume without stating it in every single step that we added everything constructed so far to B and avoid it as well. This can be done since we use only a finite set of vertices and is necessary to obtain indeed a path in the end.

If the target is w, then the edge $v'w$ is the desired path and the properties (S1)-(S8) are easy to check. So we may assume that the target is a flap K.

Again, we can cover K by a path between a root r and a vertex w' of G-distance at most 2 from w (4.39). This path will be the end of our structured path, so it remains to construct the initial segment from v' to r.

If $v' = v$, choose any neighbor g_1 of v avoiding B. In this case, we start our path with the edge vg_1 to satisfy (S8). If g_1 is green, we redefine $v' := g_1$ and continue. If g_1 is red, it is part of a flap avoiding B and we continue our path with a path covering the third power of this flap (4.39) and redefine v' as the endvertex of this path.

Thus it remains to find a structured path from $v' \neq v$ to r.

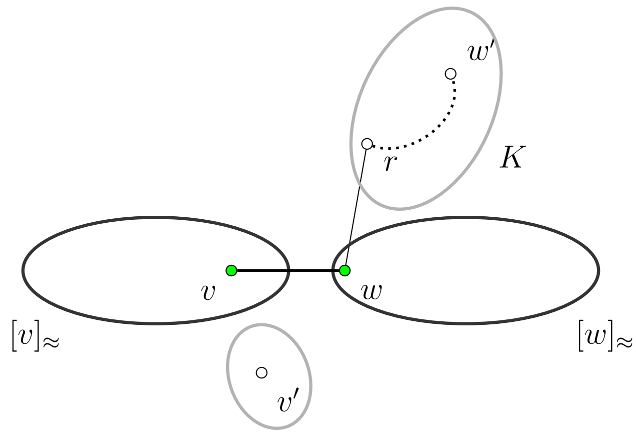

If v' has G-distance less then 2 from v, we can just use the edge $v'r$ and are done. So we may assume that v' has distance 2 from v and it remains to construct a structured path from v' to a G-neighbor r_1 of v instead. Choose $r_1 \notin B$ as an arbitrary G-neighbor of v. If r_1 is green, we can just use the edge $v'r_1$ as structured path and we are done with the construction. If r_1 is red, it is part of a flap K_1 also avoiding B. The case that this flap has no other element then r_1 is analogous to the case that r_1 is green. Thus we can assume that we cover K_1^3 according to (4.39) by a path between r_1 and a vertex w_1 of G-distance at most 2 from v. Now it remains to connect v' and w_1 by a structured path.

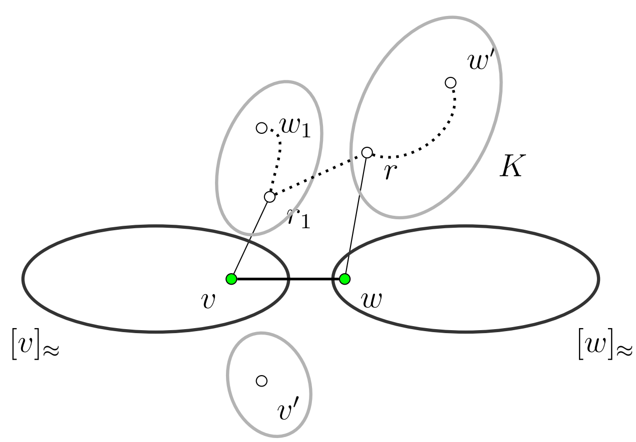

Since $[v]_\approx$ is good, there are infinitely many green vertices in $[v]_\approx \cup N_G([v]_\approx)$ left avoiding

B, let d be one of them. We define $x := d$, if $d \in [v]_\approx$. Otherwise let x be a neighbor of d in $[v]_\approx$. We know from Lemma 4.47 that v is connectable to x. Thus we can find structured paths P' from v' to a vertex x' of G-distance at most 2 from x and a structured paths P_1 from v_1 to a vertex x_1 of G-distance at most 2 from x. We choose those paths disjoint from each other (possible since we can add P_1 to our blocker before we find P') and avoiding B.

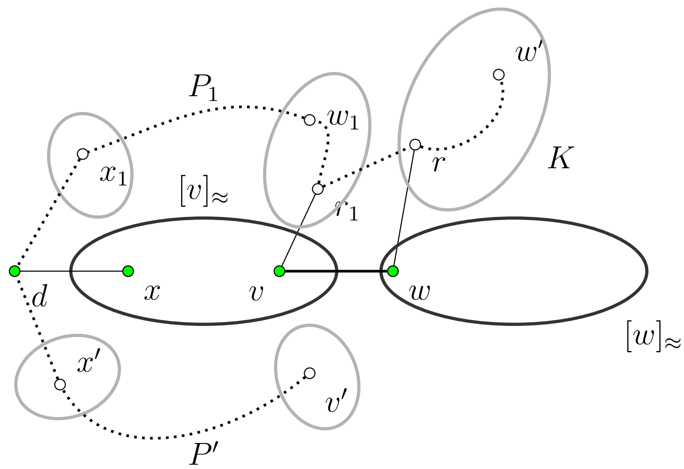

Now the endvertices x' and x_1 are both G^3-neighbors of d, since $d \in \{x\} \cup N_G(x)$. So we can connect them by adding the edges $x_1 d$ and dx' and are done with our construction. □

Remark 4.49. 'Being connectable to' is not necessary a symmetric relation. If w from the last lemma lies in a bad class and we forbid all well-coverable flaps adjacent to that class, any further structured path follows a captured arc as shown in 4.30, so we cannot connect w back to v.

Corollary 4.50. *For an edge $vw \in G^{\sim}_\infty$, for which $[v]_\approx$ and $[w]_\approx$ are good, it follows that any $v' \in [v]_\approx$ and $w' \in [w]_\approx$ are connectable.*

Consider the case that all vertices of V_∞ are connectable. The following three lemmas all use the same technique of constructing spanning paths inside a graph, by combining a sequence of structured paths using the property of being connectable. Since we can move on arbitrarily around vertices of V_∞, we can control whether we want to construct a Hamilton circle, a Hamilton ray or a Hamilton arc between two arbitrarily chosen vertices.

Lemma 4.51. *If every two vertices in V_∞ are connectable and $V_0 := \{v_0, v_1, ..., v_n\}$ is a finite subset of V_∞, then G^3 has a sovereign spanning ray starting in v_0 and respecting V_0.*

Proof. We enumerate $V_\infty \cup \mathcal{C} = \{v_0, v_1, ..., v_n, y_{n+1}, y_{n+2}, ...\}$ with $v_0, v_1, ..., v_n$ as the first elements.

Beginning with $P_0 := v_0$, we will construct a sequence of finite sovereign paths $P_1, P_2, P_3, ...$ such that each path P_i satisfies:

(i) $V(P_i)$ is the union of modules and green vertices,

(ii) the first endvertex of P_i is v_0 and the second endvertex of P_i has G-distance at most 2 from V_∞,

(iii) $P_{i-1} \subseteq P_i$,

(iv) if $i \leq n$, then P_i contains $v_0, v_1, ..., v_i$ in order, ends in v_i and respects $v_1, v_2, ..., v_{i-1}$. Further does P_i not contain $v_{i+1}, v_{i+2}, ..., v_n$ and

(v) if $i > n$, then P_i contains and respects V_0 and covers y_j for all $j \leq i+1$ (if y_j is a flap, this means that P_i contains all vertices of y_j).

We define $P_{-1} := P_0$ to satisfy (iii) for P_0. The other properties are clearly satisfied. Let P_i be given satisfying (i)-(v) with endvertices v_0 and w_i (We define $w_0 := v_0$).

Let B' be the set of vertices from P_i. It follows from (i) that B' is indeed a blocker. If $i \leq n$, we define $B := B' \cup V_0 \setminus \{v_{i+1}\}$. If $i > n$, we define $B := B'$. Note that in this case it follows from (v) that B already contains V_0.

If the next element x_{i+1} is already covered by P_i, then we define $P_{i+1} := P_i$. So we assume that x_{i+1} is not covered by P_i. If x_{i+1} (either v_{i+1} or y_{i+1}) is a flap, define $K := x_{i+1}$ and choose any vertex w in V_∞ such that K meets $N_G(w)$. If $x_{i+1} \in V_\infty$, define $w := x_{i+1}$ and K as any flap avoiding B and meeting $N_G(w)$.

If K avoids B, then per assumption there is a structured path P_+ beginning in w_i and ending in a vertex $w'(K)$ with G-distance at most 2 from w which satisfies (S5)-(S8).

It follows from $(S2)$ that the only possible case in which K does not avoid B is if B contains one green vertex d in K. Now we obtain from $(C4)$ that w has infinitely many green neighbors. Let d' be one of them avoiding B. We obtain P_+ in the same way as in the case when K avoids B apart from that we exchange d for d' in this path. To see why this exchange is possible, note that $(S2)$ implies that the structured path ends in a Hamilton path of the third power of the target K. In this case is K well-coverable, so we can assume that we cover $K - d$ according to (C') in Lemma 4.41. The endvertices of this cover have G-distance of at most 2 from w and thus G-distance at most 3 from d'. Hence we can cover $K - d + d'$ instead of K.

We define P_{i+1} as P_i extended by P_+ and the edge $w'w$ if $w = x_{i+1}$. This is indeed a path, because P_i and P_+ are disjoint from each other apart from w_i. Since any structured path is sovereign, we keep this property as well for any P_{i+1}. The properties (i)-(iii) are easy to check. If $i \leq n$, then since P_i satisfies (iv) and $V_0 \setminus \{v_{i+1}\}$ is part of B, we added no other vertex from V_0 but v_{i+1} to our path. Further follows from (S8) that we respect v_i. The vertices $v_0, v_1, ..., v_{i-1}$ are already respected by P_i and hence also by P_{i+1}. Thus (iv) is satisfied for P_{i+1}.

If $i > n$, then with the same argument does P_{n+1} already contain and respect V_0, hence the same holds for P_{i+1}. It is clear by construction that P_{i+1} covers y_{i+1} and hence by induction also all earlier elements of our enumeration, thus (v) holds as well for P_{i+1}.

Clearly the ray $R := \bigcup_{i \in \mathbb{N}} P_i$ is a sovereign spanning ray of G^3 starting in v_0 and respecting V_0. $\square$

Lemma 4.52. *If every two vertices in V_∞ are connectable and V_0 is an arbitrary finite subset of V_∞, then G^3 has a sovereign spanning double-ray respecting V_0.*

Proof. The proof is analogous to the proof of 4.51 apart from that we change (i) to "Both endvertices of P_i have G-distance at most 2 from V_∞" and extend the path alternately on both sides. $\square$

Lemma 4.53. *If every two vertices in V_∞ are connectable and $v_0, w_0 \in V_\infty$, then G^3 has a sovereign Hamilton arc between v_0 and w_0 respecting $\{v_0, w_0\}$.*

Proof. The proof is analogous to the proof of 4.51 apart from that we construct two disjoint such rays by starting at the two endvertices and extend alternately on both sides. $\square$

Proposition 4.54. *In the setup of this section does the following hold:*

(i) For every $v_0 \in V_\infty$, the graph G^3 has a sovereign Hamilton circle respecting v_0.

(ii) For every $w_1, w_2 \in V_\infty$, the graph G^3 has a sovereign Hamilton arc between w_1 and w_2 respecting w_1.

Proof. The graph G^3 has only one end and a Hamilton circle is a spanning double ray together with this end, while a Hamilton arc consists of two disjoint rays spanning $V(G)$ together with this end.

If $V_\infty/\approx$ has only one class, then there is nothing left to show (see 4.52 and 4.53).

So we may assume that $|V_\infty/\approx| \geq 2$. For both cases, we will construct two disjoint rays covering $V(G^3)$. For case (i), we will make sure that both rays start in adjacent vertices w_1 and w_2 such that we can add the edge $w_1 w_2$ to obtain a Hamilton circle. For case (ii), we will make sure that the two rays start in the given vertices w_1 and w_2.

Since V_∞ is not splitting, there is at most one class $[v_1]_\approx \in V_\infty/\approx$ with only finitely many good neighbors in G. If no such class exists, define $[v_1]_\approx$ as an arbitrary class in $V_\infty/\approx$ which is not separating $G_\infty^\approx$. (such a class exists, because $G_\infty^\approx$ is also rayless and hence a spanning tree of $G_\infty^\approx$ contains a leaf.)

We may assume that v_1 is chosen in a way such that v_1 has a neighbor $v_2 \in V_\infty \setminus [v_1]_\approx$.

Let $\mathcal{C}_0$ be the flaps in $\mathcal{C}$ with a G-neighbor in $[v_1]_\approx$ and $\mathcal{C}_1$ all other flaps. Define $G_0 := G[[v_1]_\approx \cup \mathcal{C}_0]$ and $G_1 := G \setminus V(G_0)$.

For every other strong class $[v']_\approx$, there are only finitely many flaps which have a neighbor in $[v_1]_\approx$ as well as in $[v']_\approx$. Thus in G_1 still every strong class is good. It follows from 4.47 that every two vertices in $[v_1]_\approx$ are connectable in G_0. Since every strong class in G_1 is good, it follows from 4.50 that every two vertices in $V_\infty \setminus [v_1]_\approx$ are connectable in G_1.

Now Lemma 4.51 implies that we obtain sovereign spanning rays of G_0^3 and G_1^3.

To prove (i), we choose v_1 and v_2 as their starting vertices and add the edge v_1v_2 to obtain the desired Hamilton circle. Further using the same lemma, we can make sure to respect v_0.

To prove (ii), we choose as starting vertices of the rays w_1 and w_2 and obtains the desired statement directly if one w_1 and w_2 is in G_0 and the other one is in G_1.

If not, we assume without loss of generality that $w_1, w_2 \in G_0$. Now let R_1 be a spanning ray of G_1^3 beginning in v_2 and R_0 be a spanning ray of $G_0^3 + w_2w_1$ beginning in v_1 and containing and respecting $\{v_1, w_2, w_1\}$ in the order v_1, w_2, w_1. We may assume that we used the new added edge w_2w_1 as the $w_2 - w_1-$path in the construction of the ray R_0 according to Lemma 4.51.

Now the two rays $w_2R_1v_1v_2R_2$ and w_1R_1 together with the end of G^3 are the desired sovereign Hamilton arc between w_1 and w_2 respecting w_1. □

4.2.5 Main Result

Finally we are ready to finish the proof of our main result that a countable rayless graph has a Hamilton circle in its third power if and only if no class of $V_\infty/\sim$ is splitting.

Theorem 4.55. *For a rayless graph G, its third power G^3 has a Hamilton circle if and only if no class of $V_\infty/\sim$ is splitting.*

Proof. The forward direction has been shown in 4.34.
For the backward implication, we will construct a Hamilton circle. To make this construction easier to understand, we will possibly delete some unnecessary edges of G without changing the end space of G^3.

Let us call vertices of infinite degree *big* and vertices of finite degree *small*. A *small* path is a path in G consisting of only small vertices. A *small component* is defined accordingly. Since G is rayless, each small component of G is finite.

We construct a sequence of subgraphs $G_0 \supseteq G_1 \supseteq G_2, ...$ of G, each of them with vertex set $V(G)$ (but possibly less edges then G) and minors $\bar{G}_0, \bar{G}_1, \bar{G}_2, ...$ of G such that each $\bar{G}_i$ is obtained from G_i by some contractions. Let D_i be the vertices from $\bar{G}_i$, which are obtained by contraction of a set of vertices. We call them *dummy-vertices*. It may happen in the construction that we 'contract' a single vertex. For technical reasons we still think of that singleton vertex-set as a dummy-vertex. Every dummy-vertex is a subset of $V(G)$.

Let D_i be the set of dummy-vertices of $\bar{G}_i$ and V_i be the set of vertices of G_i, which are no dummy-vertices. Every vertex in V_i is a vertex of G.

Further we construct a Hamilton circle C_i for each $\bar{G}_i^3$ and satisfy the following:

(G1) $\Omega(G_i^3) = \Omega(|G^3|)$.

(G2) Each weak class of G_i is splitting.

(M1) For every $v \in V$, there is one $i \in \mathbb{N}$ such that $v \in V_i$.

(M2) $V_i \subseteq V_{i+1}$ for all $i \in \mathbb{N}$

(M3) C_i pays attention to D_i.

(M4) C_{i+1} contains all edges from C_i between vertices of V_i.

The statement $(G1)$ means that edges we delete in our construction will not change the endspace in the third power. For the first step of our construction, we pick any weak class $[v]_\sim$ and put all its vertices in V_0.

Let $\bar{G}$ be the graph obtained from G after identifying the vertices of each other weak class (note that a weak class induces not necessarily a connected subgraph, so this identification is not always a contraction). Note that all roots of components of $G - [v]_\sim$ are small vertices, hence none of them are identified yet. Especially whenever a component of $G - [v]_\sim$ has more then one root, it still has after this identification step. Further it is easy to check that also all of the properties (i)-(iv) from Remark 4.40 are preserved under identifications. Hence each well-coverable component of $G - [v]_\sim$ is also well-coverable in $\bar{G} - [v]_\sim$.

Now we claim for each component K of $\bar{G} - [v]_\sim$ the following:

Claim 1. *There is a finite subgraph $Y_K \subseteq K$ such that each neighbor of Y_K in $K - Y_K$ is a small vertex with only big neighbors in Y_K. Further does this subgraph preserve the property of being well-coverable whenever K was well-coverable.*

To see this claim, note that in K, each big vertex comes from a weak class of K and also is K finitely separable because no two big vertices are strongly equivalent (since the would be identified in that case).

In the case that K has more than one root, choose two roots r and r'. In this case it is clearly well-coverable (see 4.40). Let X be an $r - r'$−path inside of K and X^+ be the union of X together with all small components from small vertices of X. We define Y_K as X^+ together with all big neighbors of X^+ in K.

If K has only one root r and comes from a not well-coverable component of $G - [v]_\sim$, then we add the small component R^+ of r to V_0. Now R^+ has finitely many big neighbors. In this case, we define $Y_K := R^+ \cup N_K(R^+)$.

If K has only one root r and comes from a well-coverable component of $G - [v]_\sim$, then again we add the small component R^+ of r to V_0. Define $Y'_K := R^+ \cup N_K(R^+)$. If $Y'_K - r$ has less then $d_{Y'_K}(r)$ components, it is well-coverable and we define $Y_K := Y'_K$.

If not, then according to 4.40 $Y'_K - r$ has exactly $d_{Y'_K}(r)$ components. If one of these components has size 1 and has no neighbors in $K - Y'_K$, we define again $Y_K := Y'_K$

If not, then since K was well-coverable, there are two components of $Y'_K - r$ living in the same component of $K - Y'_K$. We choose any Y'_K-path P between two components of $Y'_K - r$ with internal vertices in $K - Y'_K$. We define P^+ as the union of P together with all small components from small vertices of P and $Y_K := Y'_K \cup P^+$.

This completes the proof of Claim 1.

Claim 2. *For each component K of $\bar{G} - [v]_\sim$, there is a minor of K isomorphic to Y_K, such that the coresponding dummy vertex d_b for each big vertex b in Y_K contains almost all* *components of $K - Y_K$ that are adjacent to b in K.*

To see this claim, choose for each component K' of $K - Y_K$ one big neighbor b of K and then delete every edge between K' and $Y_K - b$.

Since K is finitely separable, there are only finitely many components of $K - Y$ which have more than one neighbor on Y, so we only deleted finitely many edges. Now each component K' of $K - Y$ has one unique corresponding big neighbor b left. We add to each big vertex b in K all those corresponding components and obtain the dummy-vertex d_b. Since after the deletion of edges, each big vertex $b \in V(K)$ had still almost all of its initial neighbors, Claim 2 is forfilled.

Let G_0 be defined as G after deleting the same set of edges and $\bar{G}_0$ be the graph obtained from $\bar{G}$ after these deletions and contractions above for each component of $\bar{G} - [v]_\sim$. Since we deleted no edges between big vertices and for each big vertex only finitely many edges adjacent to it, it is easy to check that the relation $\sim$ does not change after the deletion and thus according to 4.22 does $(G1)$ hold for G_0. For each weak class $[w]_\sim$ of G_0, only finitely many of the components of $G - [w]_\sim$ might have changed after the deletion of edges, so it is clear that $[w]_\sim$ is still splitting in G_0 and $(G2)$ holds for G_0.

When K is well-coverable with only one root r and $K - r$ has exactly $d_K(r)$ components, it follows from the Definition 4.24 that there is at least one of those components of $K - r$ with only one vertex. (see the last property of Remark 4.40). In this case, we call an independent set not containing r *sovereign*, if there is one such component of size 1 which is also not part of that set. In every other case, we use *sovereign* as a synonym for independent.

It is clear that the set of dummy-vertices are an independent set of vertices in $\bar{G}_0$. Further if a component K came from a well-coverable component of $G - [v]_\sim$ and has only one root r such that $K - r$ has exactly $d_K(r)$ components, it follows from the Definition 4.24 that there is at least one of those components of $K - r$ with only one vertex. (see the last property of Remark 4.40) One of these vertices cannot be a dummy-vertex, because else it would have had neighbors in G, which would imply that the component K was not well-coverable before contraction. Hence we obtain that the set of dummy-vertices is indeed a sovereign set of vertices in $\bar{G}_0$.

Claim 3. *There is a coloring of $[v]_\sim \cup N_{\bar{G}_0}([v]_\sim)$ satisfying the properties (C1)-(C4) of the last section. Further is each green vertex of $N_{\bar{G}_0}([v]_\sim)$ in a well-coverable component of $G - [v]_\sim$*

To see this claim, we start by coloring all vertices of $[v]_\sim$ green. Now for each component K of $\bar{G} - [v]_\sim$, we apply the following: If K has more then one root, we color one of its roots green and the other roots red. If K has only one root and comes from a not well-coverable component of $G - [v]_\sim$, then we color its root red. If K has only one root and comes from a well-coverable component of $G - [v]_\sim$, then we color its root green.

It is easy to see that this coloring satisfies the properties (C1), (C2) and (C3). In a second step, we change our coloring to satisfy the property (C4) as well: We will keep all red vertices in $[v]_\sim \cup N_{\bar{G}_0}([v]_\sim)$ and all green vertices in $[v]_\sim$, but whenever a vertex in $[v]_\sim$ has only finitely many green neighbors in $\bar{G}_0 - [v]_\sim$, we color those neighbors red instead, to make sure that (C4) holds. It is clear that (C1), (C2) hold after the color change. Since $[v]_\sim$ is not splitting per assumption, (C3) was true before this color change. However if a class $[v']_\approx \subseteq [v]_\sim$ is infinite, it contains still infinitely many green vertices. If $[v']_\approx \subseteq [v]_\sim$ and has only finitely many neighbors in $[v]_\sim$, then we changed the color for only finitely many of its neighbors of each of its finitely many vertices, thus (C3) remains true in this case as well. Now that we shown this third claim, we can apply Proposition 4.54 part (i) to obtain a Hamilton circle C_0 of $\bar{G}_0^3$ paying attention to all dummy-vertices (in other words satisfying $(M3)$).

Now to apply induction, we assume that G_i, $\bar{G}_i$, (C_i) are defined satisfying $(G1)$, $(G2)$, $(M1)$–$(M4)$. We define G_{i+1} and $\bar{G}_{i+1}$ as follows: Each dummy-vertex d in $\bar{G}_i$ comes from a unique weak class $[v_d]_\sim$. We put this class in V_i and then do the same construction as in the definition of G_0 and $\bar{G}_0$ for the graph represented by d. After we have done this for each dummy-vertex, we obtain the graphs G_{i+1} and $\bar{G}_{i+1}$.

Now to construct C_{i+1}, we have to replace the edges at each d from C_i by a path covering the new graph we obtained inside d as follows:

Let x and y be the neighbors of d in C_i such that x has distance at most 2 of d in G_i. Now we can choose big vertices x' and y' in d such that there are edges edges $xx' \in G^2$ and $yy' \in G^3$.

If $x' \neq y'$, we use 4.54 (ii) to find a new Hamilton path between x' and y' covering d' paying attention to all new dummy-vertices. Since its endvertices have G-distance at most 3 from x and y this path can be build in for the Hamilton circle G_{i+1}.

If $x' = y'$, then we find a Hamilton circle inside of the refined dummy vertex respecting x' and paying attention to all new dummy-vertices. Let x'' be the neighbor of x' witnessing that x' is respected. Since the edge xx' was in G^2, we will replace it by xx'' and find again a Hamilton path between x'' and y' whose endvertices have G-distance at most 3 from x and y covering d' paying attention to all new dummy-vertices.

After we done this replacement for each dummy-vertex d of G_i, we obtain a Hamilton circle C_{i+1} of G_{i+1}, paying attention to D_{i+1}, thus $(M2)$ is fulfilled. Also $(M2)$ clearly holds and we did not change any edges between vertices of V_i, which implies $(M4)$. Further since V_{i+1} contains at least the neighborhood of V_i, every vertex does fulfill $(M1)$ at some point.

This also implies that the edges at each vertex are defined eventually after finitely many steps and do not change later, so it is clear that a limit C of this sequence of circles is well-defined as the set of all vertices and ends of G^3 and all edges from further circles which are not adjacent with dummy-vertices.

It remains to show that C is a Hamilton circle of G^3.

First note that every vertex or end is in a C_i for some i and thus also in C.

Consider any homeomorphism $h_0 : S^1 \to C_0$. Now to obtain a homeomorphism $h_{i+1} : S^1 \to C_{i+1}$ for each G_{i+1} recursively, we change h_i on each interval on which a dummy-vertex together with two edges is replaced by an arc in the straight forward way such that the interval is mapped to the arc instead. Clearly each h_i is a homeomorphism. Now each element of S^1 will at some point be mapped to a vertex, end or edge of G^3 (and not to a dummy-vertex or an edge incident with a dummy-vertex) and hence its image does not change in a later step.

This way, the limit $h : S^1 \to C$ of the sequence $(h_i)_{i\in\mathbb{N}}$ is well-defined. To prove that h and its inverse is continuous at each element which is mapped to a vertex or inner edge point, we will show that for every such element $s \in S^1$ there is a neighborhood around s and an $i \in \mathbb{N}$ in which h and h_i coincide. That way the continuity of $h(s)$ follows from the continuity of h_i: In the case that $h(s)$ is an inner point of an edge, the statement holds clearly for the first i for which this edge is in C_i. If $h(s)$ is a vertex, we use the first i for which both edges in C_i incident with that vertex are the same as in h. Such an i exists, since every edge of C comes from some C_i.

Now it remains to show the continuity when $h(s)$ is an end ω. Let $[v]_\sim$ be the weak class corresponding to ω. Let $O := C_{G^3}(S,\omega) \cap C$ be a basic open neighborhood around ω in C. Now consider the first $i \in \mathbb{N}$, for which $\omega \in |\bar{G}_i^3|$. We define $O_i := C_{\bar{G}_i^3}(S,\omega) \cap C_i$. Per assumption is $h_i(O_i)$ an open set in S^1. Since all rays of ω live in any basic open neighborhood, it is clear that almost all vertices of $[v]_\sim \cup N_{G_i}([v]_\sim)$ are in $C_{\bar{G}_i^3}(S,\omega)$. Every arc that replaces an edge from one of the two ω-rays in C_i in a later step lives in a component K' with a big neighbor $b \in [v]_\sim$. Thus K' contains a vertex of $N_{G_i}([v]_\sim)$. Since almost all of them are not in S, there is an $S' \supseteq S$, such that for $O'_i := C_{\bar{G}_i^3}(S,\omega) \cap C_i$ and $O' := C_{G^3}(S',\omega) \cap C$ holds $h_i(O'_i) = h(O')$. This implies that $h(O')$ is an open set in S^1 for which $O' \subseteq O$.

Since S^1 is Hausdorff and C is compact it follows that $h^{-1} : C \to S^1$ is also continuous. $\square$

4.3 Hamilton circles in fourth and higher powers of trees

Our second main result of this chapter is that the fourth and higher power of any countable tree is always Hamiltonian. It seems natural that the proof for higher powers becomes easier then for the third power, which is indeed the case. However, even for graphs in which we already know that its third power is Hamiltonian, it is not immediately clear for higher powers, since the end space may be different for each power. Especially interesting is the fact that for every $n > 4$ there are examples of trees for which edges from $E(T^n) \setminus E(T^{n-1})$ are actually necessary to build a Hamilton circle. As already done for the third power of rayless graphs, our first step is to understand the end space of powers of countable trees. After that we will again construct a Hamilton circle recursively.

4.3.1 The endspace of powers of a tree

Consider any tree T. We will divide the ends of T^n into *preserved* and *new* ends. When constructing a Hamilton circle, we will deal with the preserved ends in the limit step, while we handle each new end in a similar way as in the last section.

Lemma 4.56. *Given a tree $T = (V, E)$ and $n \geq 2$. For $x, y \in V_\infty$, write $x \sim_n y$, if the distance of x and y is at most $n - 2$. Take the transitive closure of this relation (also denoted by $\sim_n$). Then the set of ends in T^n, can be written as a disjoint union $\Omega_1 \cup \Omega_2$, with the following properties:*

1. *There exists a canonical injection from Ω_1 to the set of ends of T. Each end in Ω_1 will be mapped to a subset of itself and the image of this injection is the set of ends of T, from which no ray meets a class of $V_\infty/{\sim_n}$ in infinitely many vertices. We call the ends in Ω_1* preserved ends.

2. *The set Ω_2 consists of one end for each equivalence class $[v]_{\sim_n}$ of $V_\infty/\sim_n$, where every ray of that end meets the union of that class with its first $\lfloor \frac{n}{2} \rfloor$ distance classes in T infinitely often. We call such an end* new end in $[v]_{\sim_n}$.

Proof. Let $n \geq 2$ be fixed. Given an equivalence class $[v]_{\sim_n}$, write $[[v]]_{\sim_n}$ for the set of vertices within distance $\lfloor \frac{n}{2} \rfloor$ of $[v]_{\sim_n}$ in T. For an edge $e \in T^n$, we define P_e as the unique path in T between the endvertices of e. For a subgraph $H \subseteq T^n$, we define $H(T) \subset T$ as the union of all P_e for all $e \in E(H)$.

First we note that for each $v \in V_\infty$ the set of vertices within distance $\lfloor \frac{n}{2} \rfloor$ from $\{v\}$ induces an infinite clique in T^n. If two vertices of V_∞ have distance at most $n-2$, there is an infinite matching between these cliques. It follows that for every class $[v]_{\sim_n}$ of $V_\infty/{\sim_n}$, the vertex set $[[v]]_{\sim_n}$ cannot be separated in T^n by finitely many vertices and therefore all rays meeting $[[v]]_{\sim_n}$ infinitely belong to same end in T^n.

For two vertices $v_1, v_2 \in V_\infty$ in different classes of $V_\infty/{\sim}$, each v_i together with an arbitrary ray R_i, which meets $[[v_i]]_{\sim_n}$ infinitely often, we will find a finite $R_1 - R_2-$separator in T^n: Consider the unique $v_1 - v_2$ path P in T. Without loss of generality, we can assume that v_1 and v_2 are the only vertices in V_∞ on this path. Let K be the component of $T - V_\infty$ containing P and K' be the finite subgraph of K, consisting of P and the first $2n$ distance classes of P.

We claim that $S = V(K') \cup N_T(V(K'))$ is the required separator. Note that S is finite, as $V(K')$ is finite, and every vertex of K' has finite degree. Suppose for a contradiction that there is an $R_1 - R_2-$path Q in T^n avoiding S. For each edge of this path between vertices from different components of $T - K'$, the roots from these components must have distance at most $n-2$ from each other, else there were no edge between other vertices from those components in T^n. This implies that those components are rooted at vertices of infinite degree. Hence, Q induces a sequence of vertices of infinite degree, in which successive vertices have distance at most $n-2$ from each other in T and thus $v_1 \sim_n v_2$, a contradiction.

In summary, there is indeed exactly one end for each class of $V_\infty / \sim_n$, so we can define these ends as Ω_2.

Whenever a ray of T meets a class $[[v]]_{\sim_n}$ in infinitely many vertices, then this ray belongs to the end of Ω_2 corresponding to $[v]_{\sim_n}$. It particular, it is possible that distinct normal rays of T belong to the same new end in Ω_2. We now show that whenever two distinct normal rays R_1 and R_2 of T do not belong to the same new end in Ω_2, then they are not equivalent in T^n. In other words, two distinct ends of T are merged in T^n only if both of them are contained in the same new end $\omega_2 \in \Omega_2$: If one or both of R_1 or R_2 are in a new end, then the statement is clear, so we may assume that neither R_1 nor R_2 meets any $[[v]]_{\sim_n}$ infinitely often. Let R be the unique double ray contained in $R_1 \cup R_2$. We will construct a finite separator between the tails of R in T^n:

In the case where R has a vertex a which does not lie in any $[[v]]_{\sim_n}$, define S as a together with its first $\lfloor \frac{n}{2} \rfloor + 1$ distance classes in T. By assumption, in the first $\lfloor \frac{n}{2} \rfloor$ distance classes of a are only vertices of finite degree, hence S is finite. Also, $R - S$ consists of two tails Q_1 and Q_2. Suppose for a contradiction that there is a path P between a vertex v_1 of Q_1 and v_2 of Q_2 in $T^n - S$. Since a has distance at least $\lfloor \frac{n}{2} \rfloor + 1$ from each vertex of P, no path P_e for e in $E(P)$ meets a. Hence $a \notin P(T)$, implying that $P(T)$ is not connected (as v_1 and v_2 are separated in T by a), a contradiction.

If every vertex of R belongs to some class $[[v]]_{\sim_n}$ for a vertex $v \in V_\infty$, then, since per assumption R meets no class $[[v]]_{\sim_n}$ infinitely, there are two vertices $w_1, w_2 \in E(R)$ such that there is no class $[[v]]_{\sim_n}$ containing both. Let $w_1 \in [[v_1]]_{\sim_n}, w_2 \in [[v_2]]_{\sim_n}$ and define Q_1 and Q_2 as the two disjoint tails of R, starting at w_1 and w_2. We obtain the tree T_+ from T as follows: For $i \in \{1, 2\}$, let x_i be the vertex of Q_i with Q_i-distance $\lfloor \frac{n}{2} \rfloor$ from w_i. For every vertex x in Q_i with Q_i-distance of at least $\lfloor \frac{n}{2} \rfloor$ from w_i, add infinitely many leaves as neighbors of x. In T_+^3 is each Q_i in the new end belonging to $[[x_i]]_{\sim_n}$ and $x_1 \nsim_n x_2$ in T_+, because $v_1 \nsim_n v_2$ in T, so there is a finite separator S in T_+^3 separating $[[x_1]]_{\sim_n}$ and $[[x_2]]_{\sim_n}$ and hence also Q_1 and Q_2. Because $T^3 \subseteq T_+^3$, it follows that $S|_{T^3}$ separates Q_1 and Q_2 in T^3.

Hence for each end of T, which is not contained in a new end, there is a unique end of T^3, which contains the end of T. Define the set of those ends as Ω_1.

It remains to show that these are indeed all ends, in other words that each ray of T^n lies in an end, which is either in Ω_1 or Ω_2. Let $R = v_0, v_1, v_2, ...$ be any ray of T^n and consider the subtree $R(T)$. Suppose that $R(T)$ has a vertex v of infinite degree. This means that infinitely many of the P_e for $e \in E(R)$ contain v. Because a path P_e has length at most n, one of the two paths between R and v of which this path consists has length smaller or equal to $\lfloor \frac{n}{2} \rfloor$, so at least one of its endvertices lies in $[[v]]_{\sim_n}$, so R meets $[[v]]_{\sim_n}$ in infinitely many vertices and hence is in the new end in $[v]_{\sim_n}$. If $R(T)$ is locally finite, we, we can find a Comb in $R(T)$ with spine R' and infinitely many teeth in R. Thus R' and R are equivalent and hence are in the same end. $\square$

4.3.2 Constructing the Hamilton circle

Remember why it was not always possible to find a Hamilton circle in G^3. The reason was that sometimes we obtained a captured arc around a new end because of not well-coverable components. Such a component X with root r was characterised by the fact that is was not possible to cover $X^3 - r$ by an arc between two neighbors of r. In $X^n - r$ for $n \geq 4$ this will be always possible, which can be considered as the main reason why there is always a Hamilton circle for powers higher than 3.

At first, we will show that in higher powers every finite subtree can be covered in such a good way. Remember our definition of the *bundles* from the proof of Lemma 4.36. We call a leaf *solitary* if it is in a bundle of size one.

Lemma 4.57. *For any finite rooted tree (T, r) with at least 3 vertices, there is a Hamilton circle C of T^3, which respects r and all solitary leaves of T.*

Proof. The lemma follows directly from the construction in the proof of Lemma 4.36: Covering a bundle smoothly as defined in the proof implies respecting one leaf from each bundle and especially respecting all solitary leaves. □

Lemma 4.58. *Let (T, r) be a finite rooted tree and $n \geq 4$. Then there exists a spanning path in $T^n - r$ with endvertices in the T-neighborhood of r, which respects all solitary leaves of T.*

Proof. If T is well-coverable, we even find such a path in T^3 (4.41), since the solitary leaves form an independent set. So we may assume that each component of $T - r$ is of size at least 2. Define $N_T(r) := \{v_0, v_1, ..., v_n\}$. From 4.57 we obtain for each component T_i of $T - r$ of size at least 3 a Hamilton circle of T_i^3 which respects v_i and all solitary leaves of T_i. After deleting the edge $v_i w_i$ respecting v_i, we obtain a Hamilton-path P_i of T_i^3 between v_i and one upper T-neighbor w_i of v_i. Note that w_i is no solitary leaf, since per definition are the edges respecting different vertices distinct. For a component of $T - r$ of size two with vertices v_i, w_i we define P_i as the edge $v_i w_i$. Now $v_0 P_0 w_0 w_1 P_1 v_1 w_2 P_2 v_2 w_3 P_3 v_3 ... w_n P_n v_n$ is the desired Hamilton-path. The only used edge from $T^n \setminus T^3$ is $w_0 w_1$. Every other edge is even in T^3. □

For each equivalence class $[v]_{\sim_n}$ of $V_\infty/\sim_n$, let $T_{[v]_{\sim_n}}$ be the smallest subtree of T containing $[v]_{\sim_n}$ (this is well-defined, since in a tree there is a unique path between each two vertices).

Lemma 4.59. *For two different classes $[v]_{\sim_n}$ and $[w]_{\sim_n}$ are the trees $T_{[v]_{\sim_n}}$ and $T_{[w]_{\sim_n}}$ disjoint.*

Proof. Suppose for a contradiction that $x \in T_{[v]_{\sim_n}} \cap T_{[w]_{\sim_n}}$. Each vertex x in $T_{[v]_{\sim_n}}$ lies on a path between two vertices of $[v]_{\sim_n}$. Since these paths have length at most $n - 2$, there is a

$v' \in [v]_{\sim_n}$ with distance at most $\frac{n-2}{2}$ from x. Also we find such a vertex $w' \in [w]_{\sim_n}$. If follows that v' and w' have distance at most $n-2$ from each other in T, so $[v]_{\sim_n} = [w]_{\sim_n}$. $\square$

For a subtree $U \subseteq T$, we define U^+ as follows:
For each component K of $T - U$, we add one or two vertices to U:
Call the unique vertex of K which is adjacent to U r_K and add it to the subtree. Whenever K has more than one vertex, we also add another vertex v_K from K, which is adjacent to r_K in T, chosen in the same tree $T_{[v]_{\sim_n}}$ as r_K, whenever r_K is in such a tree, and else arbitrarily. This graph is uniquely defined under isomorphism and whenever we refer to it, we assume that we fixed one possible choice.

Proposition 4.60. *Given $n > 3$ and $U = T_{[v]_{\sim_n}}$ as defined before with a root r and one of its neighbors w arbitrarily chosen in U, the graph U^{+^n} has a Hamilton circle containing rw and also the edge $r_K v_K$ for every components K of $T - U$, for which this edge exists.*

Proof. Let $\{v_0, v_1, v_2, ...\}$ be an enumeration of $V(U)$. Now, we enumerate the vertices of U and components of $U^+ - U$ as follows: Consider a sequence $(a_n)_{n \in \mathbb{N}}$ of natural numbers in which every number appears infinitely often. For each a_i in our sequence, we add v_i to our enumeration, if this a_i appears for the first time. Also we possibly add components of $U^+ - U$ with neighbor v_i to our enumeration:
Whenever there are only finitely many such components, we add all of them at the end of our current enumeration, if they were not added in an earlier step. Whenever there are infinitely many such components, we assume they are ordered in a fixed order and add the first two of them, which were not chosen before to the end of our enumeration. We call the components which are added in one such step a *convolute* of components.

Because U^{+^n} is one ended, it remains to construct a spanning double ray.

We construct a sequence of a singleton $P_0 := \{v_0\}$ and paths $P_1 \subseteq P_2 \subseteq P_3 \subseteq ...$ in U^{+^n} such that:

(P1) For each vertex of U or component of $U^+ - U$, there is a path P_i that covers it.

(P2) the endvertices from each path P_i have T-distance at most 1 from U.

(P3) For each root $r_K \in P_i$ of a component of $U^+ - U$, is the edge $r_K v_K$ also in P_i, whenever v_K exists.

We extend the path alternately from both sides. Let P_i be given with endvertices w_1 and w_2 (for P_0, let both of them be v_0) and assume without loss of generality that we have to extend the path on w_2 in this step.

Consider the next element X of our enumeration, which is not already in P_i. We will cover it in the $(i+1)$th step and thus satisfy $(P1)$.

If X is in $V(U)$, we define $b' := X$. If X is a component of $U^+ - U$, we define b' as the unique T-neighbor in U of X. Let a' be w_2 if $w_2 \in U$ and else the unique neighbor in U of the component of $U^+ - U$ containing w_2.

Then there is a unique $a' - b'$-path $a' = q_0, ..., q_l = b'$ in U. Let $p_1, p_2, ..., p_m$ be the vertices of infinite degree on this path in order. Per definition of $\sim_n$, two successive vertices of infinite degree in this sequence have T-distance at most $n-2$ from each other. For each p_j, whenever there are infinitely many components of $U^+ - U$ adjacent to p_j, P_i meets only finitely many of them and we choose one convolute of two such components K_j, K_j' of $U^+ - U$ adjacent to p_j with roots $r_{K_j}, r_{K_j'}$. We define a path $Q_j := r_{K_j}, v_{K_j}, v_{K_j'}, r_{K_j'}$, leaving out v_{K_j} or $v_{K_j'}$ if it does not exist. (Since $n \geq 4$, the edge $v_{K_j}v_{K_j'}$ is in T^n.)

In the case that there are not infinitely many components of $U^+ - U$ adjacent to p_j, p_j has infinite degree in U. Thus, we can choose two U-neighbors u and v of p_j not covered by P_i. Let $U_0, ..., U_m$ be a convolute of components of $U^+ - U$ adjacent to u and $V_0, ..., V_n$ be a convolute of components of $U^+ - U$ adjacent to v, whenever such a convolute exists. (Since u and v are not already covered by P_i, this convolutes are not covered as well.) We define a path $Q_j := uv_{U_0}r_{U_0}, v_{U_1}, r_{U_1}, ..., v_{U_m}r_{U_m}r_{V_0}v_{V_0}r_{V_1}, v_{V_1}, ..., r_{V_n}v_{V_n}v$, leaving out every vertex which does not exist. Note that all these edges are in T^n. In both cases are the endvertices of Q_j of T-neighbors of p_j, thus, because p_j and p_{j+2} have distance at most $n-2$, there is for each $j < m$ an edge in T^n between the endvertex of Q_j and the initial vertex of Q_{j+1}. Consequently, we can extend P_i by the path $w_2Q_1Q_2...Q_m$. Let r_m be the endvertex of this path. Then we add another segment to our arc, depending on how X lies in the graph:

If $X \in U$, we add the edge r_mb and are done with the construction.

If $X \notin U$, let $B_0, ..., B_n$ be the convolute of components containing X. We extend our path by $v_{B_0}r_{B_0}, v_{B_1}, r_{B_1}, ..., v_{B_n}r_{B_n}$ leaving out every v_{B_i} which does not exist. In any case, we end up with a vertex of T-distance at most 1 from U, so $(P2)$ holds. The statement $(P3)$ is also clear per construction.

After we have done countably many steps, according to $(P1)$ and $(P3)$, we constructed a double-ray covering all vertices of U^+ containing rv and also the edge r_Kv_K for every components K of $U^+ - U$, for which this edge exists. $\square$

Theorem 4.61. *For a countable tree T and any $n \geq 4$, T^n have a Hamilton circle.*

Proof. We construct a Hamilton circle of T^n.

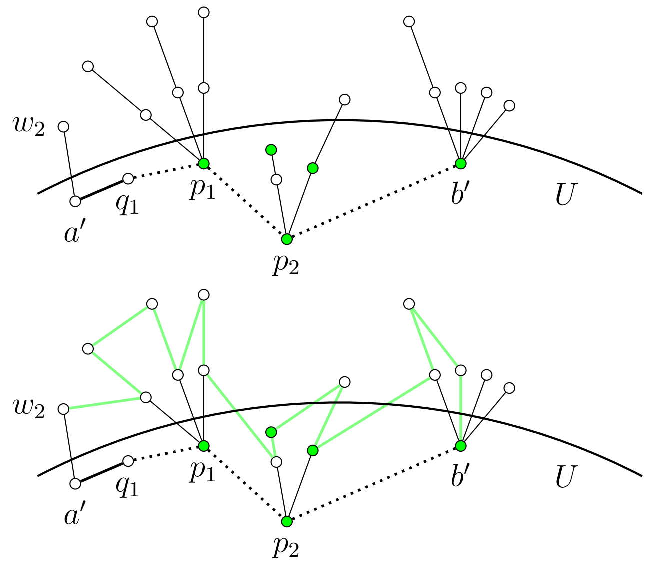

Figure 9: A path from w_2 to b'

Choose a class $[v]_{\sim_n}$ in $V_\infty/\sim_n$ arbitrarily. Starting with $T_0 := T_{[v]_{\sim_n}}$, we define a sequence of subgraphs $T_0 \subseteq T_1 \subseteq T_2 ...$ of T as follows:

If $T_i = T$ or $T_i^+ = T$, then stop the sequence. (this will only happen, if T^n has no preserved end.) If not, then given T_i, we construct T_{i+1} as follows:

For every component K_i of $T - T_i$, we add the following subtree:

Whenever r_{K_i} is in a tree $T_{[w]_{\sim_n}}$, we add $T_{[w]_{\sim_n}}$ as a whole. Because of the way we prioritised the choice of v_{k_i}, we made sure that it is also in T_{i+1}. If r_{K_i} is not in any $T_{[v]_{\sim_n}}$, then we add only r_{K_i}.

Now we will define a sequence of Hamilton circles C_i for each $T_i^{+^n}$ such that:

(i) C_i contains all edges of the form $r_{K_i}v_{K_i}$ for components K_i of $T - T_i$.

(ii) $C_{i+1} \setminus C_i$ is a disjoint union of arcs, each of them replacing an edge of the form $r_{K_i}v_{K_i}$. Further this arc replacing an edge $r_{K_i}v_{K_i}$ lies in K.

The Hamilton circle C_0 exists due to Proposition 4.60. To construct the Hamilton circle C_{i+1} from C_i it remains to replace the edge $r_K v_K$ for each component of $T - T_i$ by an arc that covers everything from that component in T_{i+1}^+:

In case that this is a tree $T_{[w]_{\sim_n}}$ we can apply Proposition 4.60 again inside that component and obtain another circle containing $r_K v_K$, so we can replace this edge by the rest of the circle. For every new component K' arising this way inside K, the new circle also uses the edge $r'_K v'_K$, so we can make sure that (i) still holds.

In case that we added from the component K only r_K to T_{i+1}, r_K has finite degree and $T_{i+1} \cap K$ consists of a star with center r_K, from which some of its leaves may be subdivided. Let $K_1, K_2, ...$ be the components of $K - r_K$ with roots $v_K = r_0, r_1, r_2, ..., r_x$ and for each of these components K_y with more than one vertex, let v_y be the second vertex chosen for T^+_{i+1}. We replace the edge $r_K v_K$ with the finite path $v_K v_0 v_1 r_1 v_2 r_2 ... v_x r_x$ leaving out every v_y that does not exist. Clearly this path uses every edge of the form $v_y r_y$ and hence again (i) is satisfied.

Since in each step, we added at least the root of every component of $T - T_i$, we made sure that each vertex is eventually in a T_i and also covered by one Hamilton circle C_i.

We define a compatible sequence of homeomorphisms $h_i : S^1 \to C_i$ as in the proof of Theorem 4.55 and its limit h' on each point on which this sequence is eventually constant. The continuity of $h'(s)$ and h'^{-1} for each element s which is mapped to a vertex, inner edge point or new end can be shown with the same argument as in the proof of Theorem 4.55 as well.

Also it follows from the fact that each h_i is a homeomorphism that h' is injective. Further the sequence $(h_i)_{i \in \mathbb{N}}$ becomes eventually constant at each point which is mapped to a vertex, inner edge point or new end and thus all these element are in the image of h'. Now we define $h(s) := h'(s)$ for all $s \in S^1$ for which $h'(s)$ is defined.

Consider now an $s \in S^1$ for which the sequence $(h_i(s))_{i \in \mathbb{N}}$ does not become eventually constant. This can only be the case, whenever each $(h_i(s))$ is an inner point of an edge which is replaced in a later step. Let $e_1, e_2, e_3, ...$ be the sequence of these edges and for each of this edges e_i, let A_i be the arc replacing it.

Since each A_i contains multiple edges, we may assume that the sequence of intervals mapped to $e_1, e_2, e_3, ...$ converges to the single point s. Now each edge e_i on this sequence is of the form $r_K v_K$ for some component K as defined above. The sequence of those r_K defines a unique ray in T, which belongs to an end ω. We define $h(s) := \omega$.

To show that h is still injective, let $s \in S^1$ be an element with $h(s) := \omega$ for which $h'(s)$ is not defined and $s' \in S^1$ be any other element. If $h(s')$ is not an end, there is nothing more to show, so we may assume that $h(s')$ is an end ω'.

If ω' is a new end in a class $[v]_{\sim_n}$ of $V_\infty / \sim_n$, then $h(s') = h'(s')$ and hence there is an $i \in \mathbb{N}$ for which $h_i(s') = \omega'$. The circle C_{i+n} covers $[v]_{\sim_n}$ and its first $\lfloor \frac{n}{2} \rfloor$ distance classes in T. Since the ray from ω in T contains a tail of outside of C_{i+n}, it cannot lie in ω'. It follows that $h(s) = \omega \neq \omega' = h(s')$.

If ω' is a preserved end, then also each $(h_i(s'))$ is an inner point of an edge which is replaced in a later step and there is another sequence $f_1, f_2, f_3, \ldots$ of these edges, each of them with an arc B_i be the arc replacing it. Now let i be chosen such that $e_i \neq f_i$. Now there are two disjoint components K_A and K_B of $T - T_i$ such that for all $j > i$ does A_j lie in K_A and B_j in K_B. It follows again that $\omega \neq \omega'$.

Now it remains to show the continuity of h at any given s for which $h(s)$ is a preserved end ω. Let $C(S, \omega) \cap C$ be any basic open set in C around ω. Since S is finite, there are only finitely many of the arcs A_i meeting S. Let A_j be the last one of them. Now define I as the interior of the interval $h_{j+1}^{-1}(A_{j+1})$. Then $h(I) \subseteq \bigcup_{i>j} A_i \subseteq C(S, \omega)$. It follows that h is continuous in s. Since S^1 is compact and C is Hausdorff it follows that $h^{-1} : C \to S^1$ is also continuous. Since the image of h is closed and contains all vertices, it is clear that it also contains all ends of T^n. $\square$

5 Nash Williams' orientation theorem for infinite graphs

5.1 Introduction

A directed multigraph is *k-arc-connected* if from any vertex v to any other vertex w of the graph there exist k arc-disjoint forwards directed paths. Clearly, the underlying undirected graph of a k-arc-connected multigraph must be $2k$-edge-connected. The classic *orientation theorem* of Nash-Williams from 1960 asserts that for finite multigraphs, also the following converse is true.

Theorem 5.1 (Nash-Williams' orientation theorem [34])**.** *Every finite $2k$-edge-connected multigraph has a k-arc-connected orientation.*

In the same paper, Nash-Williams claimed that his result also holds for infinite graphs – but the promised proof was never published and the claim was not repeated in [35]. Despite significant effort, it has remained open ever since whether the orientation theorem holds for infinite graphs as well.

So far, for arbitrary infinite graphs, only the case $k = 1$ was known, proved by Egyed by a Zorn's lemma argument already in 1941 [20].

To appreciate the difficulty of the general case, note that a priori it is not even clear whether *any* sufficiently large edge-connectivity implies the existence of a k-arc-connected orientation. This is different for finite multigraphs, where a simple argument shows that every $4k$-edge-connected multigraph has a k-arc-connected orientation: By the Nash-Williams/Tutte tree packing theorem [14, Corollary 2.4.2], any such graph has $2k$ edge-disjoint spanning trees, so after fixing a common root, we may simply orient half of the trees away from and the other half towards the root. This approach, however, is blocked for infinite graphs: there exist locally finite graphs of arbitrarily large finite (edge-)connectivity that do not even possess three edge-disjoint spanning trees [1].

Motivated by the above considerations, Thomassen has asked in 1985 whether there is a function $f : \mathbb{N} \to \mathbb{N}$ such that any $f(k)$-edge-connected multigraph has a k-arc-connected orientation [43]. This conjecture has been featured again in [4, Conjecture 8], where also a topological variation of the problem was suggested by allowing directed topological arcs in $|G|$; this topological version has been recently solved by Jannasch [25].

More than 50 years after Nash-Williams' finite orientation theorem and about 30 years after posing his own conjecture, Thomassen achieved a marvellous breakthrough towards the orientation theorem by proving that every finite $8k$-edge-connected multigraph has a

k-arc-connected orientation [44], giving $f(k) \leq 8k$. In this chapter, we show $f(k) \leq 4k$ for all graphs. Further we show for $2k$-edge-connected multigraphs with at most countably many ends, from which at most one end has odd degree that we can improve Thomassen's argument in order to get the best possible bounds, thereby establishing Nash-Williams' orientation theorem for some infinite graphs in its optimal form.

Furthermore, some steps in the proof work for arbitrary graphs, so it is possible that some of the techniques in our proof might be helpful in the future for other classes of graphs. We remark that our proof employs Mader's *lifting theorem* from 1978 [31]. There are also slightly other versions from 1992 [23] and 2016 [36], results that were certainly not available to Nash-Williams in 1960.

5.2 Boundary-linked decompositions

Let $G = (V, E)$ be a locally finite connected multigraph. The *boundary* of a set of vertices B is the collection of edges in G with one endvertex in B and the other one outside of B.

A set of vertices $B \subset V$ is called *boundary-linked* if the induced subgraph $G[B]$ together with its boundary has a collection of pairwise edge-disjoint equivalent rays $R_1, R_2, \ldots$ such that each edge in the boundary is the first edge of one of the rays R_i. If we also want to point out to which end ω those rays belong, we say that B is *ω-boundary-linked*.

Thomassen proved in [44] that for every locally finite connected multigraph $G = (V, E)$ and any given finite set of vertices $A' \subseteq V$, $V(G) \setminus A'$ can be partitioned into finitely many sets each of which is either a singleton or a boundary-linked vertex set with finite boundary in G:

Theorem 5.2. [44] *Let G be a connected, locally finite multigraph. Given any finite set of vertices $A' \subseteq V$, there is a finite set of vertices $A \supseteq A'$ such such that the vertices of $G - A$ can be partitioned into finitely many boundary-linked vertex sets with finite boundaries.*

We do not actually know in general, whether we can choose our partition always without edges between the boundary-linked sets:

Problem 5.3. *Let G be a connected, locally finite multigraph. Is it true that given any finite set of vertices $A' \subseteq V$, there is a finite set of vertices $A \supseteq A'$ such such that all components of $G - A$ are boundary-linked?*

However for multigraphs with countably many ends, we can obtain such a partition:

Theorem 5.4. *Let $G = (V, E)$ be a connected, locally finite multigraph with at most countably many ends. Given any finite set of vertices $A' \subseteq V$, there is a finite set of vertices $A \supseteq A'$ such such that all components of $G - A$ are boundary-linked.*

Proof. Let $\{\omega_1, \omega_2, ...\}$ be an enumeration of the ends of G.

Let $E_1 = \{v_1w_1, v_2w_2, ..., v_kw_k\}$ be a minimal $A' - \omega_1-$separator such that $v_1, ..., v_k$ are the endvertices of the edges from E_1 in the side of the cut that contains A'. Define $A_1 := A' \cup \{v_1, ..., v_k\}$. By minimality of E_1, the component K_0 of $G - E_1$ containing ω_1 has boundary E_1. We show that K_0 is boundary-linked: We define a sequence of connected subgraphs $K_1 \supsetneq K_2 \supsetneq K_3 ....$ To define K_i, delete all vertices from K_{i-1} incident with its boundary and then define K_i as the unique infinite component of the resulting subgraph of K_{i-1} that contains ω_1. Because E_1 was a minimal $A' - \omega_1-$separator, it follows that K_i also has a boundary of size at least k. By Menger's theorem, G has k pairwise edge-disjoint paths $P_1^i, P_2^i, ..., P_n^i$ such that P_j^i starts with $v_j w_j$ and terminates with an edge in the boundary of K_i for $j \in [n]$ and $i \in \mathbb{N}$. For every $j \in [n]$ we define a limit ray R_j from the path system $\mathcal{P}_j = \{P_j^i \colon i \in \mathbb{N}\}$ as follows: Since G is locally finite, for infinitely many i, the paths P_j^i in $\mathcal{P}_j$ have the same second edge. For infinitely many of those i, the paths P_j^i also have the same third edge, and so on. Repeating this argument, we obtain a a sequence of edges giving rise to a ray R_j starting with the edge $v_j w_j$. Clearly these rays $R_1, \ldots, R_n$ all belong to the end ω_1, witnessing that $B_1 := K_0$ is boundary-linked.

We now define $A_2, A_3, ...$ and $B_2, B_3, ...$ recursively: Since G is localy finite, $G - (A_i \cup B_1 \cup ... \cup B_i)$ has only finitely many components. If all of them are finite for some i, then we add the remaining vertices to A_i and obtain our desired partition of $V(G)$. If not, then each infinite component contains a ray (again, because G is locally finite). Let ω_k be the least end in our enumeration for which there is a ray left. With the same construction as above for A_1 and B_1, we obtain the finite vertex set A_{i+1} and a boundary linked set B_{i+1} containing all rays of ω_k. Suppose for a contradiction that this procedure does not terminate. In this case, the sets $B_1, B_2, ...$ together with all inner vertices of its boundary-edges form an open cover of the endspace $\Omega(G)$. Since G is locally finite, the endspace $\Omega(G)$ is compact, so there is a finite subcover. It follows that we already covered all ends after finitely many steps, a contradiction. $\square$

An end ω is called *even*, if there exists a finite set of vertices S such that for all finite sets of vertices $S' \supseteq S$ holds that the maximal number of edge-disjoint rays in ω starting in S' is even. Otherwise, the end is called *odd*.

Theorem 5.5. *Let $G = (V, E)$ be a connected, locally finite multigraph with only countably many ends. Given any finite set of vertices $A' \subseteq V$, there is a finite set of vertices $A \supseteq A'$ such such that every component B of $G - A$ is ω_B-boundary-linked for an end ω_B.*

Furthermore we can chose A in a way such that every component B of $G - A$ has even boundary whenever the end ω_B is even.

Proof. We do the same construction as in the proof of 5.4, apart from that whenever we construct a set B for an even end ω_B, we add a vertex set witnessing that ω_B has even degree to A_i. This way we made sure that our boundary of B has indeed even size. $\square$

5.3 Mader's lifting theorem and the lifting graph

Lifting two distinct edges vx, vy incident with a common vertex v in a multigraph G means deleting them and adding a new edge xy to G (possibly parallel to existing edges between x and y).

Suppose $G = (V + v, E)$ is a finite multigraph such that any two vertices in V are joined by k pairwise edge-disjoint paths in G. A pair of edges vx, vy is called *admissible for edge-connectivity* k, or simply *admissible* if the connectivity constant k is understood from context, if after lifting vx, vy we obtain a graph G' in which still any two vertices in V are joined by k pairwise edge-disjoint paths in G'.

We use Mader's Lifting theorem in the following version of Frank [23].

Theorem 5.6 (Mader, Frank)**.** *Suppose that $G = (V + v, E)$ is a finite connected multigraph such that any two vertices in V are joined by k pairwise edge-disjoint paths in G. If v is not incident with a bridge and $d(v) \neq 3$, there are $\lfloor \frac{d(v)}{2} \rfloor$ pairwise disjoint admissible pairs of edges incident to v.*

Two admissible pairs are called *compatible* if after lifting one of them, the second one is still admissible in the resulting graph. In this context, we also call the liftings *compatible*. Clearly no pair of edges becomes admissible after lifting another pair, if it has not been admissible before. However, the opposite is possible. Not every two liftings are compatible:

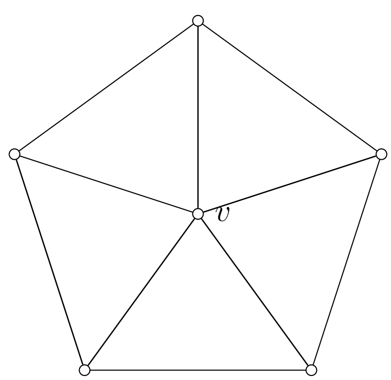

Figure 10: Every two edges at v form an admissible pair, but the liftings are not compatible, since lifting two pairs would destroy the 3-edge-connectivity.

The *lifting graph* $L(G, v, k)$ is the graph whose vertices are the edges incident with v, and two vertices e_i, e_j are adjacent if (e_i, e_j) is an admissible pair for edge-connectivity k. From this perspective, Theorem 5.6 implies that under the above assumptions on G, if $d(v)$ is even, then $L(G, v, k)$ has a perfect matching.

Substantial research on the structure of the lifting graph was done by Ok, Richter and Thomassen [36].

Theorem 5.7 (Ok, Richer and Thomassen)**.** *Let $k \geq 2$ be even, and $G = (V + v, E)$ be a finite connected multigraph such that any two vertices in V are joined by k pairwise edge-disjoint paths in G. If v is not incident with a bridge and $d(v) \geq 4$, then:*

- *If $d(v)$ is odd, then $L(G, v, k)$ is either connected or it has two components, one of them being a singleton and the other one a complete multipartite graph.*
- *If $d(v)$ is even and k is odd then $L(G, v, k)$ is either connected or it has two even components, both of them being a complete multipartite graph.*
- *If $d(v)$ is even and k is even, then $L(G, v, k)$ is a connected complete multipartite graph.*
- *If $d(v) = 5$, then $L(G, v, k)$ is either an isolated vertex plus a 4-cycle or a connected graph. If k is even and $L(G, v, k)$ is connected, then $L(G, v, k)$ is a complete multipartite graph.*

5.4 Immersions of finite graphs of prescribed connectivity

If G is a multigraph and H is a nother multigraph with vertices $x_1, x_2, \ldots, x_n$, then an *immersion* of H in G is a subgraph of G consisting of n distinguished vertices $y_1, y_2, \ldots, y_n$ and a collection of pairwise edge-disjoint paths in G such that for each edge $x_i x_j$ in H there is a corresponding path in the collection from y_i to y_j. This immersion is said to be *on* $\{y_1, \ldots, y_n\}$.

Thomassen proved in [44, Theorem 4] that for any finite set of vertices A in a $4k$-edge-connected locally finite multigraph G, there is an immersion in G of a finite Eulerian $2k$-edge-connected multigraph on A:

Theorem 5.8. [44] *Let k be a natural number, $G = (V, E)$ be a $4k$-edge-connected multigraph and $A \subseteq V$ be a finite set of vertices. Then G contains an immersion of a finite Eulerian $2k$-edge-connected multigraph with vertex set A.*

Our aim is to find a special immersion that reflects the original edge-connectivity in A.

For a multigraph H and a set $A \subseteq V(H)$, we say that A is *k-edge-connected in H*, if $\lambda_H(a, b) \geq k$ for all distinct $a, b \in A$.

Further we say for an orientation $\vec{H}$ of H that A is *k-arc-connected in $\vec{H}$*, if for every two distinct vertices x, y in A, there are k arc-disjoint directed paths in $\vec{H}$ from x to y and from y to x.

Definition 5.9. A set $A \subseteq V(G)$ in a $2k$-edge-connected multigraph G is called *immersible for edge-connectivity* $2k$ if there is a set X containing exactly one vertex of each component of $G - A$ with a boundary of odd size and if G contains an immersion of a finite multigraph H on $A \cup X$ with the following properties:

(i) $d_H(x) = 3$ for all $x \in X$ and

(ii) A is $2k$-edge-connected in H.

We call a pair (e_i, e_j) of boundary-edges of an ω-boundary-linked set $B \subseteq V$ B-admissible, if the pair (e_i, e_j) becomes admissible after contracting B.

For a multigraph $G = (V, E)$ and an ω-boundary-linked set $B \subseteq V$ with finite boundary $e_1, ..., e_q$ of size q, we define a *B-lifting* as a set of disjoint B-admissible pairs $\{(e_{i_1}, e_{j_1}), (e_{i_2}, e_{j_2}), ..., (e_{i_p}, e_{j_p})\}$ of edges in the boundary of B with $p = \frac{q}{2}$ if q is even and $p = \frac{q-3}{2}$ if q is odd such that the lifting of all pairs in this set is compatible. Note that recursively applying the lifting Theorem 5.6 implies that there is always at least one B-lifting for each such B.

Definition 5.10. We call a boundary linked set B with boundary of size q *strongly boundary-linked* in G if it satisfies the following property:

- If q is even, then there is a B-lifting, for which there is a set of edge-disjoint paths in $G[B]$, each of them connecting the edges of one of its pairs, and
- If q is odd, then there is a B-lifting, for which there is a vertex x in B and a set of $\frac{q-3}{2} + 3$ edge-disjoint paths in $G[B]$, $\frac{q-3}{2}$ of them connecting the edges of one of those pairs, and 3 of them between x and one of the three edges f_1, f_2 or f_3 that are not in the B-admissible pairs.

Proposition 5.11. [44] *Every boundary-linked set with a boundary of even size in a* 2*-edge-connected multigraph* G *is strongly boundary-linked in* G*.*

With a deeper analysis of the lifting graph [2], Amena Assem showed the following in June 2023:

Proposition 5.12. [3] *Every boundary-linked set with a boundary of odd size in a* 4*-edge-connected multigraph* G *is strongly boundary-linked in* G*.*

Proof. The statement is implied by Lemma 3.1 in [3] □

However, this result was not available by the time of my research, so we will give a direct proof in the next section only for a boundary of size 5.

Unfortunately, it is not sufficient for a set A to be immersible when all components of $G - A$ are strongly boundary linked, since after realising the linkage in a component with odd boundary, due to the vertex of degree 3, the resulting graph is no longer 4-edge-connected. However, after linking the boundary of a component of even degree as in Definition 5.10 we keep the 2-edge-connectivity of the whole graph. Thus Lemma 5.11 can be applied multiple times.

Proposition 5.13. *Let* $k \geq 2$ *be a natural number,* G *be a* $2k$*-edge-connected locally finite multigraph, and* A *be a finite set of vertices in* G*, such that every component of* $G - A$ *is boundary-linked and has a boundary of even size. Then* G *contains an immersion of a* $2k$*-edge-connected finite multigraph* H *on* A*.*

We omit the proof of Proposition 5.13, since it can be easily deduced from the proof of the following Theorem:

Theorem 5.14. *Let $k \geq 2$ be a natural number, G be a $2k$-edge-connected locally finite multigraph, and A be a finite set of vertices in G, such that every component of $G - A$ is* strongly boundary-linked *and exactly one component of $G - A$ has a boundary of odd size.*

Then the component with boundary of odd size contains a vertex x such that G contains an immersion of a finite multigraph H on $A \cup \{x\}$ with the following properties:

(i) $d_H(x) = 3$ *and*

(ii) A *is $2k$-edge-connected in H.*

Proof. Since G is locally finite and A is finite, there are only finitely many components $B_1, B_2, ..., B_n$ of $G - A$. Without loss of generality, we assume that B_n is the component with a boundary of odd size.

Starting with $G_0 = H_0 := G$, we define a sequence of immersions $G_0, G_1, ..., G_n$ of $2k$-edge-connected multigraphs $H_0, H_1, ..., H_n$, such that each H_i for $i \in \{1, ..., n-1\}$ satisfies:

(I1) $V(H_i) = V(H_{i-1}) \setminus V(B_i)$

(I2) A is $2k$-edge-connected in H_i.

To obtain G_i from G_{i-1}, we link the boundary of B_1 as in Definition 5.10 (see Lemma 5.11). After we define $V(H_i) = V(H_{i-1}) \setminus V(B_i)$ and add for each linking path an edge in H_i instead, G_i becomes an immersion of H_i. Per definition of the liftings is H_i still $2k$-edge-connected.

Now H_{n-1} is a $2k$-edge-connected multigraph on $A \cup B_n$. Since $2k \geq 4$, we can apply Lemma 5.12 on H_{i-1} to find a linkage in B_n according to Definition 5.10. Replacing B_n with that linkage, we obtain an immersion on a finite multigraph H satisfying (i) and (ii) per definition. $\square$

5.5 Linking boundaries of size five

A set $\mathcal{R}$ of rays witnessing that a set B is boundary-linked is called a *boundary-linking set.*

For each boundary-linking-set $\mathcal{R} = \{R_1, R_2, ..., R_q\}$, we define another graph $M(\mathcal{R})$ with vertex set $\{e_1, ..., e_q\}$ and edges between every two vertices e_i, e_j if $G[B]$ has a collection of infinitely many pairwise disjoint paths joining R_i, R_j having no edges in common with $R_1 \cup R_2 \cup ... \cup R_q$. In the context of a boundary-linking-set $\mathcal{R}$ of rays, we call a path between two of them having no edges in common with $R_1 \cup R_2 \cup ... \cup R_q$ an $\mathcal{R}$-path. For an $\mathcal{R}$-path P between R_i and R_j, we call the unique $e_i - e_j-$path in $R_i \cup R_j \cup P$ an $e_i - e_j-$*passage*. Note that after deleting the edges of a passage from B, it stays boundary-linked with a boundary of size $q - 2$.

We will show Theorem 5.12 for the special case of a boundary of size 5. For the rest of this section, let G be a 4-edge-connected multigraph and B be a boundary-linked set with boundary of size 5. We will show that B is strongly boundary-linked.

Consider two rays R_i, R_j in a boundary-linking set $\mathcal{R}$ with

$$V(R_i) = \{p_0, p_1, p_2, ...\}, E(R_i) = \{p_0p_1, p_1p_2, ...\},$$

$$V(R_j) = \{q_0, q_1, q_2, ...\}, E(R_j) = \{q_0q_1, q_1q_2, ...\}.$$

Two $\mathcal{R}$-paths $P_1 = p_a, ...q_b$ and $P_2 = p_c, ...q_d$ such that $a \leq c$ and $b \geq d$, are called *crossing*. The rays R_i, R_j in $\mathcal{R}$ are called *interchangeable* if there are two crossing $\mathcal{R}$-paths between them.

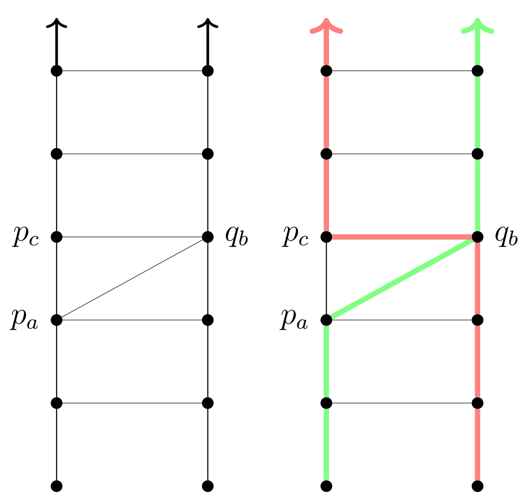

Figure 11: Two rays with crossing paths in the case $b = d$

Remark 5.15. In the notation from above, whenever R_i, R_j are interchangeable, we can exchange R_i with $p_0R_ip_aP_1q_bR_j$ and R_j with $q_0R_jq_dP_2p_cR_i$ and obtain a new boundary-linking-set $\mathcal{R}' = \{R'_1, ..., R'_q\}$ such that $R'_n = R_n$ for all $n \neq i, j$ and R'_i begins with e_i and contains a tail of R_j and vice versa.

The proof of the following lemma is similar to the second proof of Menger's Theorem in [14, Theorem 3.3.1].

Lemma 5.16. *Let $G = (V, E)$ be a k-edge-connected multigraph and $s \in V$ a vertex and $T \subseteq V$ be a set of vertices not containing s. For every set $\mathcal{P} = \{P_1, P_2, ...P_m\}$ of fewer than k edge-disjoint $s - T-$paths, we can find another set $\{P'_1, P'_2, ...P'_k\}$ of k edge-disjoint $s - T-$paths such that each P_i for $i \leq m$ has the same endvertices as P'_i.*

Proof. It is sufficient to show that we can find $m + 1$ such paths. We apply induction after $|P_1| \cup |P_2| \cup ... \cup |P_m|$. It is clear that the induction starts for the empty set of paths, since G is connected. Without loss of generality, we assume that each P_i meets T only in its endvertex. Now let t_{m+1} be any vertex from T, which is not the endvertex of any path in $\mathcal{P}$ and let P_{m+1} be any $s - t_{m+1}-$path avoiding the endvertices from $P_1, P_2, ...P_m$ on T. If $P_1, P_2, ...P_{m+1}$ are edge-disjoint, we are done. If not, let ab be the last edge of P_{m+1} on a path $P \in \mathcal{P}$. Define $T' := T \cup V(bP_{m+1} \cup bP)$ and $\mathcal{P}' := (\mathcal{P} \setminus \{P\}) \cup \{Pb\}$. Since $\mathcal{P}'$ satisfies the induction hypothesis, there is an extension $\mathcal{P}''$ of $m+1$ edge-disjoint $s - T'-$paths satisfying the lemma. Let P'' be the path of $\mathcal{P}''$ ending in b and P''_{m+1} be the path of $\mathcal{P}''$ with an endvertex y outside of the endvertices of the paths from $\mathcal{P}'$. If $y \notin bP$, we can extend P'' by bP and P''_{m+1} by yP_{m+1} (if y is not already in T). Otherwise is $y \in bP - b$ and we can extend P'' by bP_{m+1} and P''_{m+1} by yP to obtain the desired set of edge-disjoint paths. □

Lemma 5.17. *For each edge R_iR_j in $M(\mathcal{R})$ are R_i and R_j either interchangeable or there is a boundary-linking-set $\mathcal{R}'$, which can be obtained from $\mathcal{R}$ by replacing the ray R_i with a ray R'_i such that $\mathcal{R}'$ satisfies the following:*

(R1) $M(\mathcal{R}') \supseteq M(\mathcal{R})$ *and*

(R2) there is a ray $R_k \in \mathcal{R}' \setminus \{R'_i, R_j\}$ from which there is an $\mathcal{R}'$-path to R'_i and an $\mathcal{R}'$-path to R_j.

Proof. Per assumption are there infinitely many edge-disjoint $\mathcal{R}$-paths between R_i and R_j. Let P_1, P_2, P_3 be three of them. If two of them are crossing, then R_i and R_j are interchangeable and we are done. Otherwise we define for each P_k the vertex v_k as its endvertex on R_i and w_k

as its endvertex on R_j. Without loss of generality, we assume that v_1, v_2, v_3 occur in order of their index on R_i.

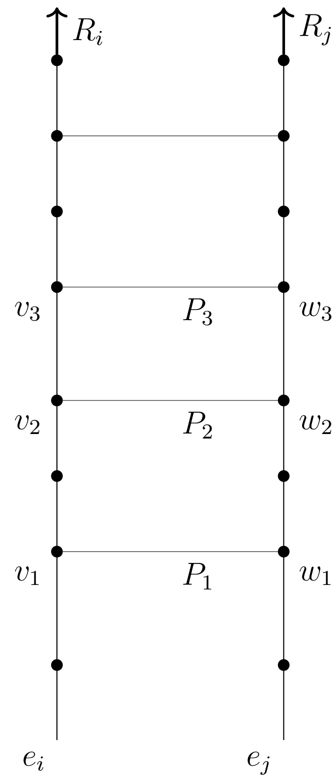

Now we define $s = v_2$ and $T = R_1 \cup R_2 \cup ... \cup R_q \cup P_1 \cup P_3 \setminus \mathring{v}_1 R_i \mathring{v}_3$ and apply Lemma 5.16 for the paths $v_2R_iv_1$, $v_2R_iv_3$ and P_2.

We obtain four edge-disjoint $s-T-$paths, three of them have endvertices v_1, v_3 w_2. Let P_1' be the $v_1 - v_2-$path and P_3' the $v_2 - v_3-$path. Also we define P_2' as the new path with the same endvertices as P_2. The fourth path P_4' has any endvertex $t \in T$. Without loss of generality, we assume that t is the only vertex of this path in T. We define R_i' as a ray starting at e_i contained in $R_i - v_1R_iv_3 + P_1' + P_3'$ and obtain $\mathcal{R}'$ from $\mathcal{R}$ by replacing R_i with R_i'. Since each ray in $\mathcal{R}'$ has a tail in common with the corresponding ray in $\mathcal{R}$, it is clear that $(R1)$ holds. It is left to show that in each possible case, either R_i and R_j are interchangeable or $(R2)$ holds:

- If $t \in (R_1 \cup R_2 \cup ... \cup R_q) \setminus (R_i \cup R_j)$, say $t \in R_k$, then P_4' is an $\mathcal{R}'$-path between R_i and R_k. Further does $P_2' \cup P_4'$ contain an $\mathcal{R}'$-path between R_j and R_k, which implies $(R2)$.

- If $t \in R_j$, then P_2' and P_4' are crossing.

- If $t \in P_1$, then $v_2 P'_4 t P_1 w_1$ and P'_2 are crossing. The case $t \in P_3$ is analogous.
- If $t \in R_i v_1$, then $P'_4 \cup P'_2$ and P_1 are crossing. The case $t \in v_3 R_i$ is analogous.

□

Proposition 5.18. *Every boundary-linked vertex set B with boundary of size at most 6 in a 4-connected multigraph is strongly boundary-linked.*

Proof. The case of an even boundary size q is already shown in Proposition 5.11. Further it is clear for $q = 3$, so it remains to show the proposition for $q = 5$.

Let again $\{e_1, e_2, e_3, e_4, e_5\}$ be the boundary of B and $\mathcal{R} = \{R_1, R_2, R_3, R_4, R_5\}$ be any fixed boundary-linking-set, each R_i with initial edge e_i.

Further let L be the lifting graph on $\{e_1, e_2, e_3, e_4, e_5\}$. From Theorem 5.7 follows that L is either connected it has two components, one of them being a singleton and the other one a complete multipartite graph. Together with Theorem 5.6 we conclude in the second case that L consists of a singleton and a 4-cycle.

It remains to find a path P in B between an admissible pair of edges such that the remaining three boundary-edges can still be connected in $B - P$.

If L and $M(\mathcal{R})$ have an edge in common, we can define P as a passage and are done.

Claim 1. *There is an edge $e_i e_j$ in L such that e_i and e_j have distance at most 2 in $M(\mathcal{R})$.*

To see this claim, let T be a spanning tree of $M(\mathcal{R})$. Since T has 5 vertices, it is either a path or a 4-star or a 3-star in which one edge is subdivided once. In the second or third case, we conclude directly from Theorem 5.6 that at least one of the admissible pairs has distance of at least 2 in T and hence also in $M(\mathcal{R})$. So lets us have a look at the case in which T is a path. If L is connected, then the middle vertex of the path has at least one neighbor in L, which has distance of at most 2 in T. If not, then the middle vertex is isolated in L and the other 4 vertices form a 4-circle which again implies that two of them have distance of at most 2 in T and hence also in $M(\mathcal{R})$.

Without loss of generality let e_2 and e_4 the two boundary edges according to Claim 1 and e_3 the boundary edge on the shortest $M(\mathcal{R})$-path between them.

If R_2 and R_3 or R_3 and R_4 are interchangeable, then according to Lemma 5.15, we could find another set of rays $\mathcal{R}'$ for which e_2 and e_4 are adjacent in $M(\mathcal{R}')$, so we are done.

In any other case, we assume after possibly applying Lemma 5.17 to the pairs (R_2, R_3) and (R_4, R_3) that there is another ray, say R_1, in $\mathcal{R}$ with $\mathcal{R}-$paths to R_2 and R_3 and another ray R_x with $\mathcal{R}$-paths to R_3 and R_4. We distinguish the two cases $R_x = R_1$ and $R_x = R_5$.

We name the vertices and edges of R_3 as follows:

$$V(R_3) = \{v_0, v_1, v_2, ...\},$$

$$E(R_3) = \{v_0v_1, v_1v_2, ...\},$$

If $R_x = R_5$, then let P_1 be an $\mathcal{R}$-path between R_1 and R_3 and P_5 be an $\mathcal{R}$-path between R_5 and R_3. Let v_m and v_n be the endvertices on R_3 from these paths with $m \leq n$. Now since $e_2e_3, e_4e_3 \in M(\mathcal{R})$, there are infinitely many edge-disjoint $\mathcal{R}$-paths between R_2 and R_3 and R_4 and R_3. Since none of those pairs are interchangeable, also their endvertices on R_3 are different. Choose such paths P_2 and P_4 edge-disjoint from $R_3v_n \cup P_1 \cup P_5$ and with endvertices v_o and v_p on R_3 such that $m \leq n < o \leq p$. Now the path contained in $R_2 \cup R_4 \cup P_2 \cup v_oR_3v_p \cup P_4$ between e_2 and e_4 is the desired path since the remaining three boundary edges can still be connected through $R_3v_n \cup P_1 \cup P_5 \cup R_1 \cup R_5$.

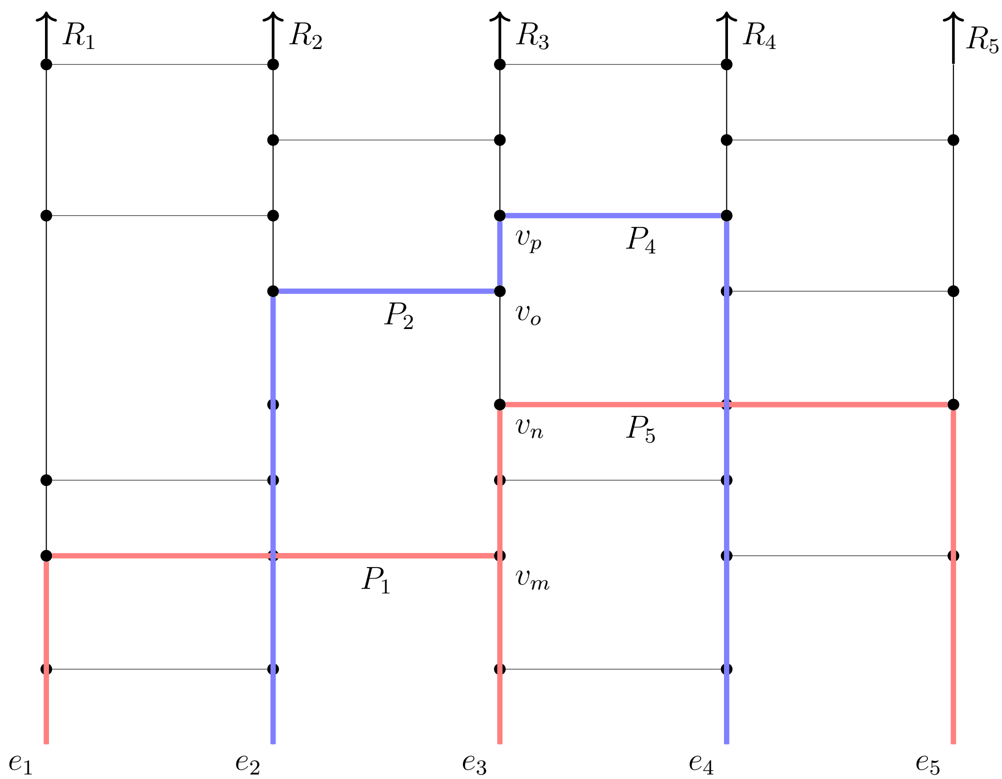

If $R_x = R_1$, then we look again at the lifting graph L. Again we assume that there is no $\mathcal{R}$-path between two rays from which their initial edges are adjacent in L. This implies that $e_1e_2, e_1e_3, e_1e_4, e_2e_3, e_3e_4 \notin E(L)$, thus $d_L(e_1) \leq 1$ and $d_L(e_3) \leq 1$. This excludes the case of L containing a 4-cycle. It is left the case that L is connected. This implies that $e_3e_5, e_1e_5 \in E(L)$, further do we already know that $e_2e_4 \in E(L)$. Since now $d_L(e_1) = d_L(e_3) = 1$, it follows that e_2e_5 or e_4e_5 is an edge of L as well. We assume without loss of generality that $e_2e_5 \in L$. Now since $M(\mathcal{R})$ is connected, e_5 has at least one neighbor in $M(\mathcal{R})$. The only possible neighbor which is not already in $E(L)$ is e_4. Thus R_5 and R_4 cannot be interchangeable (otherwise we could exchance the rays in a way that e_5 had another neighbor). Again applying Lemma 5.17 to the pair (e_5, e_4) gives us another boundary-linked set of rays $\mathcal{R}'$ such that there is a $\mathcal{R}'$-path from the ray starting at e_5 to any other ray then the ray starting at e_4. This completes the proof. $\square$

5.6 Extending orientations of Eulerian subgraphs

And indeed, as our last ingredient, we note that Nash-Williams' orientation theorem also holds in the following, slightly stronger form, improving the bounds from [44, Theorem 6].

For two vertices x, y in a multigraph G, we write $\lambda(x,y)$ for maximum number of edge-disjoint $x-y-$paths, and $\lambda^*(x,y)$ for the greatest even number $\leq \lambda(x,y)$. Further, for two vertices x, y in an oriented multigraph $\vec{G}$ define $\alpha(x,y)$ as the maximum number of edge-disjoint directed $x-y-$paths. Let us say an orientation $\vec{G}$ of a multigraph G is *connectivity preserving*, if

$$\alpha(x,y) \geq \frac{\lambda^*(x,y)}{2}$$

for any two distinct vertices $x, y \in G$.

Theorem 5.19. *Let G be a finite multigraph and $H \subseteq G$ an open or closed Eulerian subgraph. Then any consistent orientation $\vec{H}$ of H can be extended to a connectivity preserving orientation of G.*

Proof. An *odd vertex pairing* of a finite multigraph $G = (V, E)$ is a partition P of the vertices of odd-degree in G into sets of size two. Interpreting P as edges, we obtain an Eulerian multigraph $G' = (V, E')$ where $E' = E \mathbin{\dot\cup} P$. Then $H \subseteq G \subseteq G'$. Nash-Williams showed in [34, Theorem 2] that every multigraph $G = (V, E)$ has an odd-vertex pairing P such that for every two $x, y \in V$ and every bipartition (X, Y) of V with $x \in X$ and $y \in Y$ holds:

$$(\star) \qquad |E(X,Y)| - |P(X,Y)| \geq \lambda^*(x,y).$$

We claim that with such an odd-vertex pairing, any consistent orientation $\vec{G}'$ of the Eulerian multigraph G' that extends $\vec{H}$ restricts to a connectivity preserving orientation $\vec{G}$ of G as desired.

For two vertices a, b with edge-connectivity $\lambda(a,b)$, let (A, B) be a partition of G inducing a minimal edge-cut between a and b. Since $\vec{G}'$ is balanced, it follows that

$$|\vec{E}(A,B)| + |\vec{P}(A,B)| = \frac{|E(A,B)| + |P(A,B)|}{2}.$$

However, since $|\vec{P}(A,B)| \leq |P(A,B)|$, it follows that

$$\alpha(a,b) = |\vec{E}(A,B)| \geq \frac{|E(A,B)| + |P(A,B)|}{2} - |P(A,B)| = \frac{|E(A,B)| - |P(A,B)|}{2} \overset{(\star)}{\geq} \frac{\lambda^*(a,b)}{2}$$

$\square$

5.7 Main results

We are now ready to extend Nash-Williams' orientation theorem to infinite multigraphs with at most countably many ends, from which at most one end has odd degree. As mentioned in the introduction, our method of proof adapts Thomassen's [44, Theorem 7].

Theorem 5.20. *Every $2k$-edge-connected locally finite multigraph with at most countably many ends, from which at most one end has odd degree, has a k-arc-connected orientation.*

Proof. Enumerate $V = \{v_0, v_1, ...\}$. Beginning with $A_0 = \{v_0\}$ and any directed cycle $\vec{W}_0 \subseteq G$ containing v_0, we will construct a sequence of finite, 2-edge connected subgraphs $W_0 \subseteq W_1 \subseteq W_2 \subseteq \cdots$ of G with compatible orientations $\vec{W}_0 \subseteq \vec{W}_1 \subseteq \vec{W}_2 \subseteq \cdots$ and sets of vertices $A_0 \subseteq A_1 \subseteq A_2 \cdots$ such that for all $n \geq 0$:

(i) $\{v_0, \ldots, v_n\} \subseteq A_n \subseteq V(W_n)$.

(ii) For every component B of $G \setminus A_n$, all but possibly at most one exceptional vertex have in-degree equalling out-degree in $\vec{W}_n$, with the exceptional vertex having a difference of 1 between in- and out-degree.

(iii) A_n is k-arc-connected in $\vec{W}_n$.

Once the construction is complete, we claim that properties (i) and (iii) imply that any orientation $\vec{G}$ of G extending $\vec{W} := \bigcup_{i \in \mathbb{N}} \vec{W}_i$ is k-arc-connected. Indeed, for every two distinct vertices x, y in G, by (i) there is an $i \in \mathbb{N}$ with $x, y \in A_i$, and so by (iii) by there are k arc-disjoint directed paths in $\vec{W}_i$ from x to y and from y to x. Since $\vec{W}_i \subseteq \vec{W}$ as oriented subgraphs, these directed paths are directed also in $\vec{W}$, and hence in $\vec{G}$, as desired.

Thus, it remains to describe the inductive construction, and this is where property (ii) is needed. So suppose inductively that we have already constructed A_n and $\vec{W}_n$ according to (i)–(iii). Since G has countably many ends, we may apply Theorem 5.5 to the set $A'_{n+1} := V(W_n) \cup \{v_{n+1}\}$ to obtain a finite set $A_{n+1} \supseteq A'_{n+1}$ such that the components of $G - A_{n+1}$ are boundary-linked sets. Since there is at most one end of odd degree per assumption, we obtain at most one boundary linked set with odd boundary. Applying Proposition 5.13 or Theorem 5.14 yields an empty or one-elemented set X_{n+1} and also an immersion W_{n+1} on $A_{n+1} \cup X_{n+1}$ in G of a finite multigraph H for which A_n is $2k$-edge-connected in H (with $d_H(x) = 3$ for $x \in X_{n+1}$, if x exists). Since each of the paths in W_{n+1} that corresponds to an edge of H is either an edge of W_n or is internally disjoint from W_n, we may assume that $W_n \subseteq W_{n+1}$ and $W_n \subseteq H$.

Now contract A_n in H to a dummy vertex v, and call the resulting multigraph $\tilde{H}$. For each component B of $G \setminus A_n$, let $\tilde{H} \restriction B$ be the subgraph of $\tilde{H}$ induced by the dummy vertex v together with $V(B) \cap V(H)$. Property (ii) implies that all the edges of $\vec{W}_n$ inside B form a consistently oriented (open or closed) Eulerian subgraph of $\tilde{H} \restriction B$. Hence we can apply Theorem 5.19 to each $\tilde{H} \restriction B$ to extend the orientation of this subgraph to a connectivity preserving orientation of all of $\tilde{H} \restriction B$, making A_n k-arc-connected in this orientation. After doing this for every component B of $G - A_n$, we obtain an orientation $\vec{H}$ of H.

We claim that with this orientation, A_n is also k-arc-connected in $\vec{H}$: Indeed, let $E(X,Y)$ be any bond in H. If A_n lies completely on one side X or Y, then the bond restricts to a cut in some $\tilde{H} \restriction B$, and since A_n k-arc-connected in its orientation, there exist at least k edges oriented from X to Y, and also from Y to X. And if A_n meets both X and Y, then the cut restricts to a cut of $\vec{W}_n$ separating two vertices from A_n, and so by (iii) there again exist at least k edges oriented from X to Y, and also from Y to X in $\vec{W}_n$, and hence in $\vec{H}$. Together, it follows from Menger's theorem that A_n is indeed k-arc-connected in $\vec{H}$.

Finally, we now lift this orientation of $\vec{H}$ to an orientation $\vec{W}_{n+1}$ of the immersion W_{n+1} so that $\vec{W}_{n+1}$ satisfies (i)–(iii). Indeed, for each oriented edge in $\vec{H}$, we simply orient the corresponding path in the immersion W_{n+1} accordingly. Then $\vec{W}_n \subseteq \vec{W}_{n+1}$ as directed multigraphs, and (i) holds by construction. To see that property (ii) holds, note that the edges incident with a vertex v in $V(W_{n+1}) \setminus (A_{n+1} \cup X)$ belong to a collection of edge-disjoint, forwards oriented paths containing v in their interior, and hence have equal in- and out-degree. And if for a component B there is the vertex x in $B \cap X$ of degree 3 in H, then since H is 2-edge-connected and $\vec{H}$ is connectivity preserving, it follows that there is at least one ingoing and one outgoing edge at x in H, and so x has a difference of 1 between in- and out-degree in $\vec{W}_{n+1}$. Finally, property (iii) follows at once from the fact that $\vec{W}_{n+1}$ is an immersion of the multigraph $\vec{H}$ in which A_{n+1} was k-arc connected. $\square$

Corollary 5.21. *Every one-ended $2k$-edge-connected locally finite multigraph has a k-arc-connected orientation.*

Theorem 5.22. *Every $4k$-edge-connected multigraph has a k-arc-connected orientation.*

Proof. By Theorem 5.1, only the infinite case is open. Next, Thomassen has shown that every infinite $4k$-edge-connected multigraph has a decomposition into locally finite, $4k$-edge-connected subgraphs [44, §7 & §8]; hence, it suffices to prove the assertion for locally finite multigraphs. Further, by Egyed's result [20], we may assume that $k \geq 2$.

In the last theorem, the restriction to countably many ends, from which at most one end has odd degree, came from the immersion Theorem 5.14 we used. Since we have a $4k$-edge-connected multigraph, we can apply the original immersion Theorem 5.8 from Thomassen instead. Apart from that, we do the same proof as in Theorem 5.20. $\square$

6 English summary

Chapter 3: End spaces and tree-decompositions

Section 3.3: We show that the graphs for which there is a tree-decomposition displaying all ends are exactly the graphs with a normal spanning tree. We further state some topological characterisations.

Section 3.4: We introduce the concept of Envelopes. Envelopes are a powerful tool to find for any subgraph of a graph another subgraph with the same ends in the closure in $|G|$ but with finite adhesion. As a strengthening of the original envelope theorem from Max Pitz, we show that any set consisting of vertices and ends in a connected graph G has a connected, end-faithful envelope.

Section 3.5: We show that under certain topological circumstances, we can find an upwards connected tree-decomposition of finite adhesion with connected parts, which homeomorphically displays a given set of ends, such that the boundary of every part contains at most one end from another given set of ends.

Section 3.6: We characterise in several ways which sets of ends from a given graph can be displayed by a tree-decomposition. The most notably results are that these are the sets Ψ for which $|G|_\Psi$ is completely metrizable, which is also equivalent to the property of Ψ being G_δ in $|G|$.

Section 3.7: We show that a set Ξ of ends in a graph G can be distributed, whenever $V(G) \cup \Xi$ has a σ-discrete expansion in $|G|$.

Section 3.8: We show that the graphs with a tree-decomposition distributing all ends are exactly those graphs for which each end has a rank.

Section 3.9: We deduce from our research the result from Carmesin that every connected graph G has a tree-decomposition of finite adhesion with connected parts that displays precisely the undominated ends of G. Further we give a shorter proof of Carmesins result that no binary tree with uncountably many tops admits a tree-decomposition of finite adhesion distinguishing all its ends.

Chapter 4: Hamilton circles in powers of infinite graphs

Section 4.2.1: We characterize the end-space of third powers of rayless graphs with an equivalence relation on the vertices of infinite degree.

Section 4.2.2: We state and prove a sufficient condition for a rayless graph to have no Hamilton circle in their third power.

Section 4.2.3: We strengthen the original theorem for Hamilton cycles in the third power of finite graphs slightly in a way that allows us more control about some edges.

Section 4.2.4: We construct Hamilton circles for all rayless graphs with one end in their third power, which are not a counterexample as characterized before. Further we keep the control about some edges to use the result from this section for a recursive construction.

Section 4.2.5: We construct a Hamilton circle for all rayless graphs that remain after eliminating all possible counterexamples. This leads to a final characterisation.

Section 4.3.1: We characterize the end-space of fourth and higher powers of infinite trees with an equivalence relation of the vertices of infinite degree.

Section 4.3.2: We prove with a recursive construction that all fourth and higher powers of infinite trees have a Hamilton circle.

Chapter 5: Nash Williams orientation theorem for infinite graphs

Section 5.2: We show that we can find for any finite set of vertices A' in a connected locally finite multigraph G with countably many ends a finite superset A, such that the components of $G - A$ are boundary linked sets.

Section 5.3: We introduce Mader's lifting theorem and state some results about the lifting graph from Ok, Richter and Thomassen.

Section 5.4: We introduce the concept of immersions and show that we can find for a vertex set A in a multigraph G under certain circumstances a special immersion that reflects the original edge-connectivity in A.

Section 5.5: We give a construction for a linking of boundary-linked sets with boundary of size 5 as needed for the special immersion defined in the section before.

Section 5.6: We show that we can extend any open or closed eulerian subgraph in a finite multigraph to a connectivity preserving orientation of the whole graph.

Section 5.7: We show that Nash-Williams' orientation theorem holds for locally finite multigraphs with at most countably many ends, from which at most one end has odd degree. Further we show that every $4k$-edge-connected multigraph has a k-arc-connected orientation.

7 Deutsche Zusammenfassung

Kapitel 3: End spaces and tree-decompositions

Abschnitt 3.3: Wir zeigen, dass die Graphen mit einer Baumzerlegung, die alle Enden *darstellt* (displayed), genau die Graphen sind, die einen normalen Spannbaum haben. Zusätzlich beweisen wir einige topologische Charakterisierungen.

Abschnitt 3.4: Wir geben eine Einführung in das Konzept der Envelopes. Envelopes sind ein wesentliches Werkzeug, um zu jedem Teilgraphen eines Graphen einen weiteren Teilgraphen mit endlicher Adhäsion zu finden, der die selben Enden in seinem topologischen Abschluss hat. Aufbauend auf dem ursprünglichen Envelope Theorem von Max Pitz zeigen wir, dass jede Menge von Ecken und Enden in einem zusammenhängenden Graphen einen zusammenhängenden, endentreuen Envelope hat.

Abschnitt 3.5: Wir zeigen, dass unter gewissen topologischen Umständen eine Baumzerlegung mit endlicher Adhäsion und zusammenhängenden Verzweigungsmengen existiert, die *aufsteigend-zusammenhängend* (upwards connected) ist, eine vorgegebene Endenmenge homöomorph darstellt und in der in jeder Verzweigungsmenge höchstens ein weiteres Ende lebt.

Abschnitt 3.6: Wir charakterisieren die Mengen von Enden eines gegebenen Graphen, die durch eine Baumzerlegung dargestellt werden können, auf verschiedene Weisen. Insbesondere zeigen wir, dass dies die Mengen von Enden Ψ sind, für die $|G|_\Psi$ vollständig metrisierbar ist. Dies ist ebenso äquivalent dazu, dass Ψ eine G_δ-Menge in $|G|$ ist.

Abschnitt 3.7: Wir zeigen, dass eine Menge von Enden Ξ in einem Graphen G *verteilt* (distributed) werden kann, falls $V(G) \cup \Xi$ eine σ-diskrete Expansion in $|G|$ hat.

Abschnitt 3.8: Wir zeigen, dass die Graphen mit einer Baumzerlegung, die alle Enden verteilt, genau die Graphen sind, für die jedes Ende einen Rang hat.

Abschnitt 3.9: Wir folgern das Theorem von Carmesin, dass jeder Graph G eine Baumzerlegung mit endlicher Adhäsion und zusammenhängenden Zerlegungsmengen hat, die genau die undominierten Enden darstellt. Weiterhin beweisen wir, dass kein Binärbaum mit überabzählbar vielen Tops eine Baumzerlegung mit endlicher Adhäsion hat, die alle Enden unterscheidet.

Kapitel 4: Hamilton circles in powers of infinite graphs

Abschnitt 4.2.1: Wir charakterisieren den Endenraum der dritten Potenz von strahlenlosen Graphen mithilfe einer Äquivalenzrelation auf den Ecken von unendlichem Grad.

Abschnitt 4.2.2: Wir definieren und beweisen eine hinreichende Bedingung dafür, dass ein strahlenloser Graph keinen Hamiltonkreis in seiner dritten Potenz hat.

Abschnitt 4.2.3: Wir verstärken das ursprüngliche Theorem über Hamiltonkreise in der dritten Potenz von endlichen Graphen auf eine Weise, die uns mehr Kontrolle darüber erlaubt, welche Ecken für einen Hamiltonkreis verwendet werden.

Abschnitt 4.2.4: Wir konstruieren Hamiltonkreise für alle strahlenlosen Graphen mit nur einem Ende in der dritten Potenz, die nicht zu den vorher charakterisierten Gegenbeispielen gehören. Außerdem können wir in unserem Hamiltonkreis einige Kanten festlegen, die wir später für eine rekursive Konstruktion benötigen.

Abschnitt 4.2.5: Wir konstruieren einen Hamiltonkreis für alle strahlenlosen Graphen, die nicht zu den vorher charakterisierten Gegenbeispielen gehören. Damit ist unsere finale Charakterisierung vollständig.

Abschnitt 4.3.1: Wir charakterisieren den Endenraum von der vierten und höheren Potenzen von strahlenlosen Graphen mithilfe einer Äquivalenzrelation auf den Ecken von unendlichem Grad.

Abschnitt 4.3.2: Wir beweisen mit einer rekursiven Konstruktion, dass alle vierten und höheren Potenzen von abzählbaren Bäumen einen Hamiltonkreis haben.

Kapitel 5: Nash Williams' orientation theorem for infinite graphs

Abschnitt 5.2: Wir zeigen, dass wir für jede endliche Menge A' in einem zusammenhängenden, lokal endlichen Graphen G mit abzählbar vielen Enden eine endliche Obermenge A finden, sodass die Komponenten von $G - A$ boundary linked sind.

Abschnitt 5.3: Wir stellen Mader's lifting Theorem vor und geben einige Aussagen über den Lifting Graph von Ok, Richter und Thomassen an.

Abschnitt 5.4: Wir definieren Immersionen und zeigen, dass wir für eine Eckenmenge A in einem Multigraphen G unter gewissen Umständen eine spezielle Immersion finden, die den Kantenzusammenhang in A erhält.

Abschnitt 5.5: Wir konstruieren für Mengen, die *boundary-linked* mit Boundary der Größe 5 sind ein linking in dem Sinne, wie es zur Konstruktion der speziellen Immersion gebraucht wird.

Abschnitt 5.6: Wir zeigen, dass wir jede offene oder geschlossene Eulertour in einem endlichen Multigraphen zu einer kantenzusammenhangserhaltenden Orientierung des gesamten Graphens ergänzen können.

Abschnitt 5.7: Wir zeigen Nash-Williams' orientation Theorem für lokal endliche Multigraphen mit höchstens abzählbar vielen Enden, von denen höchstens eines ungeraden Grad hat. Weiterhin zeigen wir, dass jeder $4k$-kanten-zusammenhängende Multigraph eine k-arc-zusammenhängende Orientierung hat.

8 Declaration on my contributions

Chapter 3: This chapter is based on a paper of Thilo Krill, Max Pitz and me.

Building on previous work of Max Pitz ([38]), most of the detailed research was done by Thilo and me together in a great number of brainstorming sessions. I drafted Sections 3.3, 3.4 and parts of the sections 3.6, 3.7, 3.8

Chapter 4: I created this entire chapter on my own, with the exception of Section 4.3, in which the ideas are from joint work with Joshua Erde, Pascal Gollin and Max Pitz, but the draft is my own.

Chapter 5: This chapter is joint work with Max Pitz. I drafted the sections 5.2, 5.3, 5.4, 5.5 and parts of the sections 5.6 and 5.7.

References

[1] R. Aharoni and C. Thomassen, *Infinite, highly connected digraphs with no two arc-disjoint spanning trees*, Journal of graph theory **13** (1989), no. 1, 71–74.

[2] A. Assem, *Analysis of the lifting graph*, arXiv preprint arXiv:2212.03347 (2022).

[3] ———, *Towards nash-williams orientation conjecture for infinite graphs*, arXiv preprint arXiv:2306.10631 (2023).

[4] J. Barát and M. Kriesell, *What is on his mind?*, Discrete Mathematics **310** (2010), no. 20, 2573–2583.

[5] N. Bowler and J. Carmesin, *Infinite matroids and determinacy of games*, arXiv preprint arXiv:1301.5980 (2013).

[6] ———, *The ubiquity of psi-matroids*, arXiv preprint arXiv:1304.6973 (2013).

[7] C. Bürger and J. Kurkofka, *Duality theorems for stars and combs I: Arbitrary stars and combs*, Journal of Graph Theory **99** (2022), no. 4, 525–554.

[8] ———, *Duality theorems for stars and combs II: Dominating stars and dominated combs*, Journal of Graph Theory **99** (2022), no. 4, 555–572.

[9] J. Carmesin, *All graphs have tree-decompositions displaying their topological ends*, Combinatorica **39** (2019), no. 3, 545–596.

[10] J. Carmesin, *Topological cycle matroids of infinite graphs*, European Journal of Combinatorics **60** (2017), 135–150.

[11] R. Diestel, *The end structure of a graph: recent results and open problems*, Discrete Mathematics **100** (1992), no. 1–3, 313–327.

[12] ———, *The depth-first search tree structure of $TK_{\aleph_0}$-free graphs*, Journal of Combinatorial Theory, Series B **61** (1994), no. 2, 260–262.

[13] ———, *Locally finite graphs with ends: A topological approach, I–III.*, Discrete Mathematics **310–312** (2010/2011/2012).

[14] ———, *Graph Theory*, 5th ed., Springer, 2015.

[15] R. Diestel and D. Kühn, *Graph-theoretical versus topological ends of graphs*, Journal of Combinatorial Theory, Series B **87** (2003), no. 1, 197–206.

[16] R. Diestel, *End spaces and spanning trees*, Journal of Combinatorial Theory, Series B **96** (2006), 846–854.

[17] ———, *Locally finite graphs with ends: A topological approach, i. basic theory*, Discrete Mathematics **311** (2011), no. 15, 1423 –1447. Infinite Graphs: Introductions, Connections, Surveys.

[18] R. Diestel and I. Leader, *Normal spanning trees, Aronszajn trees and excluded minors*, Journal of the London Mathematical Society **63** (2001), no. 1, 16–32.

[19] R. Diestel and J. Pott, *Dual trees must share their ends*, Journal of Combinatorial Theory, Series B **123** (2017), 32–53.

[20] L. Egyed, *Ueber die wohlgerichteten unendlichen Graphen*, Math. phys. Lapok **48** (1941), 505–509.

[21] A. Elm and J. Kurkofka, *A tree-of-tangles theorem for infinite tangles*, Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg **92** (2022), 139–178.

[22] R. Engelking, *General toplogy – revised and completed ed.*, Vol. 6, Heldermann Verlag Berlin, 1989.

[23] A. Frank, *On a theorem of Mader*, Discrete Mathematics **101** (1992), 49–57.

[24] A. Georgakopoulos, *Infinite hamilton cycles in squares of locally finite graphs*, Advances in Mathematics **220** (2009), no. 3, 670–705.

[25] L. Jannasch, *Eine topologische Erweiterung des Orientierungs-Theorems von Nash-Williams auf lokal endliche Graphen*, BSc thesis, University of Hamburg (2019).

[26] H. A. Jung, *Wurzelbäume und unendliche Wege in Graphen*, Math. Nachr. **41** (1969), 1–22.

[27] J. Kurkofka and R. Melcher, *Countably determined ends and graphs* **156** (2022), 31–56. Journal of Combinatorial Theory, Series B.

[28] J. Kurkofka, *The farey graph is uniquely determined by its connectivity*, Journal of Combinatorial Theory, Series B **151** (2021), 223–234.

[29] J. Kurkofka, R. Melcher, and M. Pitz, *Approximating infinite graphs by normal trees*, Journal of Combinatorial Theory, Series B **148** (2021), 173–183.

[30] J. Kurkofka and M. Pitz, *A representation theorem for end spaces of infinite graphs*, arXiv preprint arXiv:2111.12670 (2023).

[31] W. Mader, *A reduction method for edge-connectivity in graphs*, Annals of Discrete Mathematics **3** (1978), 145–164.

[32] B. Mohar, *Hamiltonicity of infinite graphs*, 2007. `https://users.fmf.uni-lj.si/mohar/Problems/P0703_HamiltonicityInfinite.html`.

[33] M.Sekanina, *On an ordering of the set of vertices of a connected graph*, Publ. Fac. Sc. Univ. Brno (1960), 412:137 –142.

[34] C. St. J. A. Nash-Williams, *On orientations, connectivity and odd-vertex-pairings in finite graphs*, Canadian Journal of Mathematics **12** (1960), 555–567.

[35] ———, *Infinite graphs—a survey*, Journal of Combinatorial Theory **3** (1967), no. 3, 286–301.

[36] S. Ok, B. Richter, and C. Thomassen, *Liftings in finite graphs and linkages in infinite graphs with prescribed edge-connectivity*, Graphs and Combinatorics **32** (2016), no. 6, 2575–2589.

[37] M. Pitz, *Proof of Halin's normal spanning tree conjecture*, Israel Journal of Mathematics **246** (2021), no. 1, 353–370.

[38] ———, *Constructing tree-decompositions that display all topological ends*, Combinatorica **42** (2022), 763–769.

[39] N. Polat, *Ends and multi-endings, I*, Journal of Combinatorial Theory, Series B **67** (1996), 86–110.

[40] ———, *Ends and multi-endings, II*, Journal of Combinatorial Theory, Series B **68** (1996), 56–86.

[41] W. Sierpiński, *Sur une propriété topologique des ensembles dénombrables denses en soi*, Fundamenta Mathematicae **1** (1920), no. 1, 11–16.

[42] R. Thomas, *Well-quasi-ordering infinite graphs with forbidden finite planar minor*, Transactions of the American Mathematical Society **312** (1989), no. 1, 279–313.

[43] C. Thomassen, *Configurations in graphs of large minimum degree, connectivity, or chromatic number*, Proceedings of the third international conference on combinatorial mathematics, 1989, pp. 402–412.

[44] ———, *Orientations of infinite graphs with prescribed edge-connectivity*, Combinatorica **36** (2016), no. 5, 601–621.

9 Eidesstattliche Versicherung

Hiermit versichere ich an Eides statt, die vorliegende Dissertation selbst verfasst und keine anderen als die angegebenen Hilfsmittel benutzt zu haben. Darüber hinaus versichere ich, dass diese Dissertation nicht in einem früheren Promotionsverfahren eingereicht wurde.

10 Acknowledgement

I am very grateful to my supervisor Max Pitz for his support in my research over the past years and also for the one or two portions of wisdom that he gave me along the way.
Furthermore, I would like to thank my colleague Thilo Krill for the intensive and pleasant collaboration in the last year of my work at the university. We always had a positive working environment in our shared office with Florian Gut and Florian Reich. I will miss our time together.
Also I want to thank Nicola Lorenz for decorating our Office, making it an even more pleasant place to be.
Thank you to everyone working in the discrete mathematics at the University of Hamburg, who motivated me along the way. Working in a group of so many extraordinary high skilled mathematicians was inspiring.
During the time I was writing this dissertation, I experienced many challenges in my personal life. I would not have overcome many of them so well without the support of my friends. Thank you to every person who makes me feel like I am not alone in the world. Especially, I would like to thank my best friend Sina Martensen, always being there for me, whenever I needed her the most.