Limit structures and ubiquity in finite and infinite graphs

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

vorgelegt
im Fachbereich Mathematik
von

Julian Pott
aus Buchholz in der Nordheide

Hamburg
2014

Als Dissertation angenommen vom Fachbereich
Mathematik der Universität Hamburg

auf Grund der Gutachten von Prof. Dr. Reinhard Diestel
und Prof. Dr. Stephan Kreutzer

Hamburg, den 27.1.2015

Prof. Dr. Michael Hinze
Sprecher des Fachbereichs Mathematik

Contents

Dual trees must share their ends

Reinhard Diestel Julian Pott

Abstract

We extend to infinite graphs the matroidal characterization of finite graph duality, that two graphs are dual iff they have complementary spanning trees in some common edge set. The naive infinite analogue of this fails.

The key in an infinite setting is that dual trees must share between them not only the edges of their host graphs but also their ends: the statement that a set of edges is acyclic and connects all the vertices in one of the graphs iff the remaining edges do the same in its dual will hold only once each of the two graphs' common ends has been assigned to one graph but not the other, and 'cycle' and 'connected' are interpreted topologically in the space containing the respective edges and precisely the ends thus assigned.

This property characterizes graph duality: if, conversely, the spanning trees of two infinite graphs are complementary in this end-sharing way, the graphs form a dual pair.

1 Introduction

It is well known (and not hard to see) that two finite graphs are dual if and only if they can be drawn with a common abstract set of edges so that the edge sets of the spanning trees of one are the complements of the edge sets of the spanning trees of the other:

Theorem 1. *Let $G = (V, E)$ and $G^* = (V^*, E)$ be connected finite graphs with the same abstract edge set. Then the following statements are equivalent:*

(i) *G and G^* are duals of each other.*

(ii) *Given any set $F \subseteq E$, the graph (V, F) is a tree if and only if $(V^*, F^{\complement})$ is a tree.*

For infinite dual graphs G and G^* (see [1]), Theorem 1 (ii) will usually fail: when (V, F) is a spanning tree of G, the subgraph $(V^*, F^{\complement})$ of G^* will be acyclic but may be disconnected. For example, consider as G the infinite $\mathbb{Z} \times \mathbb{Z}$ grid, and let F be the edge set of any spanning tree containing a two-way infinite path, a *double ray* R. Then the edges of R will form a cut in G^*, so $(V^*, F^{\complement})$ will be disconnected.

Although the graphs $(V^*, F^{\complement})$ in this example will always be disconnected, they become arc-connected (but remain *acirclic*) when we consider them as closed subspaces of the topological space obtained from G^* by adding its end. Such subspaces are called *topological spanning trees*; they provide the 'correct' analogues in infinite graphs of spanning trees in finite graphs for numerous problems, and have been studied extensively [6, 7]. For $G = \mathbb{Z} \times \mathbb{Z}$, then, the complements of the edge sets of ordinary spanning trees of G form topological spanning trees in G^*, and vice versa (as $\mathbb{Z} \times \mathbb{Z}$ is self-dual).

It was shown recently in the context of infinite matroids [2] that this curious phenomenon is not specific to this example but occurs for all dual pairs of graphs: neither ordinary nor topological spanning trees permit, by themselves, an extension of Theorem 1 to infinite graphs, but as soon as one notion is used for G and the other for G^*, the theorem does extend. The purpose of this paper is to explain this seemingly odd phenomenon by a more general duality for graphs with ends, in which it appears as merely a pair of extreme cases.

It was shown in [3] that 2-connected dual graphs do not only have the 'same' edges but also the 'same' ends: there is a bijection between their ends that commutes with the bijection between their edges so as to preserve convergence of edges to ends. Now if G and G^* are dual 2-connected graphs with edge sets E and end sets Ω, our result is that if we specify *any* subset Ψ of Ω and consider topological spanning trees of G in the space obtained from G by adding only the ends in Ψ, then Theorem 1 (ii) will hold if the subgraphs $(V^*, F^{\complement})$ of G^* are furnished with precisely the ends in $\Omega \setminus \Psi$. (Our earlier example is the special case of this result with either $\Psi = \emptyset$ or $\Psi = \Omega$.) And conversely, if the spanning trees of two graphs G and G^* with common edge and end sets complement each other in this way for some—equivalently, for every—subset Ψ of their ends then G and G^* form a dual pair.

Here, then, is the formal statement of our theorem. A graph G is *finitely separable* if any two vertices can be separated by finitely many edges; as noted by Thomassen [9, 10], this slight weakening of local finiteness is necessary for any kind of graph duality to be possible. The Ψ*-trees* in G, for subsets Ψ of its ends, will be defined in Section 2. Informally, they are the subgraphs that induce no cycle or topological circle in the space which G forms with the ends in Ψ (but no other ends) and connect any two vertices by an arc in this space.

Theorem 2. *Let $G = (V, E, \Omega)$ and $G^* = (V^*, E, \Omega)$ be finitely separable 2-connected graphs with the same edge set E and the same end set Ω, in the sense of* [3]. *Then the following assertions are equivalent:*

(i) *G and G^* are duals of each other.*

(ii) *For all $\Psi \subseteq \Omega$ and $F \subseteq E$ the following holds: F is the edge set of a Ψ-tree in G if and only if $F^{\complement}$ is the edge set of a $\Psi^{\complement}$-tree in G^*.*

(iii) *There exists a set $\Psi \subseteq \Omega$ such that for every $F \subseteq E$ the following holds: F is the edge set of a Ψ-tree in G if and only if $F^{\complement}$ is the edge set of a $\Psi^{\complement}$-tree in G^*.*

Setting $\Psi = \emptyset$ in (ii) and (iii) as needed, we reobtain the following result from [2]:

Corollary 3. *Two 2-connected and finitely separable graphs $G = (V, E, \Omega)$ and $G^* = (V^*, E, \Omega)$ are dual if and only if the following assertions are equivalent for every $F \subseteq E$:*

(i) *F is the edge set of a spanning tree of G;*

(ii) *$F^{\complement}$ is the edge set of a topological spanning tree of G^*.* □

We shall prove Theorem 2, extended by another pair of equivalent conditions in terms of circuits and bonds, in Sections 3–4.

2 Definitions and basic facts

All the graphs we consider in this paper will be *finitely separable*, that is, any two vertices can be separated by finitely many edges.

We think of a *graph* as a triple (V, E, Ω) of disjoint sets, of *vertices*, *edges*, and *ends*, together with a map $E \to V \cup [V]^2$ assigning to every edge either one or two vertices, its *endvertices*, and another map mapping the ends bijectively to the equivalence classes of *rays* in the graph, its 1-way infinite paths, where two rays are *equivalent* if they cannot be separated by finitely many vertices. In particular, our 'graphs' may have multiple edges and loops. For the complement of F in E, and of Ψ in Ω, we write $F^{\complement}$ and $\Psi^{\complement}$, respectively.

Let $G = (V, E, \Omega)$ be a graph, and let X be the topological 1-complex formed by its vertices and edges. In X, every edge is a topological copy of $[0, 1]$ inheriting also its metric. We denote the topological interior of an edge e by $\mathring{e}$, and for a set $F \subseteq E$ of edges we write $\mathring{F} := \bigcup_{e \in F} \mathring{e}$.

Let us define a new topology on $X \cup \Omega$, to be called VTop. We do this by specifying a neighbourhood basis for every point. For points $x \in X$ we declare as open the open ϵ-balls around x in X with $0 < \epsilon < \delta$, where δ is the distance from x to a closest vertex $v \neq x$. For points $\omega \in \Omega$, note that for every finite set $S \subseteq V$ there is a unique component $C = C(S, \omega)$ of $G - S$ that contains a ray from ω. Let $\hat{C} = \hat{C}(S, \omega) \subseteq X \cup \Omega$ be the set of all the vertices and inner points of edges contained in or incident with C, and of all the ends represented by a ray in C. We declare all these sets $\hat{C}$ as open, thus obtaining for ω the neighbourhood basis

$$\big\{\hat{C}(S, \omega) \subseteq X \cup \Omega : S \subseteq V,\ |S| < \infty\big\}.$$

We write $|G|$ for the topological space on $X \cup \Omega$ endowed with this topology.[1] In topological contexts we shall also write G for the subspace $|G| \setminus \Omega$. (This has the same points as X, but a different topology unless G is locally finite.)

If ω and S are as above, we say that S *separates* ω in G from all the ends that have no ray in $C(S, \omega)$ and from all vertices in $G - C(S, \omega) - S$.

[1] This differs a little from the definition of $|G|$ in [5] when G is not locally finite.

A vertex v *dominates* an end ω if G contains infinitely many paths from v to some ray in ω that pairwise meet only in v. When this is the case we call v and ω *equivalent*; let us write $\sim$ for the equivalence relation on $V \cup \Omega$ which this generates. Note that since G is finitely separable, no two vertices will be equivalent under $\sim$: every non-singleton equivalence class consists of one vertex and all the ends it dominates. A vertex and an end it dominates have no disjoint neighbourhoods in $|G|$. But two ends always have disjoint neighbourhoods, even if they are dominated by the same vertex.

For sets $\Psi \subseteq \Omega$ of ends, we shall often consider the subspace

$$|G|_\Psi := |G| \setminus \Psi^{\complement}$$

and its quotient space

$$\widetilde{G}_\Psi := |G|_\Psi/\!\sim ,$$

whose topology we denote by Ψ-Top. For $\Psi = \Omega$ we obtain an identification space

$$\widetilde{G} := \widetilde{G}_\Omega$$

that readers may have met before; its topology is commonly denoted as ITop. We usually write $[x]_\Psi$ for the equivalence class of x in $|G|_\Psi$, and $[x]$ for its class in $\widetilde{G}$.

As different vertices are never equivalent, the vertices of G determine distinct $\sim$-classes, which we call the *vertices* of $\widetilde{G}_\Psi$. All other points of $\widetilde{G}_\Psi$ are singleton classes $\{x\}$, with x either an inner point of an edge or an undominated end in Ψ. We will not always distinguish $\{x\}$ from x in these cases, i.e., call these x also *inner point of edges* or *ends* of $\widetilde{G}_\Psi$.

Note that if Ψ contains a dominated end then $|G|_\Psi$ will fail to be Hausdorff, and if $\Psi^{\complement} \neq \emptyset$ then $\widetilde{G}_\Psi$ will fail to be compact. But we shall see that $\widetilde{G}_\Psi$ is always Hausdorff (Corollary 7), and if G is 2-connected then $\widetilde{G}$ is compact [4].

Rather than thinking of $\widetilde{G}_\Psi$ as a quotient space as formally defined above, we may think of it informally as formed from the topological space G in three steps:

- add the undominated ends from Ψ as new points, and make their rays converge to them;
- make the rays from any dominated end in Ψ converge to their unique dominating vertex;
- let the rays of ends in $\Psi^{\complement}$ go to infinity without converging to any point.

The diagram in Figure 1 shows the relationship between the spaces just defined. The subspace inclusion $\iota\colon |G|_\Psi \to |G|$ and the quotient projections $\pi\colon |G| \to \widetilde{G}$ and $\pi_\Psi\colon |G|_\Psi \to \widetilde{G}_\Psi$ are canonical, and $\sigma_\Psi\colon \widetilde{G}_\Psi \to \widetilde{G}$ is defined so as to make the diagram commute: it sends an equivalence class $[x]_\Psi \in \widetilde{G}_\Psi$ to the class $[x] \in \widetilde{G}$ containing it.

$$\begin{array}{ccc} |G|_\Psi & \xrightarrow{\iota} & |G| \\ {\scriptstyle \pi_\Psi}\downarrow & & \downarrow{\scriptstyle \pi} \\ \widetilde{G}_\Psi & \xrightarrow[\sigma_\Psi]{} & \widetilde{G} \end{array}$$

Figure 1: Spaces with ends, and their quotient spaces

Since G is finitely separable and hence no end is dominated by more than one vertex, σ_Ψ is injective: $\sigma_\Psi([x]_\Psi) = [x] \in \widetilde{G}$ is obtained from $[x]_\Psi$ simply by adding those ends of $\Psi^{\complement}$ that are dominated by a vertex in $[x]_\Psi$. As $|G|_\Psi$ carries the subspace topology induced from $|G|$, it is also easy to check that σ_Ψ is continuous. Its inverse σ_Ψ^{-1} can fail to be continuous; see Example 2 below.

The subtle differences between $|G|_\Psi$ and $\widetilde{G}_\Psi$ will often be crucial in this paper. But when they are not, we may suppress them for simplicity of notation. For example, given a subgraph H of G we shall speak of the *closure of H in $\widetilde{G}_\Psi$* and mean the obvious thing: the closure in $\widetilde{G}_\Psi$ of its subspace $\pi_\Psi(H')$, where H' is H viewed as a subspace of $|G|_\Psi \subseteq |G|$.

By a *circle* in a topological space X we mean a topological embedding $S^1 \to X$, or its image. Since circles are compact and $\widetilde{G}$ is Hausdorff, σ_Ψ maps circles in $\widetilde{G}_\Psi$ to circles in $\widetilde{G}$. Conversely, circles in $\widetilde{G}$ that use only ends in Ψ define circles in $\widetilde{G}_\Psi$; this will be shown in Lemma 11. The set of all the edges contained in a given circle in $\widetilde{G}_\Psi$ will be called a Ψ-*circuit* of G; for $\Psi = \Omega$ we just speak of *circuits* of G. We shall not consider 'circuits' of circles in $|G|$ or $|G|_\Psi$.

As with circles, we use the term *path* in topological contexts both for continuous maps from $[0,1]$, not necessarily injective, and for their images. For example, if A and B are the images of paths $\varphi, \varphi' \colon [0,1] \to \widetilde{G}$ with endpoints $x = \varphi(0)$ and $y = \varphi(1) = \varphi'(0)$ and $z = \varphi'(1)$, we write $xAyBz$ for the 'x–y path' in $\widetilde{G}$ that is the image of the concatenation of the paths φ and φ'. Note that, since $\widetilde{G}_\Psi$ is Hausdorff, every path in $\widetilde{G}_\Psi$ between two points x and y contains an x–y arc [8, p. 208].

A subspace of $\widetilde{G}_\Psi$ that is the closure in $\widetilde{G}_\Psi$ of the union of all the edges it contains is a *standard subspace of* $\widetilde{G}_\Psi$. Circles in $\widetilde{G}_\Psi$ are examples of standard subspaces; this was shown in [7] for $\widetilde{G}$, and follows for arbitrary Ψ from Lemma 6 below. A standard subspace of $\widetilde{G}_\Psi$ that contains no circle is a Ψ-*forest* of G. A Ψ-forest is *spanning* if it contains all the vertices of $\widetilde{G}_\Psi$. Note that, being closed, it then also contains all the ends of $\widetilde{G}_\Psi$. A spanning arc-connected Ψ-forest of G is a Ψ-*tree* of G.

Thus, the $\emptyset$-trees of G are precisely its (ordinary) spanning trees, while its Ω-trees are its *topological spanning trees*, the arc-connected standard subspaces of $\widetilde{G}$ that contain all the vertices of G but no topological circle.

Example 1. Let G be obtained from a double ray D by adding a vertex v adjacent to all of D. This graph G has two ends, ω and ψ say, both dominated by v. The closure in $\widetilde{G}$ of the edges of D is a circle containing the 'vertex'

$[v] = \{v, \omega, \psi\}$ of $\widetilde{G}$, even though v does not lie on D. However for $\Psi = \{\psi\}$ the closure in $\widetilde{G}_\Psi$ of the same set of edges is not a circle but homeomorphic to a half-open interval. It thus is a Ψ-tree, and even a spanning one, since v and ψ are both elements of its 'vertex' $\{v, \psi\}$ and it also contains all the other vertices of G. The closure of the edges of D in $\widetilde{G}_\emptyset$, on the other hand, is a $\emptyset$-tree but not a spanning one, since v lies in none of its points. Figure 2 shows a Ψ-tree for each choice of Ψ in this example.

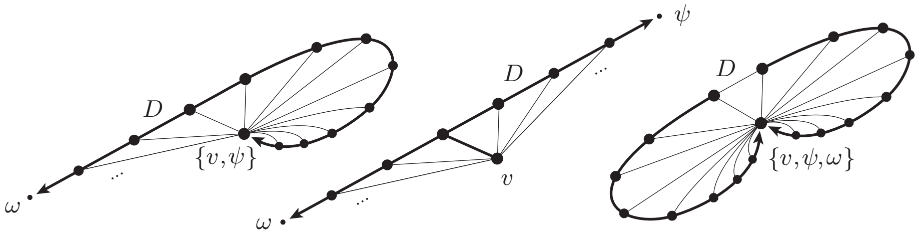

Figure 2: Ψ-trees for $\Psi = \{\psi\}$, $\Psi = \emptyset$ and $\Psi = \{\omega, \psi\}$

If G and G^* are graphs with the same edge set, and such that the bonds of G^* are precisely the circuits of G, then G^* is called a *dual* of G. If the *finite* bonds of G^* are precisely the finite circuits of G, then G^* is a *finitary dual* of G. Clearly, duals are always finitary duals. For finitely separable graphs, as considered here, the converse is also true [1, Lemmas 4.7–4.9]. If G^* is a dual of G, then G is a dual of G^* [1, Theorem 3.4]. Finally, G has a dual if and only if it is planar [1].

3 Lemmas

Our main aim in this section is to prove some fundamental lemmas about the spaces $|G|$, $|G|_\Psi$, $\widetilde{G}$ and $\widetilde{G}_\Psi$ defined in Section 2: about their topological properties, and about their relationship to each other. Throughout the section, let $G = (V, E, \Omega)$ be a fixed finitely separable graph, and $\Psi \subseteq \Omega$ a fixed set of ends.

Before we get to these topological fundamentals, let us show that Ψ-trees always exist, and prove an easy lemma about how they relate to finite circuits and bonds. As to the existence of Ψ-trees, we can even show that there are always rather special ones: Ψ-trees that are connected not only topologically through their ends, but also as graphs:

Lemma 4. *If G is connected, it has a spanning tree T whose closure in $\widetilde{G}_\Psi$ is a Ψ-tree.*

Proof. It was shown in [1, Thm. 6.3] that G has a spanning tree T whose closure $\overline{T}$ in $\widetilde{G}$ contains no circle. Let $\overline{T}_\Psi$ denote the closure of T in $\widetilde{G}_\Psi$. Then $\overline{T} = \sigma_\Psi(\overline{T}_\Psi)$. Since circles in $\widetilde{G}_\Psi$ define circles in $\widetilde{G}$ (by composition with σ_Ψ), $\overline{T}_\Psi$ contains no circle either.

For a proof that $\overline{T}_\Psi$ is arc-connected it suffices to show that every undominated end $\psi \in \Psi$ contains a ray $R \subseteq T$: then the arc $\pi_\Psi(T) \subseteq \overline{T}_\Psi$ connects the end $\{\psi\} \in \overline{T}_\Psi$ to a vertex, while all the vertices of $\overline{T}_\Psi$ are connected by T. Pick a ray $R' \in \psi$ in G, say $R' = v_0 v_1 \ldots$. By the star-comb lemma [5, Lemma 8.2.2], the connected graph $\bigcup_{n\in\mathbb{N}} v_n T v_{n+1}$ contains a subdivided infinite star with leaves in R' or an infinite comb with teeth in R'. As ψ is not dominated, we must have a comb. The back $R \subseteq T$ of this comb is a ray equivalent to R' that hence lies in ψ.

Being acirclic, arc-connected and spanning, $\overline{T}_\Psi$ is a Ψ-tree. $\square$

Lemma 5. *Assume that G is connected, and let $F \subseteq E$ be a finite set of edges.*

(i) *F is a circuit if and only if it is not contained in the edge set of any Ψ-tree and is minimal with this property.*

(ii) *F is a bond if and only if it meets the edge set of every Ψ-tree and is minimal with this property.*

Proof. (i) Assume first that F is a circuit. Then F is not contained in any Ψ-tree; let us show that every proper subset of F is. We do this by showing the following more general fact:

Every finite set F' of edges not containing a circuit extends to a spanning tree of G whose closure in $\widetilde{G}_\Psi$ is a Ψ-tree. (1)

To prove (1), consider a spanning tree T of G whose closure in $\widetilde{G}_\Psi$ is a Ψ-tree (Lemma 4). Choose it with as many edges in F' as possible. Suppose it fails to contain an edge $f \in F'$. Adding f to T creates a cycle C in $T + f$, which by assumption also contains an edge $e \notin F'$. As C is finite, it is easy to check that $T + f - e$ is another spanning tree whose closure is a Ψ-tree. This contradicts our choice of T.

Conversely, if F is not contained in any Ψ-tree, then by (1) it contains a circuit. If, in addition, it is minimal with the first property, it will in fact be that circuit, since we could delete any other edge without making it extendable to a Ψ-tree.

(ii) If F is a cut, $F = E(V_1, V_2)$ say, then the closures of $G[V_1]$ and $G[V_2]$ in $\widetilde{G}_\Psi$ are disjoint open subsets of $\widetilde{G}_\Psi \setminus \mathring{F}$, so this subspace cannot contain a Ψ-tree. Thus, F meets the edge set of every Ψ-tree.

If F is even a bond, then both V_1 and V_2 induce connected subgraphs. By Lemma 4, these have spanning trees T_i $(i = 1, 2)$ whose closures in $\widetilde{G}_\Psi$ are arc-connected and contain no circle.[2] For every edge $f \in F$, the closure $\overline{T}_\Psi$ of $T := (T_1 \cup T_2) + f$ in $\widetilde{G}_\Psi$ then is a Ψ-tree of G: it still contains no circle, because no arc in $\overline{T}_\Psi \setminus \mathring{f}$ can cross the finite cut F from which it contains no edge (as above). So F is minimal with the property of meeting the edge set of every Ψ-tree.

[2] We are applying Lemma 4 in the subgraphs $G[V_i]$. But since F is finite, the spaces $\widetilde{G[V_i]}_{\Psi_i}$ are canonically embedded in $\widetilde{G}_\Psi$.

Conversely, let us assume that F meets the edge set of every Ψ-tree, and show that F contains a bond. Let T be a spanning tree of G whose closure in $\widetilde{G}_\Psi$ is a Ψ-tree (Lemma 4), chosen with as few edges in F as possible. By assumption, T has an edge f in F. If the bond B of G between the two components of $T - f$ contains an edge $e \notin F$, then $T - f + e$ is another spanning tree whose closure is a Ψ-tree (as before) that contradicts our choice of T. So B contains no such edge e but is contained in F.

If F is minimial with the property of containing an edge from every Ψ-tree, it must be equal to the bond it contains. For by the forward implication of (ii) already proved, any other edge could be deleted from F without spoiling its property of meeting the edge set of every Ψ-tree. □

We begin our study of the spaces introduced in Section 2 by showing that finite separability extends from G to $\widetilde{G}_\Psi$:

Lemma 6. *For every two points $p, q \in \widetilde{G}_\Psi$ that are not inner points of edges there exists a finite set F of edges such that p and q lie in disjoint open sets of $\widetilde{G}_\Psi \setminus \mathring{F}$ whose union is $\widetilde{G}_\Psi \setminus \mathring{F}$.*

Proof. Let us write $p = [x]_\Psi$ and $q = [y]_\Psi$, where x and y are either vertices or undominated ends of G. We shall find a finite cut F of G, with bipartition (X, Y) of V say, such that $x \in \overline{X}$ and $y \in \overline{Y}$, where $\overline{X}$ and $\overline{Y}$ denote closures of X and Y in $|G|_\Psi$. Since F is finite, $\overline{X}$ and $\overline{Y}$ then partition of $|G|_\Psi \setminus \mathring{F}$ into disjoint open sets that are closed under equivalence, so their projections under π_Ψ partition $\widetilde{G}_\Psi \setminus \mathring{F}$ into disjoint open sets containing p and q, respectively.

If x and y are vertices, then F exists by our assumption that G is finitely separable. Suppose now that y is an end. Let us find a finite set $S \not\ni x$ of vertices that separates x from y in G. If x is another end, then S exists since $x \neq y$. If x is a vertex, pick a ray $R \in y$. If there is no S as desired, we can inductively find infinitely many independent x–R paths in G, contradicting the fact that y is undominated.

Having found S, consider the component $C := C(S, y)$ of $G - S$. For each $s \in S$ we can find a finite set $S_s \subseteq C$ of vertices separating s from y in the subgraph of G spanned by C and s, since otherwise s would dominate y (as before). Let $S' := \bigcup_{s \in S} S_s$; this is a finite set of vertices in C that separates all the vertices of S from y in G. Since G is finitely separable, there is a finite set F of edges separating S from S' in G. Choose F minimal. Then, assuming without loss of generality that G is connected, every component of $G - F$ meets exactly one of the sets S and S'. Let X be the set of vertices in components meeting S, and let Y be the set of vertices in components meeting S'. Then (X, Y) is a partition of G crossed by exactly the edges in F, and it is easy to check that F has the desired properties. □

It was proved in [7], under a weaker assumption than finite separability (just strong enough that $\widetilde{G}$ can be defined without identifying distinct vertices) that $\widetilde{G}$ is Hausdorff. For finitely separable graphs, as considered here, the proof is much simpler and extends readily to $\widetilde{G}_\Psi$:

Corollary 7. *$\widetilde{G}_\Psi$ is Hausdorff.*

Proof. Finding disjoint open neigbourhoods for distinct points $p, q \in \widetilde{G}_\Psi$ is easy if one of them is an inner point of an edge. Assume that this is not the case, let F, $\overline{X}$ and $\overline{Y}$ be defined as in Lemma 6 and its proof, and let S be the (finite) set of vertices incident with an edge in F. Then $p \subseteq \overline{X}$ and $q \subseteq \overline{Y}$. Any end $\psi \in p$ has a basic open neighbourhood $\hat{C}(S, \psi) \setminus \Psi^{\complement}$ in $|G|_\Psi$ that is a subset of $\overline{X} \setminus S$. Write O_p for the union of all these neighbourhoods, together with a small open star neighbourhood of the vertex in p if it exists. Define O_q similarly for $q \subseteq \overline{Y}$. Then $\pi_\Psi(O_p)$ and $\pi_\Psi(O_q)$ are disjoint open neighbourhoods of p and q in $\widetilde{G}_\Psi$. □

Our next aim is to select from the basic open neighbourhoods $\hat{C}(S, \omega) \setminus \Psi^{\complement}$ in $|G|_\Psi$ of ends $\omega \in \Psi$ some 'standard' neighbourhoods that behave well under the projection π_Ψ and still form neighbourhood bases of these points ω. Ideally, we would like to find for every end $\omega \in \Psi$ a basis of open neighbourhoods that are closed under $\sim$. That will not be possible, since ends $\omega' \neq \omega$ equivalent to ω can be separated topologically from ω. But we shall be able to find a basis of open neighbourhoods of ω that will be closed under $\sim$ for all points other than ω itself. Then the union of all these neighbourhoods, one for every end $\omega' \sim \omega$, plus an open star neighbourhood of their common dominating vertex, will be closed under $\sim$, and will thus be the pre-image of an open neighbourhood of $\pi_\Psi(\omega) = [\omega]_\Psi$ in $|G|_\Psi$.

Given a bond $F = E(V_1, V_2)$ of G and an end $\omega \in \Psi$ that lies in the $|G|$-closure of V_1 but not of V_2, let

$$\hat{C}_\Psi(F, \omega) \subseteq |G|_\Psi$$

denote the union of the $|G|_\Psi$-closure of $G[V_1]$ with $\mathring{F}$. For every vertex $v \in V_2$ we also call F a *v–ω bond*. Note that $\hat{C}_\Psi(F, \omega)$ depends only on F and ω: since F is a bond, $G - F$ has only two components, so V_1 and V_2 can be recovered from F and ω. Note also that every ray in ω has a tail in $\hat{C}_\Psi(F, \omega)$, so if it starts at v it must have an edge in F.

If $v \in V_2$ is an endvertex of all but finitely many of the edges in F, we say that F is *v-cofinite*. Then the set S of endvertices of F in V_2 is finite and separates ω from $V_2 \setminus S$.

Lemma 8. *Let $\omega \in \Psi$ be an end, and $v \in V$ a vertex.*

(i) *If ω is undominated, then the sets $\{ \hat{C}_\Psi(F, \omega) \mid F$ is a finite bond of $G \}$ form a basis of open neighbourhoods of ω in $|G|_\Psi$.*

(ii) *If ω is dominated by v, then the sets*

$$\{ \hat{C}_\Psi(F, \omega) \mid F \text{ is a } v\text{-cofinite } v\text{–}\omega \text{ bond} \}$$

form a basis of open neighbourhoods of ω in $|G|_\Psi$.

Proof. (i) As F is finite, so is the set S of its endvertices in V_2. Since F is a bond, $G[V_1]$ is connected. Hence $\hat{C}_\Psi(F,\omega)$ equals $\hat{C}(S,\omega) \setminus \Psi^{\complement}$, which is a basic open neighbourhood of ω in $|G|_\Psi$. Conversely, we need to find for any finite set $S \subseteq V$, without loss of generality connected,[3] a finite bond F such that $\hat{C}_\Psi(F,\omega) \subseteq \hat{C}(S,\omega)$. As no vertex dominates ω, there is a finite connected set S' of vertices of $C(S,\omega)$ that separates S from ω in G. (Otherwise we could inductively construct an infinite set of disjoint paths in $C(S,\omega)$ each starting at a vertex adjacent to S and ending on some fixed ray $R \in \omega$; then infinitely many of the starting vertices of these paths would share a neighbour in S, which would dominate ω.) As G is finitely separable, there is a finite set of edges separating S from S' in G. As both S and S' are connected, choosing this set minimal ensures that it is a bond. This bond F satisfies $\hat{C}_\Psi(F,\omega) \subseteq \hat{C}(S,\omega)$.

(ii) Although F is infinite now, the set S of its endvertices in V_2 is finite. Hence $\hat{C}_\Psi(F,\omega)$ is a basic open neighbourhood of ω in $|G|_\Psi$, as in the proof of (i). Conversely, let a finite set $S \subseteq V$ be given; we shall find a v-cofinite v–ω bond F such that $\hat{C}_\Psi(F,\omega) \subseteq \hat{C}(S,\omega)$. The sets $\hat{C}(T,\omega)$ such that $v \in T$ and both $T - v$ and T are connected in G still form a neighbourhood basis for ω in $|G|$, so we may assume that S has these properties. As in the proof of (i), there is a finite connected set S' of vertices in $C(S,\omega)$ that separates $S - v$ from ω in $G - v$, because ω is not dominated in $G - v$. As $G - v$ is finitely separable, there is a finite bond $F = E(V_1, V_2)$ of $G - v$ that separates $S - v$ from S', with $S - v \subseteq V_2$ say. Then $F' := E(V_1, V_2 \cup \{v\})$ is a v-cofinite v–ω bond in G with $\hat{C}_\Psi(F',\omega) \subseteq \hat{C}(S,\omega)$, as before. □

Let us call the open neighbourhoods $\hat{C}_\Psi(F,\omega)$ from Lemma 8 the *standard neighbourhoods* in $|G|_\Psi$ of the ends $\omega \in \Psi$. For points of $|G|_\Psi$ other than ends, let their *standard neighbourhoods* be their basic open neighbourhoods defined in Section 2.

Trivially, standard neighbourhoods of vertices and inner points of edges are closed under $\sim$. Our next lemma says that standard neighbourhoods of ends are nearly closed under $\sim$, in that only the end itself may be equivalent to points outside: to a vertex dominating it, and to other ends dominated by that vertex.

Lemma 9. *If $\hat{C} = \hat{C}_\Psi(F,\omega)$ is a standard neighbourhood of $\omega \in \Psi$ in $|G|_\Psi$, then $[x]_\Psi \subseteq \hat{C}$ for every $x \in \hat{C} \setminus [\omega]_\Psi$.*

Proof. Let S be the finite set of vertices not in $\hat{C}$ that are incident with an edge in F. Suppose, for a contradiction, that there are points $x \sim y$ in $|G|_\Psi$ such that $x \in \hat{C} \setminus [\omega]$ but $y \notin \hat{C} \setminus [\omega]$. Since the unique vertex in the $\sim_\Psi$-class of x and y lies either in $\hat{C} \setminus [\omega]$ or not, we may assume that either x or y is that vertex.

Suppose x is the vertex; then y is an end. Let R be a ray of y that avoids S. Then the finite set $S \subseteq V \setminus \{x\}$ separates x from R, a contradiction.

[3]The sets $\hat{C}(S,\omega)$ with S connected in G also form a neighbourhood basis of ω in $|G|$, since every finite set S of vertices extends to a finite connected set.

Suppose y is the vertex. If $y \notin S$ we argue as before. Suppose that $y \in S$. Note that y does not dominate ω, since $y \sim x \not\sim \omega$. But now the vertex $v \in S$ that dominates ω, if it exists, and the finitely many neighbours of $S \setminus \{v\}$ in $\hat{C}$ together separate y from every ray in x that avoids this finite set, a contradiction. □

Let us extend the notion of standard neighbourhoods from $|G|_\Psi$ to $\widetilde{G}_\Psi$. Call a neighbourhood of a point $[x]_\Psi$ of $\widetilde{G}_\Psi$ a *standard neighbourhood* if its inverse image under π_Ψ is a union $\bigcup_{y \in [x]_\Psi} U_y$ of standard neighbourhoods U_y in $|G|_\Psi$ of the points $y \in [x]_\Psi$. Neighbourhoods in subspaces of $\widetilde{G}_\Psi$ that are induced by such standard neighbourhoods of $\widetilde{G}_\Psi$ will likewise be called *standard*. All standard neighbourhoods in $\widetilde{G}_\Psi$ and its subspaces are open, by definition of the identification and the subspace topology.

Lemma 10. *For every point $[x]_\Psi \in \widetilde{G}_\Psi$ its standard neighbourhoods form a basis of open neighbourhoods in $\widetilde{G}_\Psi$.*

Proof. Given any open neighbourhood N of $[x]_\Psi$ in $\widetilde{G}_\Psi$, its inverse image W under π_Ψ is open in $|G|_\Psi$ and contains every $y \in [x]_\Psi$. By Lemma 8, we can find for each of these y a standard neighbourhood $U_y \subseteq W$ of y in $|G|_\Psi$. By Lemma 9, their union $U = \bigcup_y U_y$ is closed in $|G|_\Psi$ under $\sim$, so $U = \pi_\Psi^{-1}(\pi_\Psi(U))$. Since U is open in $|G|_\Psi$, this means that $\pi_\Psi(U) \subseteq N$ is an open neighbourhood of $[x]_\Psi$ in $\widetilde{G}_\Psi$. □

Our next topic is to compare circles in $\widetilde{G}_\Psi$ with circles in $\widetilde{G}$. We have already seen that circles in $\widetilde{G}_\Psi$ define circles in $\widetilde{G}$, by composition with σ_Ψ. The converse will generally fail: the inverse of σ_Ψ (where it is defined) need not be continuous, so a circle in $\widetilde{G}$ need not induce a circle in $\widetilde{G}_\Psi$ even if its points all lie in the image of σ_Ψ. This is illustrated by the following example.

Example 2. Consider the graph of Figure 2 with $\Psi = \{\psi\}$. The closure of the double ray D in $\widetilde{G}$ is a circle there, since in $\widetilde{G}$ the ends ω and ψ are identified. This circle lies in the image of σ_Ψ, but σ_Ψ^{-1} restricted to it fails to be continuous at the point $\{v, \omega, \psi\}$, which σ_Ψ^{-1} maps to the point $\{v, \psi\}$ of $\widetilde{G}_\Psi$.

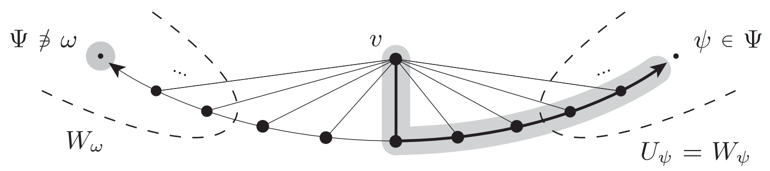

Figure 3: A circle in $\widetilde{G}$ through $p = \{v, \omega, \psi\}$ which defines for $\Psi = \{\psi\}$ a circle in $\widetilde{G}_\Psi$ through $\{v, \psi\}$.

However, the map σ_Ψ^{-1} in this example *is* continuous on the circle in $\widetilde{G}$ shown in Figure 3, which 'does not use' the end $\omega \in \Psi^{\complement}$ when it passes through the point $\{v, \omega, \psi\}$. The fact that circles in $\widetilde{G}$ do induce circles in $\widetilde{G}_\Psi$ in such cases will be crucial to our proof of Theorem 2:

Lemma 11.

(i) *Let $\rho\colon S^1 \to \widetilde{G}_\Psi$ be a circle, with image C say, and let D be the set of all inner points of edges on C. Then every end in the $|G|$-closure*[4] *of D lies in Ψ.*

(ii) *Let $\varphi\colon S^1 \to \widetilde{G}$ be a circle, with image C say, and let D be the set of all inner points of edges on C. If every end in the $|G|$-closure of D lies in Ψ, then the composition $\sigma_\Psi^{-1} \circ \varphi\colon S^1 \to \widetilde{G}_\Psi$ is well defined and a circle in $\widetilde{G}_\Psi$.*

Proof. (i) Consider an end ω in the $|G|$-closure of D. Since $|G|$ (unlike $\widetilde{G}$) is first-countable, there is a sequence $(x_i)_{i\in\mathbb{N}}$ of points in D that converges to ω in $|G|$. Suppose $\omega \in \Psi^{\complement}$. We show that the x_i have no accumulation point on C, indeed in all of $\widetilde{G}_\Psi$; this will contradict the fact that C, being a circle, is compact and contains all the x_i.

Consider a point $p \in \widetilde{G}_\Psi$, and any representative $z \in p \subseteq |G|_\Psi$. As $\omega \in \Psi^{\complement}$ we have $\lim x_i = \omega \neq z$. Therefore z has a neighbourhood W_z in $|G|$ not containing any of the x_i (other than possibly $x_i = z$, which can happen only if $p = \{x_i\}$ is a singleton class). By Lemma 8, the $|G|_\Psi$-neighbourhood $W_z \cap |G|_\Psi$ of z contains a standard $|G|_\Psi$-neighbourhood U_z of z. By Lemma 10, $\pi_\Psi\big(\bigcup_{z\in p} U_z\big)$ is a standard neighbourhood of p in $\widetilde{G}_\Psi$ that contains no x_i other than possibly p itself, so p is not an accumulation point of the x_i.

(ii) Assume that every end in the $|G|$-closure of D lies in Ψ. To show that $\sigma_\Psi^{-1} \circ \varphi$ is well defined, let us prove that $\operatorname{im}\varphi \subseteq \operatorname{im}\sigma_\Psi$. The only points of $\widetilde{G}$ not in the image of σ_Ψ are singleton $\sim$-classes of $|G|$ consisting of an undominated end $\omega \notin \Psi$. By assumption and Lemma 8, such an end ω has a standard neighbourhood in $|G| = |G|_\Omega$ disjoint from D, which π maps to a standard neighbourhood of $\{\omega\}$ in $\widetilde{G}$ disjoint from D. So $\{\omega\}$ is not in the $\widetilde{G}$-closure of D. But that closure is the entire circle C, see [7], giving $\{\omega\} \notin \operatorname{im}\varphi$. This completes the proof of $\operatorname{im}\varphi \subseteq \operatorname{im}\sigma_\Psi$. As σ_Ψ is injective, it follows that $\sigma_\Psi^{-1} \circ \varphi$ is well defined.

To show that σ_Ψ^{-1} is continuous on C, let a point $p \in C$ be given. Since p lies in $\operatorname{im}\varphi \subseteq \operatorname{im}\sigma_\Psi$, it is represented by a point x in $G \cup \Psi$; then $p = [x]$ and $\sigma_\Psi^{-1}(p) = [x]_\Psi$. By Lemma 10, it suffices to find for every standard neighbourhood u of $[x]_\Psi$ in $\operatorname{im}(\sigma_\Psi^{-1}\restriction C)$ a neighbourhood w of $[x]$ in C such that $\sigma_\Psi^{-1}(w) \subseteq u$.

By definition, u is the intersection with $\operatorname{im}(\sigma_\Psi^{-1}\restriction C)$ of a set $U \subseteq \widetilde{G}_\Psi$ whose inverse image under π_Ψ is a union

$$\pi_\Psi^{-1}(U) = \bigcup_{y\in[x]_\Psi} U_y$$

of standard neighbourhoods U_y in $|G|_\Psi$ of the points $y \in [x]_\Psi$. Our aim is to find a similar set W to define w: a set $W \subseteq \widetilde{G}$ such that for $w := W \cap C$ we have $\sigma_\Psi^{-1}(w) \subseteq u$, and such that

$$\pi^{-1}(W) = \bigcup_{y\in[x]} W_y \tag{2}$$

[4]We shall freely consider D as a subset of either $\widetilde{G}_\Psi$ or $|G|$, and similarly in (ii).

where each W_y is a standard neighbourhood of y in $|G|$.

Let us define these W_y, one for every $y \in [x]$. If $y \in G$, then $y \in [x]_\Psi$. Hence U_y is defined, and it is a standard neighbourhood of y also in $|G|$; we let $W_y := U_y$. If $y \in \Psi$, then again $y \in [x]_\Psi$, and U_y (exists and) has the form $\hat{C}_\Psi(F, y)$. We let $W_y := \hat{C}_\Omega(F, y)$ be its closure in $|G|$; this is a standard neighbourhood of y in $|G|$. Finally, if $y \in \Psi^\complement$, then $y \notin [x]_\Psi$ and U_y is undefined. We then let W_y be a standard neighbourhood of y in $|G|$ that is disjoint from D; this exists by assumption and Lemma 8. Let us call these last W_y *new*.

By Lemma 9, all these W_y are closed under equivalence in $|G| \setminus [y]$. Hence $\bigcup_{y \in [x]} W_y$ is closed under equivalence in $|G|$. Its π-image W therefore satisfies (2) and is a standard neighbourhood of $[x]$ in $\widetilde{G}$. Hence, $w := W \cap C$ is a neighbourhood of $[x]$ in C.

It remains to show that σ_Ψ^{-1} maps every point $q \in w$ to u. This is clear for $q = p = [x]$, so assume that $q \neq [x]$. By construction of W and Lemma 9, the set q lies entirely inside one of the W_y. Let us show that no such W_y can be new. Since q is a point in $w \subseteq C$, in which D is dense [7], there is no neighbourhood of q in $\widetilde{G}$ that is disjoint from D. But then q has an element z all whose $|G|$-neighbourhoods meet D. (If not, we could pick for every element of q a standard $|G|$-neighbourhood disjoint from D; then the union of all these would project under π to a standard neighbourhood of q in $\widetilde{G}$ that avoids D.) As W_y is a $|G|$-neighbourhood of $z \in q \subseteq W_y$, it thus cannot be new.

We thus have $q \subseteq W_y$ where W_y is the $|G|$-closure of U_y for some $y \in [x]_\Psi$ (or equal to U_y). In particular, $W_y \setminus U_y \subseteq \Psi^\complement$. As q lies in C, in which D is dense, we cannot have $q = \{\omega\}$ with $\omega \in \Psi^\complement$ (as earlier). So either $q = \{\psi\}$ with $\psi \in \Psi$, or q contains a vertex. In either case, $q \cap U_y \neq \emptyset$, which implies that $\sigma_\Psi^{-1}(q) \in U$. As $q \in C$, this implies $\sigma_\Psi^{-1}(q) \in u$, as desired. □

Lemma 12. *Arc-components of standard subspaces of $\widetilde{G}_\Psi$ are closed.*[5]

Proof. Let X be an arc-component of a standard subspace of $\widetilde{G}_\Psi$. If X is not closed, there is a point q in $\widetilde{G}_\Psi \setminus X$ such that every (standard) neighbourhood of q meets X. As in the proof of Lemma 11, this implies that q has a representative $y \in |G|_\Psi$ such that every standard neighbourhood U_y of y in $|G|_\Psi$ meets $\pi^{-1}(X)$, say in a point $x = x(U_y)$. Clearly, y is an end. Since $x \not\sim y$, we even have $[x]_\Psi \subseteq U_y$ by Lemma 9. Let $U_0 \supseteq U_1 \supseteq \dots$ be a neighbourhood basis for y consisting of such standard neighbourhoods U_y, and let $x_i := x(U_i)$ and $z_i := [x_i]_\Psi$ for all i. Then these x_i converge to y in $|G|_\Psi$, while $(z_i)_{i \in \mathbb{N}}$ is a sequence of points in X that converges in $\widetilde{G}_\Psi$ to $q = [y]_\Psi$.

For every $i \in \mathbb{N} \setminus \{0\}$ let A'_i be a z_i–z_0 arc in X. Define subarcs A_i of the A'_i recursively, choosing as A_i the initial segment of A'_i from its starting point z_i to its first point a_i in $\bigcup_{j<i} A_j$, where $A_0 := \{z_0\}$. (The point a_i exists by the continuity of A'_i, since $\bigcup_{j<i} A_j$ is closed, being a compact subspace of the Hausdorff space $\widetilde{G}_\Psi$.) Note that no two A_i have an edge in common.

[5]This refers to either the subspace or to the entire space $\widetilde{G}_\Psi$; the two are equivalent, since standard subspaces of $\widetilde{G}_\Psi$ are themselves closed in $\widetilde{G}_\Psi$.

Define an auxiliary graph H with vertex set $\{A_i \mid i \in \mathbb{N}\}$ and edges A_iA_j whenever j is the smallest index less than i such that $A_i \cap A_j \neq \emptyset$. Suppose first that H has a vertex A_j of infinite degree. Since the arc A_j is compact, it has a point p every neighbourhood of which meets infinitely many A_i. By Lemma 6, there is a finite set F of edges such that in $\widetilde{G}_\Psi \setminus \mathring{F}$ the points p and q have disjoint open neighbourhoods O_p and O_q partitioning $\widetilde{G}_\Psi \setminus \mathring{F}$. Then for infinitely many i we have both $A_i \cap O_p \neq \emptyset$ and $z_i \in O_q$. For all these i the arc A_i, being connected, must have an edge in the finite set F, a contradiction.

So H is locally finite. By König's infinity lemma, H contains a ray $A_{i_1}A_{i_2}\dots$ such that $i_j < i_k$ whenever $j < k$. We claim that $A := A_{i_1}a_{i_2}A_{i_2}a_{i_3}\dots q$ is an arc in $\widetilde{G}_\Psi$; this will contradict our assumption that A_{i_1} lies in the arc-component X of $\widetilde{G}_\Psi$ while q does not. We only have to show that A is continuous in q. Since every neighbourhood of q in $\widetilde{G}_\Psi$ contains the π_Ψ-image of one of our standard neighbourhoods U_n of y, it suffices to show that for every such U_n we have $A_i \subseteq \pi_\Psi(U_n)$ for all but finitely many i.

Since U_n is a standard neighbourhood of y, there exists a set F of edges such that $U_n = \hat{C}_\Psi(F, y)$ and F is either finite or v-cofinite with $v \sim y$. Let F' be obtained from F by adding to it any other edges incident with such a vertex $v \sim y$. Since none of the A_i contains such a vertex v, and distinct A_i are edge-disjoint, all but finitely many A_i lie in $(\widetilde{G}_\Psi - q) \setminus \mathring{F}'$ and have their starting vertex $z_i = [x_i]_\Psi$ in $\pi_\Psi(U_n)$, by the choice of U_n. To complete our proof, we shall show that $\pi_\Psi(U_n \setminus q) \setminus \mathring{F}'$ and its complement in $(\widetilde{G}_\Psi - q) \setminus \mathring{F}'$ are two open subsets of $(\widetilde{G}_\Psi - q) \setminus \mathring{F}'$ partitioning it: then none of those cofinitely many A_i can meet both, so they will all lie entirely in $\pi_\Psi(U_n)$.

Since U_n is a standard neighbourhood of $y \in q$, the set $U_n \setminus q$ is open in $|G|_\Psi \setminus q$ and closed under equivalence, so $\pi_\Psi(U_n \setminus q)$ is open in $\widetilde{G}_\Psi - q$ and $\pi_\Psi(U_n \setminus q) \setminus \mathring{F}'$ is open in $(\widetilde{G}_\Psi - q) \setminus \mathring{F}'$. Its complement in $(\widetilde{G}_\Psi - q) \setminus \mathring{F}'$ is open, because it is the π_Ψ-image of the ($\sim$-closed) union of the finite set S of vertices that are incident with edges in F but are not in U_n, the edges incident with them that are not in $\mathring{F}'$, and the $|G|_\Psi$-closures of the components of $G - S$ not contained in U_n. The two open sets partition all of $(\widetilde{G}_\Psi - q) \setminus \mathring{F}'$, because U_n is itself the $|G|_\Psi$-closure of a component of $G - S$ together with the edges between S and that component (which all lie in F). □

4 Proof of Theorem 2

We can now apply the lemmas from Section 3 to prove Theorem 2. One of these lemmas, Lemma 11, also implies a characterization of duality in terms of circuits and bonds. Let us include this in the statement of the theorem:

Theorem 13. *Let $G = (V, E, \Omega)$ and $G^* = (V^*, E, \Omega)$ be finitely separable 2-connected graphs with the same edge set E and the same end set Ω, in the sense of* [3]. *Then the following assertions are equivalent:*

(i) *G and G^* are duals of each other.*

(ii) *For all* $\Psi \subseteq \Omega$ *and* $F \subseteq E$ *the following holds:* F *is the edge set of a* Ψ*-tree in* G *if and only if* $F^{\complement}$ *is the edge set of a* $\Psi^{\complement}$*-tree in* G^**.*

(iii) *There exists a set* $\Psi \subseteq \Omega$ *such that for every* $F \subseteq E$ *the following holds:* F *is the edge set of a* Ψ*-tree in* G *if and only if* $F^{\complement}$ *is the edge set of a* $\Psi^{\complement}$*-tree in* G^**.*

(iv) *For all* $\Psi \subseteq \Omega$ *and* $D \subseteq E$ *the following holds:* D *is a* Ψ*-circuit of* G *if and only if* D *is a bond of* G^* *and every end in the closure*[6] *of* $\bigcup D$ *lies in* Ψ*.*

(v) *There exists a set* $\Psi \subseteq \Omega$ *such that for every* $D \subseteq E$ *the following holds:* D *is a* Ψ*-circuit of* G *if and only if* D *is a bond of* G^* *and every end in the closure*[6] *of* $\bigcup D$ *lies in* Ψ*.*

Remark. The fact that (i)–(iii) are symmetrical in G and G^*, while (iv) and (v) are not, is immaterial and only serves to avoid clutter: as noted before, it was proved in [1, Theorem 3.4] that if G^* is a dual of G then G is a dual of G^*.

We shall prove the implications (i)→(iv)→(v)→(i) first, and then the implications (i)→(ii)→(iii)→(i). The two proofs can be read independently.

(i)→(iv) Assume (i), and let $\Psi \subseteq \Omega$ and $D \subseteq E$ be given for a proof of (iv). If D is a Ψ-circuit of G, for the circle $\rho\colon S^1 \to \widetilde{G}_\Psi$ say, it is also a circuit of G with circle $\sigma_\Psi \circ \rho\colon S^1 \to \widetilde{G}$. By (i), then, D is a bond of G^*. By Lemma 11 (i), every end in the closure of $\bigcup D$ lies in Ψ.

If, conversely, D is a bond of G^*, then D is a circuit of G by (i), say with circle $\varphi\colon S^1 \to \widetilde{G}$. If every end in the closure of $\bigcup D$ lies in Ψ then, by Lemma 11 (ii), the composition $\sigma_\Psi^{-1} \circ \varphi$ is well defined and a circle in $\widetilde{G}_\Psi$. The edges it contains are precisely those in D, so D is a Ψ-circuit.

(iv)→(v) Using the empty set for Ψ in (iv) immediately yields (v).

(v)→(i) As G and G^* are finitely separable and 2-connected, [1, Lemma 4.7 (i)] implies that G^* is dual to G as soon as the *finite* circuits of G are precisely the finite bonds of G^*. This is immediate from (v).

Let us now prove the implications (i)→(ii)→(iii)→(i). When we consider edges in E topologically, we take them to include their endvertices in $\widetilde{G}_\Psi$ or in $\widetilde{G}^*_{\Psi^{\complement}}$, depending on the context. Thus, in (ii) and (iii), $\bigcup F$ will be a subspace of $\widetilde{G}_\Psi$ while $\bigcup F^{\complement}$ will be a subspace of $\widetilde{G}^*_{\Psi^{\complement}}$.

(i)→(ii) We first show that (i) implies the analogue of (ii) with ordinary topological connectedness, rather than the arc-connectedness required of a Ψ-tree:

($\star$) For all $F \subseteq E$ and $\Psi \subseteq \Omega$: F is the edge set of a connected spanning Ψ-forest of G if and only if $F^{\complement}$ is the edge set of a connected spanning $\Psi^{\complement}$-forest of G^*.

For our proof of ($\star$) from (i), let $F \subseteq E$ and $\Psi \subseteq \Omega$ be given, and assume that F is the edge set of a connected spanning Ψ-forest T of G. Let X be the closure in $\widetilde{G}^*_{\Psi^{\complement}}$ of $V(\widetilde{G}^*_{\Psi^{\complement}}) \cup \bigcup F^{\complement}$. We shall prove that X is a connected

[6] This refers to the closure in $|G|$ or, equivalently by [3], the closure in $|G^*|$.

subspace of $\widetilde{G}^*_{\Psi^\complement}$ that contains no circle. Then X cannot have isolated vertices, so it will be a standard subspace, and it is spanning by definition. Roughly, the idea is that X should be connected because T is acirclic, and acirclic because T is connected.

Let us show first that X contains no circle. Suppose there is a circle $\varphi\colon S^1 \to X$, with circuit $D \subseteq F^\complement$ say. By Lemma 11 (i) applied to G^* and $\Psi^\complement$, every end in the $|G^*|$-closure of $\bigcup D$ lies in $\Psi^\complement$. But the ends in the $|G^*|$-closure of $\bigcup D$ are precisely those in its $|G|$-closure, by (i). Hence we obtain:

The $|G|$-closure of $\bigcup D$ contains no end from Ψ. (3)

Since D is also the circuit of the circle $\sigma_{\Psi^\complement} \circ \varphi\colon S^1 \to \widetilde{G}^*$, assumption (i) implies that D is a bond in G; let $\{V_1, V_2\}$ be the corresponding partition of V. Let us show the following:

Every point $p \in \widetilde{G}_\Psi$ has a standard neighbourhood N such that $\psi_\Psi^{-1}(N)$ contains vertices from at most one of the sets V_1 and V_2. (4)

Suppose $p \in \widetilde{G}_\Psi$ has no such neighbourhood. Then p has a representative x all whose standard neighbourhoods in $|G|_\Psi$ meet V_1, and a representative y all whose standard neighbourhoods in $|G|_\Psi$ meet V_2.

If $x = y$, the point $x = y =: \psi$ is an end in Ψ. Then every standard neighbourhood of ψ in $|G|_\Psi$ contains a graph-theoretical path from V_1 to V_2, and hence an edge from D, because the subgraphs of G underlying standard neighbourhoods in $|G|_\Psi$ are connected and meet both V_1 and V_2. This contradicts (3).

So $x \neq y$. In particular, p is nontrivial, so it contains a vertex v, say in V_1. Then $v \neq y$, so $y =: \psi \in \Psi$. Pick a ray $R \in \psi$. Replacing R with a tail of R if necessary, we may assume by (3) that R has no edge in D. If all the vertices of R lie in V_1, then every standard neighbourhood of $y = \psi$ meets both V_1 and V_2, which contradicts (3) as in the case of $x = y$. So $R \subseteq G[V_2]$. Let us show that every standard neighbourhood $\hat{C}_\Psi(F', \psi)$ of ψ contains the inner points of an edge from D, once more contrary to (3).

By Lemma 8 (ii), F' is v-cofinite. Since $v \sim \psi$, there are infinitely many v–R paths $P_0, P_1, \dots$ in G that meet pairwise only in v. Since D separates v from R, each P_i contains an edge $e_i \in D$. Only finitely many of the P_i contain one of the finitely many edges from F' that are not incident with v. All the other P_i have all their points other than v in $\hat{C}_\Psi(F', \psi)$, including the inner points of e_i. This completes the proof of (4).

For every point $p \in \widetilde{G}_\Psi$ pick a standard neighbourhood N_p as in (4). Let O_1 be the union of those N_p such that $\pi_\Psi^{-1}(N_p)$ meets V_1, and O_2 the union of the others. Then O_1, O_2 are two open subsets of $\widetilde{G}_\Psi$ covering it, and it is easy to check that $O_1 \cap O_2 \subseteq \mathring{D}$. So no connected subspace of $\widetilde{G}_\Psi \setminus \mathring{D}$ contains vertices from V_1 as well as from V_2. But our connected spanning Ψ-forest T is such a subspace, since its edges lie in $F \subseteq E \setminus D$. This contradiction completes the proof that X contains no circle.

For the proof of ($\star$) it remains to show that X is connected. If not, there are open sets O_1, O_2 in $\widetilde{G}^*_{\Psi^\complement}$ that each meet X and together cover it, but intersect only outside X. It is easy to check that, since X contains all the vertices of $\widetilde{G}^*_{\Psi^{\complement}}$, both O_1 and O_2 contain such a vertex but they have none in common. For $i = 1, 2$, let V_i^* be the set of vertices of G^* representing a vertex of $\widetilde{G}^*_{\Psi^{\complement}}$ in O_i. Let C be a bond contained in the cut $E(V_1^*, V_2^*)$. Note that the edges e of this bond all lie in F: as e is connected but contained in neither O_i, it cannot lie in $O_1 \cup O_2 = X$. As F is the edge set of a Ψ-forest, $C \subseteq F$ cannot be a Ψ-circuit of G. By (i), however, C is a circuit of G, because it is a bond of G^*. By Lemma 11 (ii), therefore, there is an end $\omega \in \Psi^{\complement}$ in the $|G|$-closure of $\mathring{C}$; then ω also lies in the $|G^*|$-closure of $\mathring{C}$.

Let us show that every standard neighbourhood W of $[\omega]_{\Psi^{\complement}}$ in $\widetilde{G}^*_{\Psi^{\complement}}$ contains an edge from C, including its endvertices in $\widetilde{G}^*_{\Psi^{\complement}}$. By definition, W is the image under $\pi_{\Psi^{\complement}}$ of a subset of $|G^*|_{\Psi^{\complement}}$ that contains a standard neighbourhood U of ω in $|G^*|_{\Psi^{\complement}}$. Since ω lies in the $|G^*|$-closure of $\mathring{C}$, this U either contains an edge $e \in C$ together with its endvertices in G^*, or it contains one endvertex (in G^*) and the interior of an edge $e \in C$ whose other endvertex dominates ω in G^*. In both cases, e and its endvertices in $\widetilde{G}^*_{\Psi^{\complement}}$ lie in W.

So every standard neighbourhood of $[\omega]_{\Psi^{\complement}}$ in $\widetilde{G}^*_{\Psi^{\complement}}$ contains an edge from C, including its endvertices in $\widetilde{G}^*_{\Psi^{\complement}}$. In particular, it meets X in both O_1 and O_2, where this edge has its endvertices. So every neighbourhood of $[\omega]_{\Psi^{\complement}}$ in X meets both O_1 and O_2. This contradicts the fact that the O_i induce disjoint open subsets of X of which only one contains the point $[\omega]_{\Psi^{\complement}}$. This completes the proof of ($\star$).

It remains to derive the original statement of (ii) from ($\star$). Suppose (ii) fails, say because there is a Ψ-tree T of G, with edge set F say, such that $F^{\complement}$ is not the edge set of a $\Psi^{\complement}$-tree of G^*. By ($\star$) we know that $F^{\complement}$ is the edge set of a connected spanning $\Psi^{\complement}$-forest X in G^*, which we now want to show is even arc-connected. Suppose it is not. Since the arc-components of X are closed (Lemma 12), no arc-component of X contains all its vertices. Vertices in different arc-components are joined by a finite path in G^*, which contains an edge e whose endvertices lie in different arc-components of X. Then $X \cup e$ still contains no circle, so $F^{\complement} \cup \{e\}$ too is the edge set of a connected spanning $\Psi^{\complement}$-forest of G^*. Thus, by ($\star$), $F \smallsetminus \{e\}$ is the edge set of a connected spanning Ψ-forest of G. This can only be $T \smallsetminus \mathring{e}$, so $T \smallsetminus \mathring{e}$ has precisely two path components D_1 and D_2 but is still connected. Then D_1 and D_2 cannot both be open, or equivalently, cannot both be closed. This contradicts Lemma 12.

(ii)→(iii) Using the empty set for Ψ in (ii) immediately yields (iii).

(iii)→(i) As G and G^* are finitely separable and 2-connected, it suffices by [1, Lemma 4.7 (i)] to show that G^* is a finitary dual of G, i.e., that the finite circuits of G are precisely the finite bonds of G^*. By Lemma 5 (ii), a finite set F of edges is a bond of G^* if and only if it meets the edge set of every $\Psi^{\complement}$-tree of G^* and is minimal with this property. By (iii), this is the case if and only if F is not contained in the edge set of any Ψ-tree of G, and is minimal with this property. By Lemma 5 (i), this is the case if and only if F is a circuit of G. □

Acknowledgement

The ideas that led to the formulation of Theorem 2 were developed jointly with Henning Bruhn. We benefited greatly from his insights at this stage.

References

[1] H. Bruhn and R. Diestel, *Duality in infinite graphs*, Comb., Probab. Comput. **15** (2006), 75–90.

[2] ———, *Infinite matroids in graphs*, Discrete Math. **311** (2011), 1461–1471, Special volume on *Infinite Graphs: Introductions, Connections, Surveys* (R. Diestel, G. Hahn & B. Mohar, eds).

[3] Henning Bruhn and Maya Stein, *Duality of ends*, Comb., Probab. Comput. **19** (2010), 47–60.

[4] R. Diestel, *End spaces and spanning trees*, J. Combin. Theory (Series B) **96** (2006), 846–854.

[5] ———, *Graph Theory*, 4th ed., Springer, 2010.

[6] R. Diestel, *Locally finite graphs with ends: a topological approach*, Discrete Math. **310–312** (2010–11), 2750–2765 (310); 1423–1447 (311); 21–29 (312), arXiv:0912.4213.

[7] R. Diestel and D. Kühn, *Topological paths, cycles and spanning trees in infinite graphs*, Europ. J. Combinatorics **25** (2004), 835–862.

[8] D.W. Hall and G.L. Spencer, *Elementary topology*, John Wiley, New York 1955.

[9] C. Thomassen, *Planarity and duality of finite and infinite graphs*, J. Combin. Theory (Series B) **29** (1980), 244–271.

[10] ———, *Duality of infinite graphs*, J. Combin. Theory (Series B) **33** (1982), 137–160.

Version 7 June, 2011

Orthogonality and minimality in the homology of locally finite graphs*

Reinhard Diestel Julian Pott

Abstract

Given a finite set E, a subset $D \subseteq E$ (viewed as a function $E \to \mathbb{Z}_2$) is orthogonal to a given subspace $\mathcal{F}$ of the $\mathbb{Z}_2$-vector space of functions $E \to \mathbb{Z}_2$ as soon as D is orthogonal to every $\subseteq$-minimal element of $\mathcal{F}$. This fails in general when E is infinite.

However, we prove the above statement for the four subspaces $\mathcal{F}$ of the edge space of any 3-connected locally finite graph that are relevant to its homology: the cut space, the finite-cut space, the topological cycle space, and the finite-cycle space. This solves a problem of [4].

1 Introduction

Let G be a 2-connected locally finite graph, and let $\mathcal{E} = \mathcal{E}(G)$ be its edge space over $\mathbb{Z}_2$. We think of the elements of $\mathcal{E}$ as sets of edges, possibly infinite. Two sets of edges are *orthogonal* if their intersection has (finite and) even cardinality. A set $D \in \mathcal{E}$ is *orthogonal* to a subspace $\mathcal{F} \subseteq \mathcal{E}$ if it is orthogonal to every $F \in \mathcal{F}$. See [4], [3] for any definitions not given below.

The topological *cycle space* $\mathcal{C}(G)$ of G is the subspace of $\mathcal{E}(G)$ generated (via thin sums, possibly infinite) by the *circuits* of G, the edge sets of the topological circles in the Freudenthal compactification $|G|$ of G. This space $\mathcal{C}(G)$ contains precisely the elements of $\mathcal{E}$ that are orthogonal to $\mathcal{B}_{\text{fin}}(G)$, the finite-cut space of G. Similarly, the *finite-cycle space* $\mathcal{C}_{\text{fin}}(G)$ is the subspace of $\mathcal{E}(G)$ generated (via finite sums) by the finite circuits of G. This space $\mathcal{C}_{\text{fin}}(G)$ contains precisely the elements of $\mathcal{E}$ that are orthogonal to $\mathcal{B}(G)$, the cut space of G. Moreover, for any of the four spaces $\mathcal{F}$ just mentioned, we have $\mathcal{F}^{\perp\perp} = \mathcal{F}$. Thus the following equalities hold:

$$\mathcal{C} = \mathcal{B}_{\text{fin}}^{\perp}, \qquad \mathcal{C}_{\text{fin}} = \mathcal{B}^{\perp}, \qquad \mathcal{C}^{\perp} = \mathcal{B}_{\text{fin}}, \qquad \mathcal{C}_{\text{fin}}^{\perp} = \mathcal{B}.$$

Our aim in this note is to show that, whenever $\mathcal{F}$ is one of these four spaces, a set D of edges is orthogonal to $\mathcal{F}$ as soon as it is orthogonal to the minimal nonzero elements of $\mathcal{F}$. This is easy when $\mathcal{F}$ is $\mathcal{C}_{\text{fin}}$ or $\mathcal{B}_{\text{fin}}$:

*Electronic J. Comb. 21 (2014), #P3.36

Proposition 1. *Let $\mathcal{F}$ be a subspace of $\mathcal{E}(G)$ all whose elements are finite sets of edges. Then $\mathcal{F}$ is generated (via finite sums) by its $\subseteq$-minimal nonzero elements.*

Proof. For a contradiction suppose that some $F \in \mathcal{F}$ is not a finite sum of finitely many minimal nonzero elements of $\mathcal{F}$. Choose F with $|F|$ minimal. As F is not minimal itself, by assumption, it properly contains a minimal nonzero element F' of $\mathcal{F}$. As F is finite, $F + F' = F \setminus F' \in \mathcal{F}$ has fewer elements than F, so there is a finite family $(M_i)_{i\leq n}$ of minimal nonzero elements of $\mathcal{F}$ with $\sum_{i\leq n} M_i = F+F'$. This contradicts our assumption, as $F'+\sum_{i\leq n} M_i = F$. □

Corollary 2. *If $\mathcal{F} = \mathcal{C}_{\text{fin}}$ or $\mathcal{F} = \mathcal{B}_{\text{fin}}$, a set D of edges is orthogonal to $\mathcal{F}$ as soon as D is orthogonal to all the minimal nonzero elements of $\mathcal{F}$.* □

When $\mathcal{F} = \mathcal{C}$ or $\mathcal{F} = \mathcal{B}$, the statement of Corollary 2 is generally false for graphs that are not 3-connected. Indeed, for $\mathcal{F} = \mathcal{B}$ let G be the graph obtained from the $\mathbb{N} \times \mathbb{Z}$ grid by doubling every edge between two vertices of degree 3 and subdividing all the new edges. The set D of the edges that lie in a K^3 of G is orthogonal to every bond F of G: their intersection $D \cap F$ is finite and even. But D is not orthogonal to every element of $\mathcal{F} = \mathcal{B}$, since it meets some cuts that are not bonds infinitely.

For $\mathcal{F} = \mathcal{C}$, let B be an infinite bond of the infinite ladder H, and let G be the graph obtained from H by subdividing every edge in B. Then the set D of edges that are incident with subdivision vertices has a finite and even intersection with every topological circuit C, finite or infinite, but it is not orthogonal to every element of $\mathcal{C}$, since it meets some of them infinitely.

However, if G is 3-connected, an edge set is orthogonal to every element of $\mathcal{C}$ or $\mathcal{B}$ as soon as it is orthogonal to every minimal nonzero element of $\mathcal{C}$ or $\mathcal{B}$:

Theorem 3. *Let $G = (V, E)$ be a locally finite 3-connected graph, and $F, D \subseteq E$.*

(i) *$F \in \mathcal{C}^\perp$ as soon as F is orthogonal to all the minimal nonzero elements of $\mathcal{C}$, the topological circuits of G.*

(ii) *$D \in \mathcal{B}^\perp$ as soon as D is orthogonal to all the minimal nonzero elements of $\mathcal{B}$, the bonds of G.*

Although Theorem 3 fails if we replace the assumption of 3-connectedness with 2-connectedness, it turns out that we need a little less than 3-connectedness. Recall that an end ω of G has (combinatorial) *vertex-degree* k if k is the maximum number of vertex-disjoint rays in ω. Halin [6] showed that every end in a k-connected locally finite graph has vertex-degree at least k. Let us call an end ω of G *k-padded* if for every ray $R \in \omega$ there is a neighbourhood U of ω such that for every vertex $u \in U$ there is a *k-fan* from u to R in G, a subdivided k-star with centre u and leaves on R.[1] If every end of G is k-padded, we say that G

[1] For example, if G is the union of complete graphs $K_1, K_2, \ldots$ with $|K_i| = i$, each meeting the next in exactly one vertex (and these are all distinct), then the unique end of G is k-padded for every $k \in \mathbb{N}$.

is *k-padded at infinity.* Note that k-connected graphs are k-padded at infinity. Our proof of Theorem 3(i) will use only that every end has vertex-degree at least 3 and that G is 2-connected. Similarly, and in a sense dually, the proof of Theorem 3(ii) uses only that every end has vertex-degree at least 2 and G is 3-connected at infinity.

Theorem 4. *Let $G = (V, E)$ be a locally finite 2-connected graph.*

(i) *If every end of G has vertex-degree at least 3, then $F \in \mathcal{C}^\perp$ as soon as F is orthogonal to all the minimal nonzero elements of $\mathcal{C}$, the topological circuits of G.*

(ii) *If G is 3-padded at infinity, then $D \in \mathcal{B}^\perp$ as soon as D is orthogonal to all the minimal nonzero elements of $\mathcal{B}$, the bonds of G.*

In general, our notation follows [3]. In particular, given an end ω in a graph G, and a finite set $S \subseteq V(G)$ of vertices, we write $C(S, \omega)$ for the unique component of $G - S$ that contains a ray $R \in \omega$. The *vertex-degree* of ω is the maximum number of vertex-disjoint rays in ω. The mathematical background required for this paper is covered in [4, 5]. For earlier results on the cycle and cut space see Bruhn and Stein [1, 2].

2 Finding disjoint paths

Menger's theorem that the smallest cardinality of an A–B separator in a finite graph is equal to the largest cardinality of a set of disjoint A–B paths trivially extends to infinite graphs. Thus in a locally finite k-connected graph, there are k internally disjoint paths between any two vertices. In Lemmas 5 and 6 we show that, for two such vertices that are close to an end ω, these connecting paths need not use vertices too far away from ω.

In a graph G with vertex sets $X, Y \subseteq V(G)$ and vertices $x, y \in V(G)$, a *k-fan* from X (or x) to Y is a subdivided k-star whose center lies in X (or is x) and whose leaves lie in Y. A *k-linkage* from x to y is a union of k internally disjoint x–y paths. We may refer to a sequence $(v_i)_{i\in\mathbb{N}}$ simply by (v_i), and use $\bigcup(v_i) := \bigcup_{i\in\mathbb{N}}\{v_i\}$ for brevity.

Lemma 5. *Let G be a locally finite graph with an end ω, and let $(v_i)_{i\in\mathbb{N}}$ and $(w_i)_{i\in\mathbb{N}}$ be two sequences of vertices converging to ω. Let k be a positive integer.*

(i) *If for infinitely many $n \in \mathbb{N}$ there is a k-fan from v_n to $\bigcup(w_i)$, then there are infinitely many disjoint such k-fans.*

(ii) *If for infinitely many $n \in \mathbb{N}$ there is a k-linkage from v_n to w_n, then there are infinitely many disjoint such k-linkages.*

Proof. For a contradiction, suppose $k \in \mathbb{N}$ is minimal such that there is a locally finite graph $G = (V, E)$ with sequences $(v_i)_{i\in\mathbb{N}}$ and $(w_i)_{i\in\mathbb{N}}$ in which either (i) or (ii) fails. Then $k > 1$, since for every finite set $S \subseteq V(G)$ the unique component

$C(S,\omega)$ of $G - S$ that contains rays from ω is connected and contains all but finitely many vertices from $\bigcup(v_i)$ and $\bigcup(w_i)$.

For a proof of (i) it suffices to show that for every finite set $S \subseteq V(G)$ there is an integer $n \in \mathbb{N}$ and a k-fan from v_n to $\bigcup(w_i)$ avoiding S. Suppose there is a finite set $S \subseteq V(G)$ that meets all k-fans from $\bigcup(v_i)$ to $\bigcup(w_i)$. By the minimality of k, there are infinitely many disjoint $(k-1)$-fans from $\bigcup(v_i)$ to $\bigcup(w_i)$ in $C := C(S,\omega)$. Thus, there is a subsequence $(v'_i)_{i\in\mathbb{N}}$ of $(v_i)_{i\in\mathbb{N}}$ in C and pairwise disjoint $(k-1)$-fans $F_i \subseteq C$ from v'_i to $\bigcup(w_i)$ for all $i \in \mathbb{N}$. For every $i \in \mathbb{N}$ there is by Menger's theorem a $(k-1)$-separator S_i separating v'_i from $\bigcup(w_i)$ in C, as by assumption there is no k-fan from v'_i to $\bigcup(w_i)$ in C. Let C_i be the component of $G - (S \cup S_i)$ containing v'_i.

Since F_i is a subdivided $|S_i|$-star, $S_i \subseteq V(F_i)$. Hence for all $i \neq j$, our assumption of $F_i \cap F_j = \emptyset$ implies that $F_i \cap S_j = \emptyset$, and hence that $F_i \cap C_j = \emptyset$. But then also $C_i \cap C_j = \emptyset$, since any vertex in $C_i \cap C_j$ coud be joined to v'_j by a path P in C_j and to v'_i by a path Q in C_i, giving rise to a v'_j–$\bigcup(w_i)$ path in $P \cup Q \cup F_i$ avoiding S_j, a contradiction.

As $S \cup S_i$ separates v'_i from $\bigcup(w_i)$ in G and there is, by assumption, a k-fan from v'_i to $\bigcup(w_i)$ in G, there are at least k distinct neighbours of C_i in $S \cup S_i$. Since $|S_i| = k-1$, one of these lies in S. This holds for all $i \in \mathbb{N}$. As $C_i \cap C_j = \emptyset$ for distinct i and j, this contradicts our assumption that G is locally finite and S is finite. This completes the proof of (i).

For (ii) it suffices to show that for every finite set $S \subseteq V(G)$ there is an integer $n \in \mathbb{N}$ such that there is a k-linkage form v_n to w_n avoiding S. Suppose there is a finite set $S \subseteq V(G)$ that meets all k-linkages from v_i to w_i for all $i \in \mathbb{N}$. By the minimality of k there is an infinite family $(L_i)_{i\in I}$ of disjoint $(k-1)$-linkages L_i in $C := C(S,\omega)$ from v_i to w_i. As earlier, there are pairwise disjoint $(k-1)$-sets $S_i \subseteq V(L_i)$ separating v_i from w_i in C, for all $i \in I$. Let C_i, D_i be the components of $C - S_i$ containing v_i and w_i, respectively. For no $i \in I$ can both C_i and D_i have ω in their closure, as they are separated by the finite set $S \cup S_i$. Thus for every $i \in I$ one of C_i or D_i contains at most finitely many vertices from $\bigcup_{i\in I} L_i$. By symmetry, and replacing I with an infinite subset of itself if necessary, we may assume the following:

> *The components C_i with $i \in I$ each contain only finitely many vertices from $\bigcup_{i\in I} L_i$.* (1)

If infinitely many of the components C_i are pairwise disjoint, then S has infinitely many neighbours as earlier, a contradiction. By Ramsey's theorem, we may thus assume that

$$C_i \cap C_j \neq \emptyset \textit{ for all } i, j \in I. \tag{2}$$

Note that if C_i meets L_j for some $j \neq i$, then $C_i \supseteq L_j$, since L_j is disjoint from $L_i \supseteq S_i$. By (1), this happens for only finitely many $j > i$. We can therefore choose an infinite subset of I such that $C_i \cap L_j = \emptyset$ for all $i < j$ in I. In particular, $(C_i \cup S_i) \cap S_j = \emptyset$ for $i < j$. By (2), this implies that

$$C_i \cup S_i \subseteq C_j \textit{ for all } i < j. \tag{3}$$

By assumption, there exists for each $i \in I$ some v_i–w_i linkage of k independent paths in G, one of which avoids S_i and therefore meets S. Let P_i denote its final segment from its last vertex in S to w_i. As $w_i \in C \setminus (C_i \cup S_i)$ and P_i avoids both S_i and S (after its starting vertex in S), we also have

$$P_i \cap C_i = \emptyset. \tag{4}$$

On the other hand, L_i contains $v_i \in C_i \subseteq C_{i+1}$ and avoids S_{i+1}, so $w_i \in L_i \subseteq C_{i+1}$. Hence P_i meets S_j for every $j \geq i+1$ such that $P_i \nsubseteq S \cup C_j$. Since the $L_j \supseteq S_j$ are disjoint for different j, this happens for only finitely many $j > i$. Deleting those j from I, and repeating that argument for increasing i in turn, we may thus assume that $P_i \subseteq S \cup C_{i+1}$ for all $i \in I$. By (3) and (4) we deduce that $P_i \setminus S$ are now disjoint for different values of $i \in I$. Hence S contains a vertex of infinite degree, a contradiction. □

Recall that G is *k-padded* at an end ω if for every ray $R \in \omega$ there is a neighbourhood U such that for all vertices $u \in U$ there is a k-fan from u to R in G. Our next lemma shows that, if we are willing to make U smaller, we can find the fans locally around ω:

Lemma 6. *Let G be a locally finite graph with a k-padded end ω. For every ray $R \in \omega$ and every finite set $S \subseteq V(G)$ there is a neighbourhood $U \subseteq C(S,\omega)$ of ω such that from every vertex $u \in U$ there is a k-fan in $C(S,\omega)$ to R.*

Proof. Suppose that, for some $R \in \omega$ and finite $S \subseteq V(G)$, every neighbourhood $U \subseteq C(S,\omega)$ of ω contains a vertex u such that $C(S,\omega)$ contains no k-fan from u to R. Then there is a sequence $u_1, u_2, \ldots$ of such vertices converging to ω. As ω is k-padded there are k-fans from infinitely many u_i to R in G. By Lemma 5(i) we may assume that these fans are disjoint. By the choice of $u_1, u_2, \ldots$, all these disjoint fans meet the finite set S, a contradiction. □

3 The proof of Theorems 3 and 4

As pointed out in the introduction, Theorem 4 implies Theorem 3. It thus suffices to prove Theorem 4, of which we prove (i) first. Consider a set $F \neq \emptyset$ of edges that meets every circuit of G evenly. We have to show that $F \in \mathcal{C}^\perp$, i.e., that F is a finite cut. (Recall that $\mathcal{C}^\perp$ is known to equal $\mathcal{B}_{fin}$, the finite-cut space.) As F meets every finite cycle evenly it is a cut, with bipartition (A,B) say. Suppose F is infinite. Let $\mathcal{R}$ be a set of three disjoint rays that belong to an end ω in the closure of F. Every R–R' path P for two distinct $R, R' \in \mathcal{R}$ lies on the unique topological circle $C(R,R',P)$ that is contained in $R \cup R' \cup P \cup \{\omega\}$. As every circuit meets F finitely, we deduce that no ray in $\mathcal{R}$ meets F again and again. Replacing the rays in $\mathcal{R}$ with tails of themselves as necessary, we may thus assume that F contains no edge from any of the rays in $\mathcal{R}$. Suppose F separates $\mathcal{R}$, with the vertices of $R \in \mathcal{R}$ in A and the vertices of $R', R'' \in \mathcal{R}$ in B say. Then there are infinitely many disjoint R–$(R' \cup R'')$ paths each meeting F

at least once. Infinitely many of these disjoint paths avoid one of the rays in B, say R''. The union of these paths together with R and R' contains a ray $W \in \omega$ that meets F infinitely often. For every R''–W path P, the circle $C(W, R'', P)$ meets F in infinitely many edges, a contradiction. Thus we may assume that F does not separate $\mathcal{R}$, and that $G[A]$ contains $\bigcup \mathcal{R}$.

As ω lies in the closure of F, there is a sequence $(v_i)_{i\in\mathbb{N}}$ of vertices in B converging to ω. As G is 2-connected there is a 2-fan from each v_i to $\bigcup \mathcal{R}$ in G. By Lemma 5 there are infinitely many disjoint 2-fans from $\bigcup(v_i)$ to $\bigcup \mathcal{R}$. We may assume that every such fan has at most two vertices in $\bigcup \mathcal{R}$. Then infinitely many of these fans avoid some fixed ray in $\mathcal{R}$, say R. The two other rays plus the infinitely many 2-fans meeting only these together contain a ray $W \in \omega$ that meets F infinitely often and is disjoint from R. Then for every R–W path P we get a contradiction, as $C(R, W, P)$ is a circle meeting F in infinitely many edges.

To prove (ii), let $D \subseteq E$ be a set of edges that meets every bond evenly. We have to show that $D \in \mathcal{B}^\perp$, i.e., that D has an (only finite and) even number of edges also in every cut that is not a bond.

As D meets every finite bond evenly, and hence every finite cut, it lies in $\mathcal{B}_{fin}^\perp = \mathcal{C}$. We claim that

D is a disjoint union of finite circuits. $\quad (\star)$

To prove $(\star)$, let us show first that every edge $e \in D$ lies in some finite circuit $C \subseteq D$. If not, the endvertices u, v of e lie in different components of $(V, D \setminus \{e\})$, and we can partition V into two sets A, B so that e is the only A–B edge in D. The cut of G of all its A–B edges is a disjoint union of bonds [3], one of which meets D in precisely e. This contradicts our assumption that D meets every bond of G evenly.

For our proof of $(\star)$, we start by enumerating D, say as $D =: \{e_1, e_2, \dots\} =: D_0$. Let $C_0 \subseteq D_0$ be a finite circuit containing e_0, let $D_1 := D_0 \setminus C_0$, and notice that D_1, like D_0, meets every bond of G evenly (because C_0 does). As before, D_1 contains a finite circuit C_1 containing the edge e_i with $i = \min\{j \mid e_j \in D_1\}$. Continuing in this way we find the desired decomposition $D = C_1 \cup C_2 \cup \dots$ of D into finite circuits. This completes the proof of $(\star)$.

As every finite circuit lies in $\mathcal{B}^\perp$, it suffices by $(\star)$ to show that D is finite. Suppose D is infinite, and let ω be an end of G in its closure. Let us say that two rays R and R' *hug* D if every neighbourhood U of ω contains a finite circuit $C \subseteq D$ that is neither separated from R by R' nor from R' by R in U. We shall construct two rays R and R' that hug D, inductively, as follows.

Let $S_0 = \emptyset$, and let R_0, R_0' be disjoint rays in ω. (These exist as G is 2-connected [6].) For step $j \geq 1$, assume that let S_i, R_i, and R_i' have been defined for all $i < j$ so that R_i and R_i' each meet S_i in precisely some initial segement (and otherwise lie in $C(S_i, \omega)$) and S_i contains the ith vertex in some fixed enumeration of V. If the jth vertex in this enumeration lies in $C(S_{j-1}, \omega)$, add to S_{j-1} this vertex and, if it lies on R_{j-1} or R_{j-1}', the initial segement of that ray up to it. Keep calling the enlarged set S_{j-1}. For the following choice of S

we apply Lemma 6 to S_{j-1} and each of R_{j-1} and R'_{j-1}. Let $S \supseteq S_{j-1}$ be a finite set such that from every vertex v in $C(S,\omega)$ there are 3-fans in $C(S_{j-1},\omega)$ both to R_{j-1} and to R'_{j-1}. By ($\star$) and the choice of ω, there is a finite circuit $C_j \subseteq D$ in $C(S,\omega)$. Then C_j can not be separated from R_{j-1} or R'_{j-1} in $C(S_{j-1},\omega)$ by fewer than three vertices, and thus there are three disjoint paths from C_j to $R_{j-1} \cup R'_{j-1}$ in $C(S_{j-1},\omega)$.

There are now two possible cases. The first is that in $C(S_{j-1},\omega)$ the circuit C_j is neither separated from R_{j-1} by R'_{j-1} nor from R'_{j-1} by R_{j-1}. This case is the preferable case. In the second case one ray separates C_j from the other. In this case we will reroute the two rays to obtain new rays as in the first case. We shall then 'freeze' a finite set containing initial parts of these rays, as well as paths from each ray to C_j. This finite fixed set will not be changed in any later step of the construction of R and R'. In detail, this process is as follows.

If $C(S_{j-1},\omega)$ contains both a C_j–R_{j-1} path P avoiding R'_{j-1} and a C_j–R'_{j-1} path P' avoiding R_{j-1}, let Q and Q' be the initial segments of R_{j-1} and R'_{j-1} up to P and P', respectively. Then let $R_j = R_{j-1}$ and $R'_j = R'_{j-1}$ and

$$S_j = S_{j-1} \cup V(P) \cup V(P') \cup V(Q) \cup V(Q').$$

This choice of S_j ensures that the rays R, R' constructed form the R_i and R'_i in the limit will not separate each other from C_j, because they will satisfy $R \cap S_j = R_j \cap S_j$ and $R' \cap S_j = R'_j \cap S_j$.

If the ray R_{j-1} separates C_j from R'_{j-1}, let $\mathcal{P}_j$ be a set of three disjoint C_j–R'_{j-1} paths avoiding S_{j-1}. All these paths meet R_{j-1}. Let $P_1 \in \mathcal{P}_j$ be the path which R_{j-1} meets first, and $P_3 \in \mathcal{P}_j$ the one it meets last. Then $R_{j-1} \cup C_j \cup P_1 \cup P_3$ contains a ray R_j with initial segment $R_{j-1} \cap S_{j-1}$ that meets C_j but is disjoint from the remaining path $P_2 \in \mathcal{P}$ and from R'_{j-1}. Let $R'_j = R'_{j-1}$, and let S_j contain S_{j-1} and all vertices of $\bigcup \mathcal{P}_j$, and the initial segments of R_{j-1} and R'_{j-1} up to their last vertex in $\bigcup \mathcal{P}$. Note that R_j meets C_j, and that P_2 is a C_j–R'_j path avoiding R_j.

If the ray R'_{j-1} separates C_j from R_{j-1}, reverse their roles in the previous part of the construction.

The edges that lie eventually in R_i or R'_i as $i \to \infty$ form two rays R and R' that clearly hug D.

Let us show that there are two disjoint combs, with spines R and R' respectively, and infinitely many disjoint finite circuits in D such that each of the combs has a tooth in each of these circuits. We build these combs inductively, starting with the rays R and R' and adding teeth one by one.

Let $T_0 = R$ and $T'_0 = R'$ and $S_0 = \emptyset$. Given $j \geq 1$, assume that T_i, T'_i and S_i have been defined for all $i < j$. By Lemma 6 there is a finite set $S \supseteq S_{j-1}$ such that every vertex of $C(S,\omega)$ sends a 3-fan to $R \cup R'$ in $C(S_{j-1},\omega)$. As R and R' hug D there is a finite cycle C in $C(S,\omega)$ with edges in D, and which neither of the rays R or R' separates from the other. By the choice of S, no one vertex of $C(S_{j-1},\omega)$ separates C from $R \cup R'$ in $C(S_{j-1},\omega)$. Hence by Menger's theorem there are disjoint $(R \cup R')$–C paths P and Q in $C(S_{j-1},\omega)$. If P starts on R and Q starts on R' (say), let $P' := Q$. Assume now that P and Q start on the same ray R or R', say on R. Let Q' be a path from R' to $C \cup P \cup Q$ in $C(S_{j-1}, \omega)$ that avoids R. As Q' meets at most one of the paths P and Q, we may assume it does not meet P. Then $Q' \cup (Q \setminus R)$ contains an R'–C path P' disjoint from P and R. In either case, let $T_j = T_{j-1} \cup P$, let $T'_j = T'_{j-1} \cup P'$, and let S_j consist of S_{j-1}, the vertices in $C \cup P \cup P'$, and the vertices on R and R' up to their last vertex in $C \cup P \cup P'$.

The unions $T = \bigcup_{i \in \mathbb{N}} T_i$ and $T' = \bigcup_{i \in \mathbb{N}} T'_i$ are disjoint combs that have teeth in infinitely many common disjoint finite cycles whose edges lie in D. Let A be the vertex set of the component of $G - T$ containing T', and let $B := V \setminus A$. Since T is connected, $E(A, B)$ is a bond, and its intersection with D is infinite as every finite cycle that contains a tooth from both these combs meets $E(A, B)$ at least twice. This contradiction implies that D is finite, as desired. □

References

[1] H. Bruhn and M. Stein. On end degrees and infinite circuits in locally finite graphs. *Combinatorica*, 27:269–291, 2007.

[2] Henning Bruhn and Maya Stein. Duality of ends. *Comb., Probab. Comput.*, 19:47–60, 2010.

[3] R. Diestel. *Graph Theory.* Springer, 4th edition, 2010.

[4] R. Diestel. Locally finite graphs with ends: a topological approach. *Discrete Math.*, 310–312: 2750–2765 (310); 1423–1447 (311); 21–29 (312), 2010–11. arXiv:0912.4213.

[5] R. Diestel and P. Sprüssel. The homology of locally finite graphs with ends. *Combinatorica*, 30:681–714, 2010.

[6] R. Halin. A note on Menger's theorem for infinite locally finite graphs. *Abh. Math. Sem. Univ. Hamburg*, 40:111–114, 1974.

Extending cycles locally to Hamilton cycles

Matthias Hamann Florian Lehner* Julian Pott

Abstract

A Hamilton circle in an infinite graph is a homeomorphic copy of the unit circle S^1 that contains all vertices and all ends precisely once. We prove that every connected, locally connected, locally finite, claw-free graph has such a Hamilton circle, extending a result of Oberly and Sumner to infinite graphs. Furthermore, we show that such graphs are Hamilton-connected if and only if they are 3-connected, extending a result of Asratian. *Hamilton-connected* means that between any two vertices there is a Hamilton arc, a homeomorphic copy of the unit interval $[0, 1]$ that contains all vertices and all ends precisely once.

1 Introduction

The proofs of many classical sufficient conditions for the existence of a Hamilton cycle can be outlined as follows. Start with an arbitrary cycle, extend the cycle by some additional vertices and iterate this extension procedure until the cycle covers all vertices. It is often the case that the extension happens locally, that is, most of the original cycle—in fact everything outside a bounded distance from some newly added vertex—remains unchanged.

While such a strategy will obviously give a Hamilton cycle for finite graphs, the situation is more complicated with infinite graphs, particularly because it is not entirely clear what an infinite analogue of a Hamilton cycle should be.

Considering spanning rays or spanning double rays as infinite analogues of Hamiltonian cycles has yielded some results (e. g., see Thomassen [23]) but it has the obvious drawback that a graph with more than two ends can never be Hamiltonian. However, there is a different approach suggested by Diestel and Kühn [11, 12]. They define a circle in an infinite graph G to be a homeomorphic image of the unit circle in the end compactification of G. This approach has not only been successful in generalizing Hamiltonicity results to locally finite graphs, it has also yielded generalizations of many theorems about the cycle space (see [9] for an overview).

Unfortunately, the extension strategies mentioned above do not immediately give a Hamiltonian circle in this sense. There will be a limit object if the extension procedure only alters the cycle locally, but it is not guaranteed that this limit object will be a circle. In particular, ensuring injectivity at the ends of G can be challenging. So far, there are some results on Hamilton circles in infinite graphs, see [3, 5, 7, 15, 18]

*The author acknowledges the support of the Austrian Science Fund (FWF), project W1230-N13.

In this paper we present a strategy that ensures that the limit will be a Hamilton circle. We call it a k-local skip-and-glue strategy and it roughly states that every finite 2-regular subgraph can be extended to a larger 2-regular graph with an extension of bounded size, or to a 2-regular graph with fewer components. In Section 3 we define this k-local skip-and-glue strategy and prove that a locally finite graph with such a strategy contains a Hamilton circle.

We then proceed to prove that locally finite claw-free graphs always admit such a strategy. So we shall show that the following two theorems can be extended to locally finite graphs and affirmatively answers questions of Stein [22, Question 5.1.3] and Bruhn, see Stein [22, Question 5.1.4].

Theorem 1.1. [21, Theorem 1] *Every finite connected locally connected claw free graph on at least three vertices is Hamiltonian.*

Theorem 1.2. [1, Theorem 3.4] *Every finite connected locally connected claw-free graph is Hamilton-connected if and only if it is* 3*-connected.*

We also give some corollaries of the two theorems whose infinite but locally finite counterparts are corollaries to the infinite versions of those theorems. Similar questions on Hamilton circles in infinite graphs are currently investigated by Heuer [17].

2 The topological space $|G|$

Let $G = (V, E)$ be a locally finite graph. A *ray* is a one-way infinite path. Two rays are *equivalent* if they lie eventually int he same component of $G - S$ for every finite vertex set $S \subseteq V$. This is an equivalence relation whose equivalence classes are the *ends* of G. For $S \subseteq V$ and an end ω, let $C(S, \omega)$ be the component of $G - S$ that contains some ray, and hence a tail of every ray, in ω and let $\Omega(S, \omega)$ be the set of ends with at least one ray in $C(S, \omega)$.

The space $|G|$ is a topological space on G with its ends such that it coincides on G with its 1-complex and such that the sets

$$C(S, \omega) \cup \Omega(S, \omega) \cup E'(S, \omega)$$

for all finite $S \subseteq V$ and ends ω form a basis for the open neighbourhoods around each end ω, where $E'(S, \omega)$ is any union of half edges $(z, y]$, one for every edge xy with $x \in S$ and $y \in C(S, \omega)$, with z an inner point of xy. It can be proved (see [10]) that $|G|$ is the Freudenthal compactification [14] of the 1-complex G.

A *standard subspace* of $|G|$ is a subspace of the form $\overline{U} \cup \mathring{F}$, where (U, F) is a subgraph of G. We will need the following two lemmas on standard subspaces:

Lemma 2.1. *A standard subspace of $|G|$ is topologically connected if and only if one of the following statements holds.*

(i) *It contains an edge from every finite cut of G which meets both sides* [8, Lemma 8.5.5].

(ii) *It is arc-connected* [13, Theorem 2.6].

A *circle* in $|G|$ is a homeomorphic image of S^1 and an arc in $|G|$ is a homeomorphic image of $[0, 1]$. A circle that contains every vertex and every end of G is a *Hamilton circle* and an arc whose endpoints are vertices and that contain every vertex and every end of G is a *Hamilton arc*. We call G *Hamilton-connected* if there is a Hamilton arc between each two vertices of G.

The *degree* of an end ω in a standard subspace $X \subseteq |G|$ is the supremum of the cardinalities of sets of vertex disjoint arcs ending in ω. This notion of degree for the whole space $|G|$ coincides with the notion of vertex-degree of G.

Combining [11, Theorem 7.1] and [2, Theorem 5], we obtain the following theorem (see [9, Theorem 2.5]):

Theorem 2.2. *Let G be a locally finite graph and $F \subseteq E(G)$. Then the following statements are equivalent:*

(i) *Every vertex and every end has even degree in $\overline{F}$.*

(ii) *Every finite cuts meets F in an even number of edges.*

The following charachterization of subspaces that are circles comes in extremly handy.

Lemma 2.3. [4, Proposition 3] *A standard subspace X of $|G|$ is a circle if and only if it is topologically connected and every vertex and end of G in X has degree* 2*.*

3 Skip-and-glue extensions

Let H be a subgraph of a graph G. A vertex $v \in G$ has *depth* $d(v, G-H)$ in H, that is the distance from v to anything outside of H.

Let F be a 2-regular finite sugraph of G. Let P be a path whose end vertices v and w are adjacent in F but that is otherwise disjoint from F. Then the 2-regular subgraph $(F \cup P) - vw$ is the *glue extension* of F *by* P *over* the edge vw. A path $R \subseteq F$ is *skippable* if its two end vertices x and y are adjacent in G but not in F. The edge xy is a *bypass* of R. A 2-regular graph F' obtained from F by successively replacing skippable paths by their bypasses is called a *reduction*.

A glue extension of a reduction F' of F by P that covers F is a *skip-and-glue extension of F by P*. Its *depth* is the maximum depth of any vertex of $F \cap P$ in its component of F. It is a *proper* skip-and-glue extension *via u* if there is a vertex u in $P \setminus F$. A skip-and-glue extension of F via u is *k-local* if its *length*, the length of P, is at most k and u is adjacent to F.

For a finite 2-regular subgraph $F \subseteq G$ with at least two components, a *glue fusion* of F is a 2-regular subgraph F' such that F' coincides with $F - e - f$ on $V(D)$ e, f are edges from distinct components of F and such that the edges in $F' \setminus F$ form two disjoint paths connecting the two end vertices of e with those of f. In particular, the glue fusion F' has less components than F: precisely two components of F are *fused* to one of F'. We call D a *skip-and-glue fusion* of F if it is a glue fusion of a reduction F' of F that covers F. This skip-and-glue fusion is *k-local* if the length of each of the two non-trivial paths in $D \setminus F'$ is at most k. A k-local skip-and-glue fusion is *centred around* a vertex $u \in V(F)$ if each vertex of the two non-trivial paths in $F' \setminus F$ has distance at most k to u and if u lies in one of the two fused components.

A *k-local skip-and-glue strategy (with respect to S)* is a function which assigns to every pair (F, u) of a finite 2-regular subgraph $F \subseteq G$ (covering S) all whose components contain at least k vertices and a vertex $u \in N(F)$ a k-local skip-and-glue extension of F via u and which assigns to every pair (F, u) of a finite 2-regular disconnected subgraph $F \subseteq G$ (covering S) all whose components contain at least k vertices and vertex $u \in V(F)$ that has a neighbour in a different component of F a k-local skip-and-glue fusion of F centred around u.

Note that the image of a pair (F, u) of a k-local skip-and-glue strategy is not necessarily a graph all whose components contain at least k vertices.

Lemma 3.1. *Let G be a locally finite connected graph containing a cycle. If there is a k-local skip-and-glue strategy for G, then G is 2-connected.*

Proof. Let C be a cycle in G and suppose that some vertex v separates G. Let P be a path with one end vertex in C and one end vertex in a component of $G - v$ not containing C. Let C' be a cycle containing C and a maximal number of vertices from P. As there is a k-local skip-and-glue strategy for G, all vertices of P are contained in C'. This contradicts that v is a separating vertex as it does not separate C' but P meets two components of $G - v$. □

Lemma 3.2. *Let G be a locally finite 2-connected graph. Then every neighbourhood of every end contains a cycle of length at least n for all $n \in \mathbb{N}$.*

Proof. For an end ω of G and a neighbourhood U of ω and some $n \in \mathbb{N}$, let R_1, R_2 be two vertex disjoint rays in ω that lie in U, which exist as G is 2-connected [16]. Let $P_1, P_2, \dots$ be an infinite sequence of vertex disjoint finite R_1–R_2 paths. Such a sequence exists, as in every neighbourhood of ω there is one such path, and for any collection of finitely many paths there is a neighbourhood of ω avoiding those. Let P_i and P_j be two of these path whose end vertices on R_1 have distance at least n. The unique cycle in $R_1 \cup P_j \cup R_2 \cup P_i$ has length at least n. □

We will see later that a locally finite graph which satisfies the conditions of Theorem 1.1 always admits a 4-local skip-and-glue strategy, hence the following lemma can be used to extend the theorem to locally finite graphs.

Theorem 3.3. *Let $k \in \mathbb{N}$ and let $G = (V, E)$ be a locally finite connected graph.*

(I) *If G has a k-local skip-and-glue strategy and contains a cycle of length at least k, then G is hamiltonian.*

(II) *If G contains a v–w path P whose end vertices have depth at least $k+1$ in P and $G + vw$ has a k-local skip-and-glue strategy with respect to $V(P)$, then G contains a v–w arc that is hamiltonian.*

Proof. To prove (I) we will define a sequence of finite cycles $(C_i)_{i \in \mathbb{N}}$ such that

(i) $V(C_{i+1})$ contains $V(C_i)$,

(ii) Every edge with depth at least $2k$ in C_i is contained in C_{i+1} if and only if it is contained in C_i.

(iii) every vertex is contained in some C_i,

(iv) for every end ω of G and every finite set $V' \subseteq V$ there is a finite cut F separating ω from V' and an index i_0 such that $|E(C_{i_0}) \cap F| = 2$ and $E(C_{i_0}) \cap F = E(C_i) \cap F$ for every $i > i_0$.

Let C be the limit of the sequence C_i, that is, C is the set of all those edges that are contained in the C_i eventually. Let us first show that $\overline{C}$ is a Hamilton circle as soon as (i) to (iv) are satisfied.

Together, (i) and (iii) imply $\bigcup V(C_i) = V(G)$. By (iii) and as G is locally finite there is for every vertex v an index j such that C_j contains all vertices with distance $k+1$ or less from v. By (ii) every vertex has degree 2 in C and $V(C) = V(G)$. Let F be a finite cut with bipartition (A, B). Then there is some $j \in \mathbb{N}$ such that all vertices of F lie in C_j. As C_{j+1} is connected and meets both A and B, it must contain an even number of edges from F. These edges are contained in C by (ii). As C meets every finite cut in an even number of edges, every vertex and every end of $\overline{C}$ has even degree by Theorem 2.2 and $\overline{C}$ is topologically connected by Lemma 2.1 (i). Additionally, Lemma 2.1 (ii) implies that the standard subspace $\overline{C}$ is arc-connected and hence its degree is at least 2. We already saw that every vertex has degree 2. Since every vertex lies in $\overline{C}$, so does every end. By (iv) we find for every end ω a sequence of cuts $(F_i)_{i\in\mathbb{N}}$ such that the components of $G - F_i$ that contain ω converge to ω and such that each of these cuts contains precisely two edges of C. Thus, the degree of ω is at most 2 and hence precisely 2. This implies that the standard subspace $\overline{C}$ is a circle by Lemma 2.3. It is a Hamilton circle, as it contains every vertex and every end.

We define the C_i recursively. Besides this sequence, we shall define a second sequence $(\Lambda_i)_{i\in\mathbb{N}}$ of labellings $\Lambda_i \colon V \to \mathbb{N}$. For $i, q \in \mathbb{N}$ let

$$X_q^i := \{v \in V \mid \Lambda_i(v) = q\}.$$

Then the sequences will satisfy the following properties for every $1 \leq q$:

(Λ1) X_0^i is finite.

(Λ2) $C_i[X_q^i]$ is a subpath of C_i.

(Λ3) Any non-empty cut $E(X_q^i, V \setminus X_q^i)$ intersects with C_i precisely twice.

(Λ4) Every vertex in X_q^i not in C_i has distance more than $2k$ from $V \setminus X_q^i$.

(Λ5) Every vertex of C_i of depth $2k$ or more in C_i lies in X_0^i.

Let C_0 be any finite cycle of length at least k, which exists by assumption. For $j \in \mathbb{N}$ assume that all C_i and Λ_i with $i \leq j$ have been defined and that they satisfy (i), (ii), and (Λ2) to (Λ5). Let $B(i) \subseteq G$ for $i \in \mathbb{N}$ be the restriction of G to the vertices with distance at most i from C_j. Let $\mathcal{D} = \{D_1, \ldots, D_n\}$ be the set of infinite components of $G \setminus B(2k)$. Let $\mathcal{D}_f$ be the set of finite components of $G \setminus B(2k)$. Due to Lemma 3.1, we know that G is 2-connected. So according to Lemma 3.2, we find a cycle within every component $D_\ell \in \mathcal{D}$ of length at least k that has distance at least $2k$ to the boundary of D_ℓ, adding successively all vertices with distance at most k to this initial cycle via k-local skip-and-glue extensions we have a cycle C^ℓ in D_ℓ with at least k vertices of depth at least k in C^ℓ. Let $\mathcal{C} = \{C^1, \ldots, C^n\}$. Let $x_1, \ldots, x_m$ be an enumeration of some vertices of $\bigcup \mathcal{D}$ with

$$\left(B\left((n+5)\,k\right) \cap \bigcup \mathcal{D}\right) \cup \bigcup \mathcal{C} \subseteq G[x_1, \ldots, x_m]$$

such that for every $i \leq m$ each component of $G[x_1, \ldots, x_i]$ has a vertex in some C^ℓ. By our k-local skip-and-glue strategy, there is a sequence

$$C_j \cup \bigcup \mathcal{C} = F_0, \ldots, F_m = F$$

of finite 2-regular graphs such that $F_{i+1} = F_i$ if $x_{i+1} \in V(F_i)$, and such that F_{i+1} is obtained from F_i by a k-local skip-and-glue extension via x_{i+1} if $x_{i+1} \notin V(F_i)$. Note that all these extensions have length at most k and thus $C_j \subseteq F$. Clearly, F has at most $n+1$ components as F_0 has precisely $n+1$. Furthermore, the choice of the vertices x_i gives us that the vertices of F in D_i induce a connected graph in G for each $i \leq n$.

We proceed with a finite sequence $F = F^0, \ldots, F^p = F'$ of k-local skip-and-glue fusions: If some vertex u of $F^i \cap \bigcup \mathcal{D} \subseteq G - B(2k)$ has a neighbour in a different component of F^i, let F^{i+1} be a k-local skip-and-glue fusion of F^i centred around u. Since the number of components of F is at most $n+1$ and reduces by at least 1 with each skip-and-glue fusion we have $p \leq n+1$. Then we have the following properties:

(1) Every $D_i \in \mathcal{D}$ contains vertices from exactly one component of F' and all vertices within $B((n+5)k) \cap D_i$ lie in this component.

(2) $C_j = F' \cap B(m)$

(3) F' has at most $n+1$ components.

Next, we construct a finite sequence of k-local skip-and-glue extensions via vertices in $B((n+5)k) \cup \bigcup \mathcal{D}_f$ such that its first element is F' and its last element E contains every vertex of $B((n+5)k) \cup \bigcup \mathcal{D}_f$. As every extension has length at most k, we have the following properties:

(4) All vertices with depth at least k in C_j lie in a common component of E.

(5) Only one component of E meets $D_i \setminus B(3k)$ for every $i \leq n$.

(6) $V(E)$ is connected in G.

(7) E has at most $n+1$ components.

Let $E = E_0, \ldots, E_\ell = C_{j+1}$ be a sequence of 2-regular subgraphs of G such that E_ℓ is connected and E_{i+1} is a k-local skip-and-glue fusion centred around a vertex in the unique component U_i of E_i that contains all vertices with depth at least k in C_j. Note that this is well-defined as $V(U_i)$ is contained in the fused cycle U_{i+1}. The properties (i), (ii), and (iii) are direct consequences of the construction of C_{j+1} as $V(B((n+5)k)) \subseteq V(C_{j+1})$.

To proof (iv) let us define a finite sequence $\Lambda^0, \ldots, \Lambda^\ell = \Lambda_{j+1}$ of labellings $\Lambda^i \colon V \to \mathbb{N}$ (with the same ℓ as above). For $i, q \in \mathbb{N}$ let

$$Y_q^i := \{v \in V \mid \Lambda^i(v) = q\}.$$

We shall construct these labellings such that they satisfy the following conditions for $i \in \mathbb{N}$, $q > 0$, and $1 \leq p \leq \ell$. Note that $\ell \leq n$.

(Λ6) Y_0^i is finite.

(Λ7) $E_i[Y_q^i]$ is either a component of E_i or a subpath of U_i.

(Λ8) The intersection of any non-empty cut $E(Y^i_q, V \setminus Y^i_q)$ with E_i contains either none or precisely two edges.

(Λ9) Every vertex in Y^i_q not in E_i has distance more than $(n+2-p)k$ from $V \setminus Y^i_q$.

(Λ10) Every vertex of C_j of depth $2k$ or more in C_j lies in Y^i_0.

Let $K_0, \dots, K_p$ be the components of E_0 such that K_0 contains all vertices with depth at least k in C_j, cp. (4). Let us define a labelling $\Lambda^0 \colon V \to \mathbb{N}$. For a vertex $x \in V(K_i)$, let $\Lambda^0(x) = i$. For every component $D_i \in \mathcal{D}$, there is a unique label $\ell \in \mathbb{N}$ with $\Lambda^0(x) = \ell$ for all $x \in V(D_i \setminus B(3k))$, cp. (5). Let Λ^0 map every $y \in V(D_i \setminus E)$ to ℓ. Clearly, Λ^0 satisfies (Λ7) to (Λ10).

Assume Λ^p has been defined for $0 \leq p < \ell$. The cycle U_{p+1} is the union of a reduction R_p of U_p, a reduction R of some other component U of E_p and two disjoint paths P, Q minus two edges $e_p \in E(R_p)$ and $e \in E(R)$. If both end vertices of e_p have the Λ^p-label 0 let λ be the unique label of the vertices of U. Otherwise, let λ be the largest Λ^p-label at an end vertex of e_p. Let W be the set consisting of the vertices on P, Q, and R as well as all the vertices labelled with the label of U. Let Λ^{p+1} equal Λ^p outside of W and let Λ^{p+1} map every vertex of W to λ.

Clearly, properties (Λ6) and (Λ7) are kept throughout the construction. Moreover, (Λ8) is a direct consequence of (Λ7). Since all fusions are k-local we have (Λ9). Every vertex of depth at least k in C_j lies in all U_i and thus the fusions are centred around vertices with depth at most k in C_j as those vertices have neighbours outside of U_i. Since every vertex with depth $2k$ in C_j lies on the reduction of U_i for all k-local fusions of E_i, we have (Λ10). Thus we have (Λ1) to (Λ5) as $E_\ell = C_{j+1}$ is connected, $\Lambda^\ell = \Lambda_{j+1}$, and $\ell \leq n$.

It remains to prove (iv). Let us consider any finite set $S \subseteq V$ and any end ω of G. By (i) and (iii), we find an index j such that every vertex of S has depth more than $2k$ in C_j. So by (Λ5), it lies in X^{j+1}_0. The end ω lies in some infinite component K of $G - B((n+5)k)$. By (Λ4), all vertices of K have the same Λ_{j+1}-label q. As $\mathcal{B} = E(X^{j+1}_q, V \setminus X^{j+1}_q)$ is non-empty, it intersects with C_{j+1} precisely twice due to (Λ3). By (Λ4), the vertices incident with $\mathcal{B}$ have depth $2k$ or more in C_{j+1}. So (ii) implies that $\mathcal{B}$ meets C_i in precisely these two edges for every $i \geq j+1$. This shows (iv) and completes the proof of (I).

To prove (II) we pick a sequence $P + vw = F_0, \dots, F_n$ of cycles such that F_{i+1} is a k-local skip-and-glue extension of F_i and F_n covers every vertex of distance at most $2k+1$ from vw. Clearly, vw is an edge of F_n. Following the proof of (I) there is a sequence $F_n = C_0, C_1, \dots$ of finite cycles satisfying (i) to (iv). Thus for their limit C its closure $\overline{C}$ is a Hamilton circle and C contains vw by (ii). This completes the proof of (II) as $C - vw$ is a Hamilton arc of G with end vertices v and w. □

4 Locally connected graphs

For a subgraph F of a graph G, let us call a vertex $x \in V(F)$ *skippable* if x has degree 2 in F and its two neighbours y, z are adjacent in G but not in F. We call yz an *x-bypass*. Note that, if F is 2-regular, x is skippable if and only if yxz is a skippable path in F.

Lemma 4.1. *Every connected locally connected claw-free graph G has a 5-local skip-and-glue strategy.*

Furthermore, for $x, y \in V(G)$ the graph $G + xy$ has a 5-local skip-and-glue strategy with respect to $V(P)$ for some x–y path P that contains all vertices of distance at most 3 from x or y.

Proof. Let $G = (V, E)$ be a locally finite connected locally connected claw-free graph on at least three vertices. Let F be a finite subgraph of G with maximum degree 2 all whose components contain at least four vertices. Let C be any component of F. Note that C is either a cycle or a path. If C is a path we require its end vertices to have depth at least 3 in C and add the edge between those end vertices to G. Note that in the later construction, none of these new edges appears as they are to deep in their respective component. For the readability of the proof we omit their presence from now on and consider G to be the graph with these edges added and F to be 2-regular. With a slight stretch of terminology, we consider G to be claw-free, although it is G without these additional edges that is claw-free.

On each component C we choose one order $\leq_C$ of the two available canonical cyclic orders. For every vertex $v \in V(C)$ denote by v^- its predecessor with respect to $\leq_C$ and by v^+ its successor.

Let $u \notin V(F)$ be a vertex of G that has a neighbour v on F. To show that there is a k-local skip-and-glue strategy we have to provide a k-local skip-and-glue extension of F via u. As G is locally connected $N(v)$ is connected; it contains u, v^-, and v^+. Thus there is a shortest u–$\{v^-, v^+\}$ path $P = p_0p_1 \dots p_k$ with all its vertices in $N(v)$. We may assume that $p_0 = u$ and $p_k = v^+$ by choosing the other canonical order for the component containing u if necessary. Clearly, P does not contain v^-. By minimality P is induced and thus if $k \geq 4$, we have the claw $G[v, p_0, p_2, p_4]$ in G. Hence $k \leq 3$ and the length of P is at most 3.

If all inner vertices of P which lie on F are skippable and no two consecutive vertices of P are adjacent on F, then we can replace every skippable path $p_i^- p_i p_i^+$ by the edge $p_i^- p_i^+$. Furthermore, we replace the edge vv^+ by the path $vuPv^+$ to obtain a 4-local skip-and-glue extension of F via u. This covers the case that u is adjacent to v^- or v^+.

Thus we may assume that P has an inner vertex and since G is claw-free it holds that $v^-v^+ \in E(G)$. Since the component of F containing v has at least four vertices, v is skippable. If some inner vertex p_i of P lies on F and is not skippable, then the path $p_i^- p_i p_i^+$ is not skippable. Thus, we have $p_i^- p_i^+ \notin E$.

It remains to construct a k-local skip-and-glue extension if either some vertex of P is not skippable or some edge of P lies on F. Let p_i be the first vertex on P that is either not skippable or incident with an edge of $P \cap F$. As $G[v, p_i, p_i^-, p_i^+]$ is not a claw, there is an edge vp_i^- or vp_i^+ in the first case. As P lies in the neighbourhood of v, there is an edge vp_i^- or vp_i^+ in the second case, too. As v is skippable, we reduce F by replacing v^-vv^+ by the edge v^-v^+ and $p_j^- p_j p_j^+$ by the edge $p_j^- p_j^+$ for all $j < i$ where p_j is skippable. If v and p_i^- are adjacent, then we replace the edge $p_ip_i^-$ by the path $p_i^- vuPp_i$. Otherwise, v and p_i^+ are adjacent and we replace the edge $p_ip_i^+$ by the path $p_iPuvp_i^+$. Note that $i \leq 2$. Thus, we obtain a 4-local skip-and-glue extension of F via u in both cases.

Next, let us assume that F has more than one component each containing at least four vertices. Let $v \in V(F)$ be a vertex in a component K of F with a neighbour in some other component of F. Following the above argument for each vertex $u \in N(v) \cap V(F \setminus K)$ we have a 4-local skip-and-glue extension K' of K by P via u over some edge $e \in E(K)$ such that the inner vertices of P lie in $N(v) \cup \{v\}$ and u is the only vertex in P from $F \setminus K$.[1]

Similarly, there is a path Q from $\{u^-, u^+\}$ to one of the two neighbours of u on K', which are its neighbours on P.

For the following construction let us choose such P and Q with some minimality conditions:

(i) Let P be shortest possible.

(ii) With respect to (i) let Q be shortest possible.

Clearly, $|P| = 2$ if and only if there is an edge e in K with both its end vertices adjacent to u. Note that no inner vertex of Q lies on P, as this is in contradiction to the minimality of P.

Suppose that there is an inner vertex q of Q in K' that is either not skippable or incident with an edge of Q that lies in K'. As seen above $q \in V(K) \setminus V(P)$. Let a, b be the two neighbours of q on K'. As $G[u, q, a, b]$ is not a claw, there is an edge ua or ub if q is not skippable. As Q lies in the neighbourhood of u, there is an edge ua or ub if q is incident with an edge of Q on K', too. By symmetry we may assume that a is adjacent to u. If $qa \in E(K)$ we have a contradiction to the minimality of Q as we could have chosen $P = qua$ and shortened Q to end in q. Thus we may assume that $qa \in E(K') \setminus E(K)$. By the construction of K' the edge qa is a z-bypass for some $z \in V(P)$. This contradicts the minimality of P as $quPz$ is shorter since both u and z are inner vertices of P.

Thus every inner vertex of Q in K' is skippable and no edge of Q is an edge of K'.

If all inner vertices of Q are skippable in F and in K' and no edge of Q lies in $F \setminus K$ or K', let F' be the reduction of F where the inner vertices of Q and P in F are replaced by their bypasses and let L be the set consisting of the edge e and the edges from u to the end vertices of Q. Then $(F' \cup P \cup Q) - L$ is a skip-and-glue fusion of F. It is 5-local as Q has length at most 3 and every component of $P - L$ has length at most 2. Clearly, it is centred around v.

Thus we may assume that there is an inner vertex on Q in $F \setminus K$ that is not skippable or incident with an edge of $Q \cap (F \setminus K)$. First note that every vertex of Q is adjacent to a neighbour of u in F if u is not skippable in F. This implies, by the minimal choice of Q and as Q contains an inner vertex, that u is skippable in F.

Let q be the last inner vertex of Q in $F \setminus K$ that is not skippable or incident with an edge of $Q \cap (F \setminus K)$. Thus its subpath qQp with $p \in V(P)$ does not contain any such vertices. Let a, b be the two neighbours of q on F. As $G[u, q, a, b]$ is not a claw, there is an edge ua or ub if q is not skippable. As Q lies in the neighbourhood of u, there is an edge ua or ub if q is incident with an edge of Q on F, too. By symmetry we may assume that a is adjacent to u.

Let F' be the reduction of F where the inner vertices of Q and P in F are replaced by their bypasses and let $L = \{e, up, aq\}$. Then $(F' \cup P \cup qQp) - L + au$ is a skip-and-glue fusion of F. It is 4-local as qQp has length at most 2 and the

[1]Indeed, we get a path P completely contained in $N(v) \cup \{v\}$, and such a path of minimal length contains only one vertex from $F \setminus K$. We forget that its end vertices are contained in $N(v) \cup \{v\}$ for later convenience.

component of $P - up$ containing p has length at most 2 and the other component of $P - up$ has length at most 3. Clearly, it is centred around v. □

The following is a corollary extracted from the previous proof.

Corollary 4.2. *Every connected locally connected locally finite claw-free graph on at least k vertices contains a cycle of length at least k.* □

Now we can combine our previous results to prove our first main theorem.

Theorem 4.3. *Every connected locally connected locally finite claw-free graph on at least three vertices has a Hamilton circle.*

Proof. Let G be a connected locally connected locally finite claw-free graph on at least three vertices. By Corollary 4.2 the graph G contains a cycle of length $\min\{|G|, 5\}$. Due to Lemma 4.1, there is a 5-local skip-and-glue strategy for G. Thus, the assertion is a direct consequence of Theorem 3.3. □

Let us now turn our attention to Hamilton arcs. We deduce from the proof of Theorem 1.2 in [1] the following proposition that is valid for all locally finite graph.

Proposition 4.4. *Let G be a connected locally connected locally finite claw-free graph. For each two $x, y \in V(G)$ that do not disconnect G there is an x-y path of length at least 3 that contains $N(x) \cup N(y)$.* □

In the proof that the result of Proposition 4.4 implies that the graph has a Hamilton arc, Asratian showed in [1] the following:

Proposition 4.5. *Let G be a connected locally connected locally finite claw-free graph. For each two $x, y \in V(G)$ that do not disconnect G there is an x-y path of length at least 3 that contains all vertices of distance at most 2 to either x or y.* □

Using our terminology, Proposition 4.5 implies the existence of some x–y path P such that the depth in P of its end vertices is at least 3. This enables us to prove our second main theorem.

Theorem 4.6. *In every connected locally connected locally finite claw-free graph on at least 3 vertices, any two vertices that do not disconnect the graph are connected by a Hamilton arc.*

Proof. Let G be a connected locally connected locally finite claw-free graph and let $x, y \in V(G)$ be distinct vertices that do not disconnect G. As mentioned before, Proposition 4.5 implies that we find an x–y path P such that x and y have depth 3 in P. Since $G+xy$ has a 5-local skip-and-glue strategy with respect to $V(P)$ by Lemma 4.1 there is a sequence $P + xy = C_0, \dots, C_n$ of cycles in $G + xy$ such that C_{i+1} is a 5-local skip-and-glue extensions of C_i and such that C_n contains all vertices of distance at most 6 from x or y. Thus the assertion follows from Theorem 3.3. □

As a corollary of the previous theorem, we obtain the following theorem:

Theorem 4.7. *A connected locally connected locally finite claw-free graph on at least four vertices is Hamilton-connected if and only if it is 3-connected.* □

5 Further sufficient conditions for the existence of a Hamilton circle

In this section, we deduce some corollaries from the main theorems of Section 4. To shorten this section, we say that a graph G satisfies ($\star$) if the following statements are true:

(i) G has a Hamilton circle.

(ii) For each two vertices $u, v \in V(G)$ that do not separate G, there is a Hamilton u–v arc in $|G|$.

(iii) G is Hamilton-connected and $|V(G)| \geq 4$ if and only if it is 3-connected.

It is well-known that line graphs are claw-free. Thus, we directly obtain the following corollary (whose finite version for Hamilton cycles is due to Oberly and Sumner [21, Corollary 1]):

Corollary 5.1. *Let G be a locally finite connected locally connected line graph on at least three vertices. Then G satisfies ($\star$).* □

The proof that the assumptions of the following corollary imply that $L(G)$ is locally connected is the same as for finite graphs. Thus, we obtain the following corollary (whose finite version for Hamilton cycles is due to Oberly and Sumner [21, Corollaries 2 and 3]):

Corollary 5.2. *Let G be a locally finite connected graph on at least three vertices such that either every edge lies on a triangle or G is locally connected. Then $L(G)$ satisfies ($\star$).* □

For the following two corollaries the proofs that their assumptions imply that $L(L(G))$, $L(G^2)$, respectively, is locally connected is the same as for finite graphs, see [21, Corollaries 4 and 5]. The finite version for Hamilton cycles of Corollary 5.3 is due to Chartrand and Wall [6]) and that of Corollary 5.4 is due to Nebeský [20].

Corollary 5.3. *Let G be a locally finite connected graph with minimum degree at least 3. Then $L(L(G))$ satisfies ($\star$).* □

Note that for a graph G, its *square* G^2 has $V(G)$ as its set of vertices and two distinct vertices are adjacent in G^2 if their distance in G is at most 2.

Corollary 5.4. *Let G be a locally finite connected graph on at least three vertices. Then $L(G^2)$ satisfies ($\star$).* □

The last corollary of the results of Section 4 carries over the result by Matthews and Sumner [19, Corollary 1] for Hamilton cycles in finite graphs to locally finite graphs. Again, their proof carries over almost verbally.

Corollary 5.5. *Let G be a locally finite connected graph on at least three vertices. If G^2 is claw-free, then it satisfies ($\star$).* □

References

[1] A.S. Asratian, *Every 3-connected, locally connected, claw-free graph is Hamilton-connected*, J. Graph Theory **23** (1996), 191–201.

[2] E. Berger and H. Bruhn, *Eulerian edge sets in locally finite graphs*, Combinatorica **31** (2011), 21–38.

[3] R.C. Brewster and D. Funk, *On the hamiltonicity of line graphs of locally finite,* 6*-edge-connected graphs*, jgt **71** (2012), 182–191.

[4] H. Bruhn and M. Stein, *On end degrees and infinite circuits in locally finite graphs*, Combinatorica **27** (2007), 269–291.

[5] H. Bruhn and X. Yu, *Hamilton cycles in planar locally finite graphs*, SIAM J. Discrete Math. **22** (2008), 1381–1392.

[6] G. Chartrand and C.E. Wall, *On the Hamiltonian index of a graph*, Studia Sci. Math. Hungar. **8** (1973), 43–48.

[7] Q. Cui, J. Wang, and X. Yu, *Hamilton circles in infinite planar graphs*, J. Combin. Theory Ser. B **99** (2009), 110–138.

[8] R. Diestel, *Graph Theory*, 4th ed., Springer, 2010.

[9] R. Diestel, *Locally finite graphs with ends: a topological approach*, Discrete Math. **310–312** (2010–11), 2750–2765 (310); 1423–1447 (311); 21–29 (312), arXiv:0912.4213.

[10] R. Diestel and D. Kühn, *Graph-theoretical versus topological ends of graphs*, J. Combin. Theory (Series B) **87** (2003), 197–206.

[11] ______, *On infinite cycles I*, Combinatorica **24** (2004), 68–89.

[12] ______, *On infinite cycles II*, Combinatorica **24** (2004), 91–116.

[13] ______, *Topological paths, cycles and spanning trees in infinite graphs*, Europ. J. Combinatorics **25** (2004), 835–862.

[14] H. Freudenthal, *Über die Enden topologischer Räume und Gruppen*, Math. Zeitschr. **33** (1931), 692–713.

[15] A. Georgakopoulos, *Infinite hamilton cycles in squares of locally finite graphs*, Advances in Mathematics **220** (2009), 670–705.

[16] R. Halin, *A note on Menger's theorem for infinite locally finite graphs*, Abh. Math. Sem. Univ. Hamburg **40** (1974), 111–114.

[17] K. Heuer, *personal communication*, 2013.

[18] F. Lehner, *On spanning tree packings of highly edge connected graphs*, arXiv:1109.6787, 2011.

[19] M.M. Matthews and D.P. Sumner, *Hamiltonian results in $K_{1,3}$-free graphs*, J. Graph Theory **8** (1984), 139–146.

[20] L. Nebeský, *On the line graph of the square and the square of the line graph of a connected graph*, Časopis Pěst. Mat. **98** (1973), 285–287.

[21] D.J. Oberly and D.P. Sumner, *Every connected, locally connected nontrivial graph with no induced claw is Hamiltonian*, J. Graph Theory **3** (1979), 351–356.

[22] M. Stein, *Extremal infinite graph theory*, Discrete Math. **311** (2011), 1472–1496.

[23] C. Thomassen, *Hamiltonian paths in squares of infinite locally finite blocks*, Annals of Discrete Mathematics **3** (1978), 269–277.

Transitivity conditions in infinite graphs*

Matthias Hamann Julian Pott

Abstract

We study transitivity properties of connected graphs with more than one end. We completely classify the distance-transitive such graphs and, for all $k \geq 3$, the k-CS-transitive such graphs.

1 Introduction

A *k-distance-transitive graph* is a graph G such that for every two pairs (x_1, x_2) and (y_1, y_2) of vertices with distances $d(x_1, x_2) = d(y_1, y_2) \leq k$ there is an automorphism α of G with $x_i^\alpha = y_i$ for $i = 1, 2$, where x_i^α is the image of x_i under α. A graph is called *distance-transitive* if it is k-distance-transitive for all $k \in \mathbb{N}$. Macpherson [11] and Ivanov [8] independently classified the connected locally finite distance-transitive graphs. They are exactly the graphs $X_{k,l}$, the infinite graphs of connectivity 1 such that each block is a K^k, a complete graph on k vertices, and every vertex lies in l distinct blocks. Here, k and l are integers, but we shall use the notation of $X_{\kappa,\lambda}$ also when κ or λ are infinite cardinals.

Answering a question of Thomassen and Woess [16], Möller [13] showed that the 2-distance-transitive locally finite connected graphs with more than one end are still only the graphs $X_{k,l}$.

For graphs that are not locally finite, little is known. Our first main result is the following common generalization of the theorems of Macpherson and Möller to arbitrary connected graphs with more than one end:

Theorem 1.1. *Let G be a connected infinite graph with more than one end. The following properties are equivalent:*

(i) G is distance-transitive;

(ii) G is 2-distance-transitive;

(iii) $G \cong X_{\kappa,\lambda}$ for some cardinals κ and λ with $\kappa, \lambda \geq 2$.

A graph is called *n-transitive* or also *n-arc-transitive* if it has no cycle of length at most n and for every two paths $x_0 \ldots x_m$ and $y_0 \ldots y_m$ with $0 \leq m \leq n$ it admits an automorphism α with $x_i^\alpha = y_i$ for all i.

Thomassen and Woess [16] characterized the locally finite connected graphs with more than one end that are 2-transitive. These are precisely the r-regular trees for some $r \in \mathbb{N}$. As a consequence of Theorem 1.1 we get the following characterization of all such graphs, not necessarily locally finite, which we prove at the end of Section 3.

*Combinatorica 32 (6) (2012) 649-688

Corollary 1.2. *If G is a connected 2-transitive graph with more than one end, then G is a λ-regular tree for some cardinal $\lambda \geq 2$.*

In the second part of this paper we investigate graphs with the property that the existence of an isomorphism φ between two finite induced subgraphs implies that there is an automorphism ψ of the entire graph mapping one of the subgraphs to the other. This area divides into two parts: In one part φ has to induce ψ on these subgraphs, while in the other part they may differ. More precisely, a graph G is *k-CS-transitive* if for every two connected isomorphic induced subgraphs of order k some isomorphism between them extends to an automorphism of G. On the other hand, G is called *k-CS-homogeneous* if *every* isomorphism between two induced connected subgraphs of order k of G extends to an automorphism of G. A graph is *CS-transitive* if it is k-CS-transitive for all $k \in \mathbb{N}$, and *CS-homogeneous* if it is k-CS-homogeneous for all $k \in \mathbb{N}$. Furthermore, a graph is *end-transitive* if its automorphism group acts transitively on the set of its ends.

Gray [6] classified the connected locally finite 3-CS-transitive graphs with more than one end and showed that these graphs are end-transitive. He asked whether all locally finite k-CS-transitive graphs for $k \geq 3$ are end-transitive. We give a positive answer to his question, and also show that the ends of k-CS-transitive graphs of arbitrary cardinality have at most two orbits under the action of the automorphism group of the graph.

Since 1-CS-transitive graphs are the transitive graphs and 2-CS-transitive graphs are the edge-transitive graphs, there is not much hope to classify them. Thus we investigate $k \geq 3$. We shall give a complete classification of these k-CS-transitive graphs with more than one end. This is formulated in Theorem 1.3.

In order to state our characterization we have to introduce some classes of graphs. For a graph H let $X_{\kappa,\lambda}(H)$ be the graph which arises from the graph $X_{\kappa,\lambda}$ by replacing each vertex with a copy of H and adding all edges between two copies replacing adjacent vertices of $X_{\kappa,\lambda}$.

For $\kappa \geq 3$, let Y_κ denote a connected graph that has two different kinds of blocks, single edges and blocks that are complete graphs of order κ, and in which every vertex lies in exactly one block of each kind.

Let H_1, H_2 be graphs, and let $\kappa, \lambda \geq 2$ be cardinals. We construct the graph $Z_{\kappa,\lambda}(H_1, H_2)$ as follows. Let T be an infinite tree, viewed as a bipartite graph with bipartition A, B, and assume that the vertices in A have degree κ and the vertices in B have degree λ. We replace every vertex from A by an isomorphic copy of H_1 and every vertex from B by an isomorphic copy of H_2. We add all edges between vertices that belong to graphs that replaced adjacent vertices. The resulting graph is a $Z_{\kappa,\lambda}(H_1, H_2)$.

We also need some finite *homogeneous*[1] graphs. These are graphs such that any isomorphism between two finite induced subgraphs (not necessarily connected) extends to an automorphism of the whole graph. These graphs were determined by Gardiner [5]. Interestingly, Ronse [15] showed that the class of finite homogeneous graphs coincides with its 'transitive' counterpart, the class of graphs such that for any two isomorphic induced subgraphs (not necessarily connected) there *exists* an isomorphism between them that extends to an automorphism of the whole graph.

[1] *ultrahomogeneous* in [5]

The classes of finite homogeneous graphs featuring in our characterization will be the classes, denoted as $\mathcal{E}_{k,m,n}$, that occur in Enomoto's article [4] on combinatorially homogeneus graps. Each of these classes consists of all finite homogeneous graphs with the property that every vertex has at most m neighbours, every subgraph of order at least n is connected, and no two non-adjacent vertices have $k-2$ or more common neighbours. Furthermore, we exclude the complete graphs and the complements of complete graphs from $\mathcal{E}_{k,m,n}$ for technical reasons.

Now we are able to state our second main result, the classification of all connected k-CS-transitive graphs with more than one end if k is at least 3.

Theorem 1.3. *Let $k \geq 3$. A connected graph with more than one end is k-CS-transitive if and only if it is isomorphic to one of the following graphs*[2]*:*

(1) $X_{\kappa,\lambda}(K^1)$ with arbitrary κ and λ;

(2) $X_{2,\lambda}(K^n)$ with arbitrary λ and $n < \frac{k}{2}+1$;

(3) $X_{\kappa,2}(\overline{K^m})$ with arbitrary κ and $m < \frac{k+2}{3}$;

(4) $X_{2,2}(E)$ with $E \in \mathcal{E}_{k,m,n}$, $m \leq k-2$, $n < \frac{k-|E|}{2}+2$, and $2|E|-2 < k$;

(5) Y_κ with arbitrary κ (if k is odd);

(6) $Z_{2,2}(\overline{K^m}, K^n)$ with $2m+n \leq k+1$ (if k is even);

(7) $Z_{\kappa,\lambda}(K^1, K^n)$ with $n \leq k-1$, arbitrary κ, λ with $\kappa = 2$ or $\lambda = 2$ (if k is even);

(8) $Z_{2,2}(K^1, E)$ with $E \in \mathcal{E}_{k,m,n}$, $m \leq k-2$, $n \leq \frac{k}{2}+1$ (if k is even).

Gray [6] characterized the connected locally finite 3-CS-homogeneous graphs with more than one end. As a corollary of Theorem 1.3 we obtain in Section 7 the following classification of connected k-CS-homogeneous graphs for $k \geq 3$ with more than one end.

Corollary 1.4. *Let $k \geq 3$. A connected graph with more than one end is k-CS-homogeneous if and only if it is isomorphic to $X_{\kappa,\lambda}(H)$ for one of the following values of κ, λ and graphs H:*

(1) arbitrary κ and λ and $H \cong K^1$;

(2) $\kappa = 2$, arbitrary λ, $n < \frac{k}{2}$ and $H \cong K^n$;

(3) arbitrary κ, $\lambda = 2$, $m < \frac{k}{3}$ and $H \cong \overline{K^m}$;

(4) $\kappa = \lambda = 2$, $H \in \mathcal{E}_{k,m,n}$ for $m \leq k-2$, $n < \frac{k-|E|}{2}+1$, and $2|E| < k$.

Gray and Macpherson [7] classified the countable CS-homogeneous graphs[3]. Such graphs, connected and with more than one end, are those described in our Theorem 1.1 for countable cardinals κ, λ. As a further corollary of Theorem 1.3 we can extend their classification to arbitrary connected graphs with more than one end.

[2] By the definition of these graphs, κ and λ are at least 2 and κ is at least 3 in case (5).

[3] They call these graphs *connected-homogeneous* graphs.

Corollary 1.5. *For connected graphs with at least two ends the notions of being distance-transitive, CS-transitive, or CS-homogeneous coincide. (These graphs are described in Theorem 1.1.)* □

Let us say a word about the techniques we use for our proofs. The proofs of the corresponding theorems for locally finite graphs are all based on Dunwoody's *structure trees* corresponding to finite edge cuts that are invariant under the action of the automorphism group of the graph. This structure tree theory is described in the book of Dicks and Dunwoody [1]; see Möller [12, 14] and Thomassen and Woess [16] for introductions. Since those edge cuts must be finite, these structure trees can in general only be applied to locally finite graphs.

Recently, Dunwoody and Krön [3] developed a similar structure tree theory based on vertex cuts, providing a similarly powerful tool for the investigation of graphs that are not necessarily locally finite. We use this new theory in our proofs.

2 The structure tree

Throughout this paper we use the terms and notation from [2] if not stated otherwise. In particular, a *ray* is a one-way infinite path and a *double ray* is a two-way infinite path. Two rays in a graph G are *equivalent* if there is no finite vertex set S in G such that the two rays lie eventually in distinct components of $G-S$. (For an induced subgraph H and a subset S of the vertex set of G, we use $H-S$ to denote the induced subgraph $G[V(H)\setminus S]$ and $H+S$ to denote $G[V(H)\cup S]$.) The equivalence of rays is an equivalence relation whose classes are the *ends* of G.

Let G be a connected graph and $A,B\subseteq V(G)$ two vertex sets. The pair (A,B) is called a *separation* of G if

(i) $A\cup B=V(G)$ and

(ii) $E(G[A])\cup E(G[B])=E(G)$.

The *order* of a separation (A,B) is the cardinality of its *separator* $A\cap B$ and the *wings* of (A,B) are the induced subgraphs $G[A\setminus B]$ and $G[B\setminus A]$. With $(A,\sim)$ we refer to the separation $(A,(V(G)\setminus A)\cup N(V(G)\setminus A))$. A *cut* is a separation (A,B) of finite order with non-empty wings such that the wing $G[A\setminus B]$ is connected and such that no proper subset of $A\cap B$ separates the wings of (A,B)[4]. A *cut system* $\mathcal{S}$ is a non-empty set of cuts of G satisfying the following properties.

(1) If $(A,B)\in\mathcal{S}$ then there is an $(X,Y)\in\mathcal{S}$ with $X\subseteq B$.

(2) Let $(A,B)\in\mathcal{S}$ and C be a component of $G[B\setminus A]$. If there is a separation $(X,Y)\in\mathcal{S}$ with $X\setminus Y\subseteq C$, then the separation $(V(C)\cup N(C),\sim)$ is also in $\mathcal{S}$.

(3) If $(A,B)\in\mathcal{S}$ with wings X,Y and $(A',B')\in\mathcal{S}$ with wings X',Y' then either there is a component C in $X\cap X'$ and a component D in $Y\cap Y'$, or there is a component C in $Y\cap X'$ and a component D in $X\cap Y'$ such that both $(V(C)\cup N(C),\sim)$ and $(V(D)\cup N(D),\sim)$ are $\mathcal{S}$-separations.

[4] Dunwoody and Krön [3] call the corresponding vertex set $A\setminus B$ a *cut* and the set $B\setminus A$ the $*$-*complement* of this cut.

The cuts in $\mathcal{S}$ are also called *$\mathcal{S}$-separations* and an *$\mathcal{S}$-separator* is a vertex set that is the separator of an $\mathcal{S}$-separation.

Two vertices, vertex sets or subgraphs X, Y of G are *separated* by a separation (A, B)—not necessarily in $\mathcal{S}$—if either $X \subseteq A$ and $Y \subseteq B$, or $Y \subseteq A$ and $X \subseteq B$. They are separated *properly* if both, X and Y, meet components C and D of their corresponding wings such that every vertex in $A \cap B$ is adjacent to a vertex in C and a vertex in D. A vertex set S *separates* X and Y (*properly*) if there is a separation (A, B) with separator S that separates X and Y (properly). A vertex set or subgraph is *separated properly* by a separation (or its separator) if it contains two vertices that are separated properly by this separation.

Two separations $(A, B), (A', B')$ are *$\mathcal{S}$-nested* if there is one wing of each of them, W, W' say, such that both separators $A \cap B$ and $A' \cap B'$ are disjoint from $W \cup W'$ and such that there is no component C of $W \cap W'$ with $(C \cup N(C), \sim) \in \mathcal{S}$.[5] If it is clear which cut system we are referring to we may drop its identifier and speak of nested only. The cut system $\mathcal{S}$ is *nested* if each two $\mathcal{S}$-separations are ($\mathcal{S}$-)nested. If $\mathcal{S}$ is nested, then no $\mathcal{S}$-separation (A, B) separates any other $\mathcal{S}$-separator S properly, since S meets at most one wing of (A, B).

A cut in the cut system $\mathcal{S}$ is *minimal* if no other cut in $\mathcal{S}$ has smaller order. A *minimal cut system* is a cut system all of whose cuts are minimal and thus have the same order. If $\mathcal{S}$ is a minimal cut system, then the *order* $\operatorname{ord}(\mathcal{S})$ of $\mathcal{S}$ is the order of any of its cuts.

Remark 2.1. Let G be a transitive connected graph and let $\mathcal{S}$ be a nested cut system of G. Then any component of $G - S$ for an $\mathcal{S}$-separator S, is the wing of an $\mathcal{S}$-separation [3, Corollary 3.10]. In particular, for any two (nested) $\mathcal{S}$-separations (A, B) and (A', B') there is a wing of each of them, W, W' say, such that $W \subseteq W'$ or $W' \subseteq W$.

An *($\mathcal{S}$-)block* is a maximal induced subgraph X such that

(i) for every $(A, B) \in \mathcal{S}$ there is $V(X) \subseteq A$ or $V(X) \subseteq B$ but not both, that is X is not separated by any $\mathcal{S}$-separation;

(ii) there is some $(A, B) \in \mathcal{S}$ with $V(X) \subseteq A$ and $A \cap B \subseteq V(X)$.

Let $\mathcal{B}$ be the set of $\mathcal{S}$-blocks and let $\mathcal{W}$ be the set of $\mathcal{S}$-separators. If $\mathcal{S}$ is nested and minimal let $\mathcal{T}(\mathcal{S})$ be the graph with vertex set $\mathcal{W} \cup \mathcal{B}$ and edges WB ($W \in \mathcal{W}$ and $B \in \mathcal{B}$) if and only if $W \subseteq B$. Then $\mathcal{T} = \mathcal{T}(\mathcal{S})$ is called the *structure tree* of G and $\mathcal{S}$.

It is the same structure tree that is used by Dunwoody and Krön [3] but we use a different notation for the underlying cut system. They substantiate the term 'structure tree' in one of their theorems.

Theorem 2.2 ([3, Theorem 6.5])**.** *Let G be a connected graph, and let $\mathcal{S}$ be a nested minimal cut system of G. Then the structure tree of G and $\mathcal{S}$ is a tree.* □

We remark that this implies for every $\mathcal{S}$-separation (X, Y) that (if $\mathcal{S}$ is minimal and nested) there is an edge WB in $\mathcal{T}$ such that W is the $\mathcal{S}$-separator

[5] This means that there is no '$\mathcal{S}$-important' part of G that lies in $W \cap W'$; Dunwoody and Krön [3] call the vertex set of $W \cap W'$ an *isolated corner*.

$X \cap Y$ and $V(B) \subseteq X$. On the other hand, it follows from (2) of the definition of a cut system that for any $\mathcal{S}$-block B and any $\mathcal{S}$-separator $S \subseteq B$, there is an $\mathcal{S}$-separation (X, Y) with separator S such that $V(B) \subseteq X$.

In our proofs we use a certain kind of minimal cut system that was introduced by Dunwoody and Krön [3, Example 2.5].

Example 2.3. Let G be a connected infinite graph with at least two ends. Let $\mathcal{S}$ be the set of all cuts (A, B) such that both $G[A]$ and $G[B]$ contain a ray. Then $\mathcal{S}$ is a cut system.

We need a fundamental property of cut systems that is shown in [3, Theorem 8.6] by Dunwoody and Krön. Since we do not use the whole theorem, we only state the part that is applied in this paper.

Theorem 2.4. *Let G be a connected graph with at least two ends and let $\mathcal{C}$ be the cut system of G from Example 2.3. There is a nested cut system $\mathcal{S} \subseteq \mathcal{C}$ consisting only of minimal $\mathcal{C}$-separations that is invariant under* $\mathrm{Aut}(G)$ *such that if two ends are separated by a minimal cut in $\mathcal{C}$, then they are separated by a cut in $\mathcal{S}$.* □

For a connected graph G, a cut system is called *basic* if it is maximal with the following properties: it is nested, minimal and $\mathrm{Aut}(G)$-invariant, all of its separators lie in the same $\mathrm{Aut}(G)$-orbit, both wings of each cut contain a ray and the order of any cut is minimal with regard to separating two ends of G. We may state a useful corollary of Theorem 2.4 which we shall use in the later proofs without further mentioning.

Corollary 2.5. *Every connected graph with at least two ends has a basic cut system.* □

Let us investigate some properties of basic cut systems.

Lemma 2.6. *For a basic cut system $\mathcal{S}$ of a connected graph G with at least two ends and any $\mathcal{S}$-separator S, every component of $G - S$ that contains a ray is a wing of an $\mathcal{S}$-separation.*

Proof. For this proof we invoke [3, Lemma 3.9] which says that no separator of a nested cut system separates any other separator of that cut system properly. Let C be a component of $G - S$ containing a ray. We show that the separation $(V(C) \cup S, \sim)$ lies in $\mathcal{S}$. If there is an $\mathcal{S}$-separation (X, Y) whose separator S' meets C, then $S' \subseteq V(C) \cup S$ as $\mathcal{S}$ is nested and no two vertices of an $\mathcal{S}$-separator are separated properly by any $\mathcal{S}$-separator. Thus either X or Y is contained in $V(C) \cup S$ and $(V(C) \cup S, \sim)$ and (X, Y) are nested.

If there is an $\mathcal{S}$-separation (X, Y) whose separator S' does not meet C, then one wing of (X, Y) is disjoint from C and from S and thus (X, Y) and $(V(C) \cup S, \sim)$ are nested. Thus $(V(C) \cup S, \sim)$ is nested with all $\mathcal{S}$-separations. Clearly, there is a ray in $G - C$ as S is a separator of a basic cut system. Thus $\mathcal{S} \cup (V(C) \cup S, \sim)$ is nested, minimal and $\mathrm{Aut}(G)$-invariant, all of its separators lie in the same $\mathrm{Aut}(G)$-orbit, both wings of each cut contain a ray and the order of any cut is minimal with regard to separating two ends of G. As $\mathcal{S}$ is basic and thus maximal with these properties, it contains the cut $(V(C) \cup S, \sim)$. □

Together with Lemma 3.9 in [3] this implies the following lemma.

Lemma 2.7. *Let G be a connected graph with at least two ends and let $\mathcal{S}$ be a basic cut system of G. For any $\mathcal{S}$-separator S the components of $G - S$ that do not contain a ray are disjoint from any $\mathcal{S}$-separator.* □

For a basic cut system this lemma yields the following remark.

Remark 2.8. Let S, S' be two distinct $\mathcal{S}$-separators of a basic cut system $\mathcal{S}$ of a connected graph G. Then S' meets precisely one component of $G - S$ and this component contains a ray.

Lemma 2.9. *Let $\mathcal{S}$ be a basic cut system of a connected graph G with more than one end. Then, every finite vertex set separating two $\mathcal{S}$-separators separates two ends.*

In particular, less than $\text{ord}(\mathcal{S})$ *vertices do not separate any two $\mathcal{S}$-separators.*

Proof. Let S be a finite vertex set separating two distinct $\mathcal{S}$-separators S_1 and S_2. As $\mathcal{S}$ is nested and according to Remark 2.8, there is a component C_1 of $G - S_1$ containing an end ω_1 but no vertex of S_2 as well as a component C_2 of $G - S_2$ containing an end ω_2 and no vertex of S_1. Let $C = C_1 \cap C_2$. If C contains a vertex v, then for every $s \in S_1 \setminus S_2$ there is a v-s path with its inner vertices in C_1 as S_1 is minimal end separating. By the choice of C_1, this path contains no vertex from S_2. This implies that C_2 contains s contrary to the choice of C_2. Thus C is empty and $\omega_1 \neq \omega_2$.

Suppose that S does not separate ω_1 and ω_2. Then ω_1 and ω_2 live in the same component of $G - S$ and thus there is a double ray R with one tail in ω_1 and another one in ω_2 avoiding S. Every such double ray meets S_1 and S_2 as shown above. Hence R contains an S_1-S_2 path contradicting that S separates these two $\mathcal{S}$-separators. The last assertion holds, as $\mathcal{S}$ is basic, particularly as no vertex set of cardinality less than $\text{ord}(\mathcal{S})$ separates any two ends. □

The following lemma is proved in [3, Lemma 4.1]. We state it here as it nicely shortens some proofs.

Lemma 2.10. *For any k, every pair of vertices in a connected graph is separated properly by only finitely many distinct separators of order k.* □

2.1 Basic cut systems of special graphs

In Theorem 1.1 and 1.3 several classes of graphs arise. Let us give descriptions of basic cut systems and their structure trees for each of them.

The *building blocks* of $X_{\kappa,\lambda}(H)$ and $Z_{\kappa,\lambda}(H_1, H_2)$ are the isomorphic copies of H, H_1, and H_2 that are used for the construction of these graphs. For a Y^κ the copies of K^κ and the bridges are its *building blocks*.

Let G be isomorphic to $X_{\kappa,\lambda}(H)$ for $\kappa, \lambda \geq 2$ and a finite graph H. In this case there is a unique basic cut system of G. Its separators are the building blocks of the $X_{\kappa,\lambda}(H)$, and its separations are of the form $(V(C) \cup S, \sim)$, where S is any of the separators and C any component of $G - S$. Any block consists of the union of a maximal set of pairwise completely adjacent building blocks. The structure tree is a (semi-regular) tree of degrees κ and λ where the blocks have degree κ and the separators have degree λ.

Let G be isomorphic to Y_κ for $\kappa \geq 3$, then G is vertex transitive and every vertex is a separator of G that separates ends. The unique basic cut system has every single vertex as a separator and separations as in the example above. The blocks are precisely the building blocks. The structure tree is the κ-regular tree with every edge subdivided three times. The vertices of degree κ are the blocks corresponding to the K^κ and the vertices with distance two to them are the blocks corresponding to the K^2. The separators are precisely the vertices of the tree that are adjacent to a vertex of degree κ. The automorphism group has two orbits on the blocks. One orbit contains the building blocks of cardinality 2 and the other orbit those of cardinality κ. This shows that even though the automorphism group acts transitively on the separators it may not act transitively on the blocks.

Let G be isomorphic to $Z_{\kappa,\lambda}(H_1, H_2)$ for $\kappa, \lambda \geq 2$ and non-empty finite graphs H_1, H_2. In this case there may be two distinct basic cut systems, this happens only if $|H_1| = |H_2|$ and either $H_1 \ncong H_2$ or $\kappa \neq \lambda$. Then one may choose $i, j \in \{1, 2\}$ with $i \neq j$ arbitrarily and there is a basic cut system $\mathcal{S}$ of G with the building blocks corresponding to H_i as the $\mathcal{S}$-separators and the building blocks corresponding to H_j plus all its neighbours in G as the $\mathcal{S}$-blocks. If $H_1 \cong H_2$ and $\kappa = \lambda$, then $G \cong X_{2,\lambda}(H_1)$ and the basic cut system is as discussed above. If $|H_i| < |H_j|$ for $i, j \in \{1, 2\}$, then the building blocks corresponding to H_i are precisely the $\mathcal{S}$-separators and any building block corresponding to H_j plus all its neighbours is an $\mathcal{S}$-block. In both cases all cuts are of the form $(V(C) \cup S, \sim)$ where C is a component of the graph minus a separator S. The structure tree is a semi-regular tree with degrees κ and λ, where if H_1 corresponds to the separators they have degree κ and the blocks have degree λ and if H_2 corresponds to the separators the degrees swap.

3 Distance-transitive graphs

In this section we classify the connected distance-transitive graphs with more than one end (Theorem 1.1). Let us give a short outline of the proof, in particular of the implication that every connected 2-distance-transitive graph with more than one end is an $X_{\kappa,\lambda}$ for some cardinals κ and λ. Considering a basic cut system of such graphs, we show that its blocks are complete graphs and that any two of its separators are disjoint. We finish the proof by showing that all separators of the given cut system have cardinality 1 and have to lie in the same number of blocks and that each block consists of the same number of separators.

Proof of Theorem 1.1. Since the graphs $X_{\kappa,\lambda}$ are indeed distance-transitive and distance-transitive graphs are 2-distance-transitve by definition, it suffices to prove that every connected 2-distance-transitive graph with at least two ends is an $X_{\kappa,\lambda}$ for cardinals $\kappa, \lambda \geq 2$.

Let G be a connected 2-distance-transitive graph with more than one end. Let $\mathcal{S}$ be a basic cut system of G and let $\mathcal{T}$ be the structure tree of G and $\mathcal{S}$. In particular, for every separation $(A, B) \in \mathcal{S}$ and every automorphism α of G, the cuts $(A, B), (A^\alpha, B^\alpha)$ are nested and (A^α, B^α) lies also in $\mathcal{S}$. Furthermore, both wings of any cut in $\mathcal{S}$ contain a ray. As every 2-distance-transitive graph is vertex transitive by definition and thus every vertex lies in an $\mathcal{S}$-separator, which implies that every vertex lies in an $\mathcal{S}$-block.

Let us show first that all $\mathcal{S}$-blocks are complete graphs. Suppose not and let X be such an $\mathcal{S}$-block that is not complete. Let x, y be two non-adjacent vertices in X and let P be a shortest x-y path in G. To get to a contradiction let us find a block containing three consecutive vertices of P. If P is contained in X, let $Y = X$ and $a, b \in V(X \cap P)$ with $d(a, b) = 2$. If P is not contained in X, then there is an $\mathcal{S}$-separator, separating X and a vertex on P properly. Let S be such an $\mathcal{S}$-separator with maximal distance from X in $\mathcal{T}$. Then there is a component C of $G - S$ that avoids X and contains a vertex v from P for which $(V(C) \cup S, \sim)$ lies in $\mathcal{S}$, according to Lemma 2.6. Let Y be the neighbour of S in $\mathcal{T}$ contained in $C + S$, that is Y is the $\mathcal{S}$-block in $C + S$ containing S. The two neighbours of v on P lie in C or S and all vertices of $P \cap C$ lie in Y by the choice of S. Thus, three consecutive vertices on P, the vertex v and its two neighours a and b, lie in Y, and as P is an induced path a and b are not adjacent and $d(a, b) = 2$.

As $\mathcal{S}$ is a cut system, for every $\mathcal{S}$-separation (A, B) every vertex s in $A \cap B$ has a neighbour c in $A \setminus B$ and d in $B \setminus A$ such that these neighbours are separated properly by (A, B). As $cd \notin E(G)$, we have $d(c, d) = 2$. Since G is 2-distance-transitive, there is an automorphism α of G with $c^\alpha = a$ and $d^\alpha = b$. This contradicts the fact that Y is an $\mathcal{S}$-block as it is separated properly by (A^α, B^α) as c^α and d^α which are both contained in Y have to lie in distinct wings of (A^α, B^α). Thus all $\mathcal{S}$-blocks are complete.

Let us continue by showing that two distinct $\mathcal{S}$-separators S, S' are disjoint. Let (A, B) be an $\mathcal{S}$-separation and $\alpha \in \mathrm{Aut}(G)$ such that $S = A \cap B$ and $S^\alpha = S'$. These choices are valid since $\mathcal{S}$ is basic. As (A, B) and (A^α, B^α) are nested and G is transitive, we know by Remark 2.1 that there are wings, one of each of these two separations, W, W' say, that are disjoint. Suppose $S \cap S'$ is not empty and let $s \in S \cap S'$, $s' \in S' \setminus S$ and $w \in W$, $w' \in W'$ both adjacent to s. As all blocks are complete, s and s' are adjacent. Furthermore, w and s' are not adjacent, since they are separated by S. Thus, there is an automorphism β of G mapping (w, w') to (w, s'), since $ww', ws' \notin E(G)$. This is a contradiction according to Lemma 2.10 which says that there are only finitely many separators of cardinality $\mathrm{ord}(\mathcal{S})$ separating w and w' properly: The existence of β implies that there is the same finite number of $\mathcal{S}$-separators separating w from w' and w from s' properly. This does not hold since all $\mathcal{S}$-separators separating w and s' properly lie in the component of $G - S'$ that contains w and thus these separators also separate w and w' properly. On the other hand S^α separates w and w' properly while it does not separate w and s' properly. Thus, any two distinct $\mathcal{S}$-separators are disjoint.

In the next step let us show that all $\mathcal{S}$-separators have cardinality 1. Suppose not, then there are at least two vertices in some $\mathcal{S}$-separator S and, as all $\mathcal{S}$-blocks are complete, there is an edge e in $G[S]$. On the other hand, there is an edge e' that has precisely one of its end vertices in S. Since G is 2-distance-transitive it is also 1-distance-transitive and thus there is an automorphism α of G that maps e to e'. This is a contradiction, since S and S^α are neither disjoint nor the same. Thus all $\mathcal{S}$-separators have cardinality 1.

As G is 1-distance-transitive any two $\mathcal{S}$-blocks have the same order and 0-distance-transitivity implies that for every vertex the set of $\mathcal{S}$-blocks it lies in has the same cardinality λ. The order κ of an $\mathcal{S}$-block is at least 2, since there are edges in G and every $\mathcal{S}$-separator lies in at least two different $\mathcal{S}$-blocks. Thus G is isomorphic to $X_{\kappa,\lambda}$ for two cardinals $\kappa, \lambda \geq 2$. □

Next, we briefly deduce Corollary 1.2 from Theorem 1.1.

Proof of Corollary 1.2. A 2-transitive graph is also 2-distance-transitive and, if it has at least two ends, then it is an $X_{\kappa,\lambda}$ for cardinals $\kappa, \lambda \geq 2$. If $\kappa \geq 3$, then there is a path of length 2 in every block whose (adjacent) endvertices can be mapped onto vertices with distance 2 in distinct blocks. Since no adjacent vertices can be mapped onto vertices with distance 2 by any isomorphism, we know that $\kappa = 2$. The graphs $X_{2,\lambda}$ with $\lambda \geq 2$ are precisely the λ-regular trees. □

4 The local structure for some finite subgraphs

In some k-CS-transitive graphs the previously introduced finite homogeneous graphs play a role as building blocks. Enomoto [4] gave a combinatorially characterization of these homogeneous graphs. We apply a corollary of his result [4, Theorem 1] in our proofs.

For a subgraph X of a graph G let $\Gamma(X) = \bigcap_{x \in V(X)} N(x)$, which is the set of all vertices in G that are adjacent to all the vertices in X. A graph G is *combinatorially homogeneous* if $|\Gamma(X)| = |\Gamma(X')|$ for any two isomorphic induced subgraphs X and X'. Furthermore, a graph G is *l-S-transitive* if for every two isomorphic induced subgraphs of order l there is an automorphism of G mapping one onto the other.

Theorem 4.1. *[4, Theorem 1] Let G be a finite graph. The following properties of G are equivalent.*

(1) G is homogeneous;

(2) G is combinatorially homogeneous;

(3) G is isomorphic to one of the following graphs:

- *(a) a disjoint union of isomorphic complete graphs;*
- *(b) a complete t-partite graph K_r^t with r vertices in each partition class and with $2 \leq t, r$;*
- *(c) C_5;*
- *(d) $L(K_{3,3})$ (the line graph of $K_{3,3}$).* □

Whenever we need finite homogeneous graphs as building blocks for k-CS-transitive graphs we use Corollary 4.2 to handle them.

Corollary 4.2. *Let $k \geq 3$, $m \leq k-2$, and $n \leq \frac{k}{2}$ be positive integers. Let G be a finite graph with maximum degree at most m that is neither complete nor the complement of a complete graph. If G is l-S-transitive for all $l \leq k-1$, if any induced subgraph of G on at least n vertices is connected, and if any two non-adjacent vertices do not have $k-2$ common neighbours, then G is (combinatorially) homogeneous and isomorphic to one of the following graphs:*

(1) t disjoint K^r with $2 \leq t$, $1 \leq r-1 \leq m$, and $tr \leq n-1$;

(2) K_r^t with $2 \leq t$, $2 \leq r \leq n-1$, and $(t-1)r \leq \min\{m, k-3\}$;

(3) C_5 *with* $2 \leq m$ *and* $4 \leq n$*;*

(4) $L(K_{3,3})$ *with* $4 \leq m$ *and* $6 \leq n$*.*

Proof. Let us prove first that G is combinatorially homogeneous. If X and X' are isomorphic induced subgraphs of G both of order at most $k-1$, then l-S-transitivity for $l = |X|$ implies that there is an automorphism φ of G with $X^\varphi = X'$. Thus, we have $\Gamma(X)^\varphi = \Gamma(X')$ and $|\Gamma(X)| = |\Gamma(X')|$. If X and X' are isomorphic induced subgraphs of order at least k, then both $\Gamma(X)$ and $\Gamma(X')$ are empty because the maximum degree of G is at most $k-2$. This implies that G is combinatorially homogeneous and that we can apply Theorem 4.1 which provides that, ignoring the boundaries, there are no other cases as (1) to (4). The specific boundaries for each case can be checked easily. For example, in case (2) the '$k-3$' in the inequality $(t-1)r \leq \min\{m, k-3\}$ ensures that K_r^t does not contain two non-adjacent vertices with $k-2$ common neighbours if $m = k-2 = (t-1)r$. □

Let $\mathcal{E}_{k,m,n}$ be the class of all those graphs that satisfy the assumptions of Corollary 4.2 for the values k, m and n.

5 k-CS-transitivity for special graphs

This section is dedicated to showing that any graph on the list in Theorem 1.3 is indeed k-CS-transitive for the specific values of k.

Let G be a graph and $k \geq 3$. A graph H is *good for G* if for any two induced isomorphic copies H' and H'' of H in G there is an automorphism of G mapping H' onto H''. Clearly, a graph is k-CS-transitive if and only if all of its connected induced subgraphs of order k are good for it.

Lemma 5.1. *Let $k \geq 3$ and let G belong to one of the classes (1) to (8) of Theorem 1.3. The complete graph on k vertices is good for G.*

Proof. If G contains a complete graph on k vertices, then it is isomorphic to $X_{\kappa,\lambda}(K^1)$, $X_{2,\lambda}(K^n)$, $X_{\kappa,2}(\overline{K^m})$, Y_κ, $Z_{\kappa,2}(K^1, K^n)$, or $Z_{2,\lambda}(K^1, K^n)$ with the corresponding values for m and n.

- In $X_{\kappa,\lambda}(K^1)$ and Y_κ any complete graph on k vertices lies completely in some K^κ.
- In $X_{2,\lambda}(K^n)$, as $2n < k+2$, any complete graph on k vertices consists of precisely two building blocks or precisely two building blocks without one vertex depending on the parity of k.
- In $X_{\kappa,2}(\overline{K^m})$ any complete graph on k vertices has no two vertices in the same building block, and all its vertices in building blocks that are pairwise completely adjacent.
- In $Z_{\kappa,2}(K^1, K^n)$ and $Z_{2,\lambda}(K^1, K^n)$, as $n \leq k-1$, any complete graph on k vertices consists of precisely two adjacent building blocks.

In all these cases K^k is good for G by the construction of G. □

Lemma 5.2. *Let $k \geq 3$ and let G belong to one of the classes (1) to (8) of Theorem 1.3. Every connected graph on k vertices with diameter 2 is good for G.*

Proof. Let X be a connected induced subgraph of G on k vertices with diameter 2. If $G \cong Y_\kappa$ then X is isomorphic to some K^{k-1} with one edge attached. For any two such graphs in G, there is an automorphism of G mapping one to the other. Thus we may assume that $G \ncong Y_\kappa$.

If X is contained in a single building block, then—by cardinality and as it is neither complete nor the complement of a complete graph—$G \cong Z_{2,2}(K^1, E)$. Again by cardinality X lies in a building block corresponding to $E \cong K_r^t$ with $2 \leq t$ and $2 \leq r \leq \frac{k}{2}$ and $(t-1)r \leq k-3$. As $(t-1)r \leq k-3$ holds, there are at least 3 vertices of X in any of its necessarily t partition classes. This implies that for any (complete multipartite) induced subgraph Y of G isomorphic to X there is one building block containing Y, since—because of its diameter—it is contained in at most three building blocks and no building block corresponding to K^1 is contained in any complete multipartite induced subgraph of G that consists of t classes, each of which has cardinality at least 3. By the construction of G there is an automorphism α of G mapping the building block containing X to the building block containing Y and, as E is homogeneous, with $X^\alpha = Y$.

Therefore we may assume that X meets at least two building blocks. If X meets precisely two building blocks, then by cardinality $G \cong X_{2,2}(E)$ for some graph $E \in \mathcal{E}_{k,m,n}$ with $m \leq k-2$, $n < \frac{k-|E|}{2} + 2$ and $2|E| - 2 < k$, or k is even and $G \cong Z_{2,2}(K^1, E)$ for some graph $E \in \mathcal{E}_{k,m,n}$ with $m \leq k-2$ and $n \leq \frac{k}{2} + 1$.

In the first case, since $2|E| - 2 < k$, either X covers both building blocks it meets (if k is even) or it misses precisely one vertex in one of these building blocks (if k is odd). As E is homogeneous X is good for G.

In the second case there is one vertex v with $k-1$ neighbours that is the building block corresponding to K^1. As $n \leq \frac{k}{2} + 1$ we know that $X - v$ is connected. On the other hand Y contains a vertex with degree $k-1$ and thus is not contained in a single building block. Let $v' \in V(Y)$ be a vertex in a building block of G corresponding to K^1. As X and Y are isomorphic and any two vertices of degree $k-1$ in Y lie in the same $\mathrm{Aut}(Y)$-orbit it holds that $Y - v'$ is connected, and thus $Y - v'$ lies in a single building block of G. As above there is an automorphism α of G mapping the building blocks containing X to the building blocks containing Y with $(X - v)^\alpha = (Y - v')$.

Thus we may assume that X meets at least three building blocks. Let $B \subseteq G$ be a building block that is adjacent to all vertices of $X \setminus B$, which exists by the small diameter of X. If a separator in X does not contain every vertex of $X \cap B$, then it must contain at least all the vertices in $X \setminus B$ as every vertex of X in B is adjacent to every vertex of X not in B. Furthermore, the existence of a separator that separates $X \cap B$ properly implies that B is not complete. If the number $|X \cap B|$ is smaller than $|X \setminus B|$, then $X \cap B$ is the unique smallest separator and for every isomorphic induced copy Y of X in G precisely the vertices of $X \cap B$ are mapped to *the* smallest separator S in Y. We may assume that Y meets three building blocks, as it, and thus X is good for G otherwise. Since S is a smallest separator, we have $S = Y \cap D$ for the unique building block D of G that is adjacent to all vertices of $Y \setminus D$. Each of the smallest separators of these graphs either contains an edge, contains two non adjacent vertices, or is a single vertex. In all these cases B and D correspond to the same kind of building block by the construction of G. Since the building blocks are homogeneous and B is mapped to D by some automorphism of G, every isomorphism from X to Y extends to an automorphism of G. Thus we may assume that $X \cap B$ is not the unique smallest separator of X and also it is not complete.

Let us finish the remainder of the proof on a case by case analysis. The previous arguments cover (1), (2), (5), and (7) of Theorem 1.3. In (3) as $m < \frac{k+2}{3}$ and $k \geq 3$ it holds that $m < \frac{k}{2}$ and thus if there is a building block B, that separates X, then it is unique and $X \cap B$ is the smallest separator in X. If there is no such separating building block, then all building blocks that meet X are pairwise adjacent and X is a complete multipartite graph with at least three partition classes. As vertices of X lie in the same building block if and only if they are not adjacent, X is good for G.

In (4) there is a unique building block $B \cong E$ adjacent to all vertices in $X \setminus B$ and B separates X. If $X \cap B$ is not the smallest separator in X, then $\frac{k}{2} \leq |X \cap B|$ and as $2|E| - 2 < k$ it holds that $|B| < \frac{k}{2} + 1$ and thus $X \cap B = B$. The building block $B \cong E$ is connected, since $n < \frac{k-|E|}{2} + 2$. All connected graphs in $\mathcal{E}_{k,m,n}$ are 2-connected and thus any separator of X not containing $X \cap B$ contains $X \setminus B$ and at least two vertices from B and hence has at least $\frac{k}{2} + 1$ vertices. Again $X \cap B$ is the unique smallest separator in X, which completes this case.

For the case (6) that $G \cong Z_{2,2}(\overline{K^m}, K^n)$, if $n \neq 1$, then $m < \frac{k}{2}$ and thus $X \cap B$ is the smallest separator in X, as it is either complete or lies in a building block corresponding to $\overline{K^m}$ of order less than $\frac{k}{2}$. If $n = 1$, then B is either complete and the smallest separator or B is not complete and the two building blocks adjacent to B together with B cover X. Thus $|B| + 2 \geq k \geq 2m$ and this implies that $m = 2$ and $k = 4$. Since B is not complete it holds that $B \cong \overline{K^2}$ and $X \cong C_4$. Then it is easy to see that X can be mapped to every other copy of C_4 in $G \cong Z_{2,2}(\overline{K^2}, K^1)$ by some automorphism of G.

In (8) $G \cong Z_{2,2}(K^1, E)$ and we may assume that X meets two building blocks corresponding to K^1 and one other building block $B \cong E$, as otherwise the separating building block is complete, consists of only one vertex and is the unique smallest separator of X. Thus every induced subgraph Y of G isomorphic to X is good for G or meets precisely three building blocks, and—by the same arguments as above—two of these building blocks that Y meets correspond to the K^1. Any pair of non-adjacent vertices in X with $k - 2$ common neighbours in X, can be mapped to any other such pair by an automorphism of X. By the construction of G there is an automorphism α of G mapping the two building blocks corresponding to K^1 in X onto those in Y. As E is homogeneous and $X \cap B$ and $Y \cap B^\alpha$ are isomorphic, there is an automorphism of G mapping X onto Y.

Thus X is good for G in all cases. □

Lemma 5.3. *Let $k \geq 3$ and let G belong to one of the classes (1) to (8) of Theorem 1.3. Every connected graph on k vertices with diameter at least* 3 *is good for G.*

Proof. Let X and Y be isomorphic connected induced subgraphs of G on k vertices with diameter at least 3 and let α be an isomorphism from X to Y. If

X is a path, then there is an automorphism of G mapping X to Y according to the construction of G. Thus we may assume that X is not a path.

If $G \cong Y^\kappa$, then there is a maximal clique $K \subseteq X$ with at least 3 vertices. By the construction of Y^κ there is an automorphism α' of G that maps the building block containing K to the building block of G containing K^α and that is an extension of α.

If G is not isomorphic to a Y_κ, let P be a longest induced path in X whose diameter in X is at least 3. We show that every vertex v on P that lies in a building block corresponding to a finite graph B is mapped onto a vertex $v^\alpha \in V(Y)$ that also lies in a building block corresponding to B. This is easy in all the cases that have only one kind of building block. In particular, we have to proof this property in the cases (6), (7), and (8) of Theorem 1.3.

The path P meets at least four building blocks of G, since there is no building block B in any of the possible graphs with an induced path of length 3, except for the C_5, in which case $k > 5$ and X meets a building block adjacent to B and the diameter of $X \cap B$ in X is 2. As X is connected and not a path, there is a vertex v in $X - P$ that is adjacent to P. The cardinality of $N(v) \cap V(P)$ is $1, 2$, or 3, as P is induced and thus meets every building block in at most one vertex. In particular, these neighbours of v have distance at most 2 on P. Let us show that these cases determine in which kinds of building blocks the neighbours of v on P lie.

If v has only one neighbour p on P, then p is not a leaf of P by the maximiality of P. Furthermore, the vertices v and p do not lie in the same building block, as v would be adjacent to the same vertices on P as p otherwise. If $G \cong Z_{\kappa,\lambda}(K^1, K^n)$, then p lies in a building block corresponding to K^1 if and only if $\kappa > 2$, and in one corresponding to K^n if and only if $\lambda > 2$. If $G \ncong Z_{\kappa,\lambda}(K^1, K^n)$, then G belongs to one of the cases (6) or (8) and the vertex v lies in a building block that contains a leaf of P and thus two non-adjacent vertices. Hence p lies in a complete building block.

If v has two neighbours p_1, p_2 on P, and $d_P(p_1, p_2) = 2$ then the vertex on P adjacent to p_1 and p_2 lies in the same building block as v. This building block corresponds to the complement of a complete graph or a graph from $\mathcal{E}_{k,m,n}$ as it contains two non-adjacent vertices. If $d_P(p_1, p_2) = 1$ then one of p_1 or p_2 is a leaf of P and v lies together with this leaf in a common building block that corresponds to K^n or a graph from $\mathcal{E}_{k,m,n}$.

If v has three neighbours p_1, p_2, p_3 on P, then they induce a path of length 2 in P and v lies in the same building block as the middle vertex of that path of length 2 which is a building block corresponding to K^n or a graph from $\mathcal{E}_{k,m,n}$.

In all these cases, it is determined in which kind of building blocks of G the neighbours of v lie. Thus there is (at least) one vertex w on P such that w and w^α lie in building blocks that are in the same $\mathrm{Aut}(G)$-orbit of G. As in (6), (7), and (8) every second vertex on P lies in building blocks of the same $\mathrm{Aut}(G)$-orbit, it holds that for every w' on P the vertices w' and w'^α lie in the same kind of building block of G.

Using this path P, let us recursively construct an automorphism of G that maps X to Y. The arguments above show as all building blocks are homogeneous that there exists an automorphism α_0 of G with $\alpha_0|_P = \alpha|_P$ and that every such automorphism satisfies that p and p^{α_0} lie in building blocks that correspond to the same graph for every vertex $p \in V(P)$.

To define the automorphism α_l of G for $l \geq 1$ let α_i be defined for $i < l$. First, let W be the set of vertices in G with distance at most $l-1$ to the building blocks that contain P. The graphs X and Y induce graphs $X_1, \ldots, X_n$ and $Y_1, \ldots, Y_n$ with $X_j^\alpha = Y_j$ for all $1 \leq j \leq n$ in the components of $G-W$ and $G-W^{\alpha_{l-1}}$, respectively. Let α_l be an automorphism of G with $w^{\alpha_l} := w^{\alpha_{l-1}}$ for $w \in W$, that maps the component of $G-W$ containing X_j to the component of $G-W^{\alpha_{l-1}}$ containing Y_j for all $j \leq n$ so that the vertices of X adjacent to W are mapped precisely to those vertices of Y adjacent to $W^{\alpha_{l-1}}$. Since the diameter of X is less than k, the automorphism α_k of G maps X onto Y. □

Combining these lemmas we obtain the following corollary.

Corollary 5.4. *Let $k \geq 3$ and let G belong to one of the classes (1) to (8) of Theorem 1.3. Every connected graph on k vertices is good for G.*
In particular, G is k-CS-transitive. □

6 The global structure of k-CS-transitive graphs

This section contains the substantial part of the proof of Theorem 1.3. We show that for $k \geq 3$ every connected k-CS-transitive graph with at least two ends is isomorphic to one of the graphs described in Theorem 1.3. At first, we provide some general properties for basic cut systems of such graphs. Later on we distinguish two fundamentally different cases: in Subsection 6.1 we look at those graphs that are covered by the separators of a basic cut system and in Subsection 6.2 at those that are not.

Lemma 6.1. *Let $k \geq 3$. If G is a connected k-CS-transitive graph with at least two ends, then for G and any of its basic cut systems their structure tree has no leaves.*

Proof. Let $\mathcal{S}$ be a basic cut system of G and let $\mathcal{T}$ be the structure tree of G and $\mathcal{S}$. Suppose that $\mathcal{T}$ has a leaf X. By the construction of a structure tree, X is an $\mathcal{S}$-block. Let $(A, B) \in \mathcal{S}$ be a cut with $V(X) \subseteq A$ and $A \cap B \subseteq V(X)$. By the construction of $\mathcal{T}$, we know that X is adjacent to all $\mathcal{S}$-separators that are contained in X. This implies that $A \cap B$ is the only $\mathcal{S}$-separator in X and $V(X) = A$. In particular, no vertex of $A \setminus B = V(X - B)$ lies in an $\mathcal{S}$-separator as $\mathcal{S}$ is nested. Since there is a ray in $G[A]$, the block X is infinite. There is no vertex in X that has distance $k+1$ to B, as otherwise an induced path in $G[A]$ starting at $v \in A \cap B$ could be mapped into $X - B$ by an automorphism of G. The image of $A \cap B$ under this automorphism is not an $\mathcal{S}$-separator as it contains a vertex from $X - B$. This contradicts the $\mathrm{Aut}(G)$-invariance of the basic cut system $\mathcal{S}$. Thus there are vertices of infinite degree in X. Let $x \in V(X)$ be a vertex with infinite degree and minimal distance to B with this property. Let N be an infinite set of neighbours of x with $d(v, B) > d(x, B)$ for all $v \in N$. By the infinite version of Ramsey's Theorem (see for example [2, Theorem 9.1.2]) there is either a $K^{\aleph_0}$ or an infinite independent set in $G[N]$. Suppose there is an independent set of cardinality $k-1$ in N. As $d(v, B) > d(x, B)$ for all $v \in N$, there is a neighbour u of x with $d(u, B) < d(x, B)$ if $d(x, B) \geq 1$ or with $u \in B \setminus A$ if $x \in A \cap B$ such that u is not adjacent to any vertex in N. Any $k-2$ independent vertices in N together with x and u induce a subgraph that could be mapped onto a subgraph induced by x and $k-1$ independent vertices in N. The former subgraph is either properly separated by an $\mathcal{S}$-separator while the latter is not, or it is closer to any $\mathcal{S}$-separator than the latter one. Thus there is no independent set of cardinality $k-1$ in N and there is a $K^{\aleph_0}$ in $G[N]$. Again, this yields to a contradiction. Indeed, let H be a complete graph on k vertices in $G[N]$, and let $v \in V(H)$. Then there is no automorphism of G that maps $H - v + x$ to H as $H - v + x$ contains only vertices of distance at least $d(x, B) + 1$ to the unique $\mathcal{S}$-separator in X, which is a contradiction to the k-CS-transitivity of G. □

Lemma 6.2. *Let $k \geq 3$, let G be a connected k-CS-transitive graph with at least two ends, and let $\mathcal{S}$ be a basic cut system of G. Then any ray in the structure tree of G and $\mathcal{S}$ contains infinitely many pairwise disjoint $\mathcal{S}$-separators.*

In particular, G does not have finite diameter.

Proof. Let $\mathcal{T}$ be the structure tree of G and $\mathcal{S}$ and let R be a ray in $\mathcal{T}$. The only neighbours of $\mathcal{S}$-blocks in $\mathcal{T}$ are $\mathcal{S}$-separators. Thus, infinitely many different (finite) $\mathcal{S}$-separators lie on R.

Suppose that there is a vertex x in G that lies in infinitely many of the separators on R. Let S_0 be the first separator on R that contains x, and let X be an $\mathcal{S}$-block adjacent to S_0 in $\mathcal{T}$ that does not lie on the tail S_0R of R with initial vertex S_0 which exists as $\mathcal{T}$ has no leaf and thus S_0 has at least two neighbours in $\mathcal{T}$. Let $(A, B) = (V(C) \cup S_0, \sim)$ be the $\mathcal{S}$-separation with separator S_0 for the component C of $G - S_0$ that meets X. Then all separators on R that contain x lie in B. As (A, B) is a cut, there exists a neighbour $y \in A \setminus B$ of x.

Any $\mathcal{S}$-separator separates two $\mathcal{S}$-separators in G properly if and only if it separates them properly in the structure tree. Hence, if a vertex of G lies in two separators, then it also lies in any separator that appears in the structure tree on the unique path between those two. Every $\mathcal{S}$-separator $S \subseteq B$ on R contains x, as it lies between S_0 and one of the infinitely many other separators containing x on R. There is a neighbour y_S of x such that S separates y and y_S properly. By Lemma 2.10 we know that the number of $\mathcal{S}$-separators separating v and w properly is finite for all vertices $v, w \in V(G)$. This implies that there is an infinite set $\mathcal{F}$ of $\mathcal{S}$-separators on R such that for any two distinct separators $S, S' \in \mathcal{F}$ the vertices $y_S, y_{S'}$ are distinct. Thus, $U := \{y_S \in V(G) \mid S \in \mathcal{F}\}$ is infinite. By the infinite version of Ramsey's Theorem there is either a $K^{\aleph_0}$ or an infinite independent set in $G[U]$. In the first case, let $K \subseteq G[U]$ be such an infinite complete graph. The (finite) $\mathcal{S}$-separators do not separate K properly and hence there are infinitely many $\mathcal{S}$-separators separating y from K properly. As K is infinite and all separators in $\mathcal{F}$ have the same cardinality, there exists a vertex $v \in V(G)$ that lies outside of infinitely many separators in $\mathcal{F}$. Each of the infinitely many separators S in $\mathcal{F}$ for which y_S lies in $V(K)$ and that does not contain v separate y and v properly, as every such separator separates y and y_S properly and $vy_S \in E(G)$. This is a contradiction as y and v are separated properly by infinitely many separators of cardinality $\mathrm{ord}(\mathcal{S})$.

Thus there is an infinite independent set $U' \subseteq U$ completely adjacent to x. Remember that y is not adjacent to any vertex in U'. We choose a subset V_1 of U' of cardinality $k-1$. There is a maximal number n of separators of cardinality $\mathrm{ord}(\mathcal{S})$ that separate any two vertices of V_1 properly as for each of the finitely many pairs of vertices in V_1 there is only a finite number of separators of cardinality $\operatorname{ord}(\mathcal{S})$ that separates it properly. Let V_2 be another subset of U' of cardinality $k-2$ that contains a vertex that is separated by more than n separators of cardinality $\operatorname{ord}(\mathcal{S})$ from y properly: pick a separator S in $\mathcal{F}$ such that on the S_0-S path on R there are more than n other $\mathcal{S}$-separators and let $y_S \in V_2$. By k-CS-transitivity there is an automorphism of G that maps $G[V_2 \cup \{x, y\}]$ onto $G[V_1 \cup \{x\}]$ as both these induced subgraphs are stars with $k-1$ leaves. This automorphism has to fix x and map $V_2 \cup \{y\}$ onto V_1. As y and y_S are separated properly by more than n separators of cardinality $\operatorname{ord}(\mathcal{S})$, their respective images in V_1 are separated properly by just as many such separators. This contradicts the choice of n.

Thus no vertex of G lies in infinitely many $\mathcal{S}$-separators on R and we conclude that there are infinitely many pairwise disjoint $\mathcal{S}$-separators on R. Two $\mathcal{S}$-separators S_1, S_2 that have n disjoint $\mathcal{S}$-separators on their S_1-S_2 path in $\mathcal{T}$ have distance at least n in G. As by Lemma 6.1 every structure tree of a basic cut system of G contains a ray, this implies the second assertion. □

The next lemma provides a fundamental tool in the proof of Theorem 1.3.

Lemma 6.3. *Let $k \geq 3$, let G be a connected k-CS-transitive graph with at least two ends, let $\mathcal{S}$ be a basic cut system of G, and let S be an $\mathcal{S}$-separator. For every ray R in the structure tree $\mathcal{T}$ of G and $\mathcal{S}$ that starts at S, there are $\operatorname{ord}(\mathcal{S})$ disjoint induced rays $R_1, \ldots, R_{\operatorname{ord}(\mathcal{S})}$ in G starting at S such that for every $i \leq \operatorname{ord}(\mathcal{S})$ the ray R_i intersects with all $\mathcal{S}$-separators on R.*

Proof. On R there are infinitely many disjoint $\mathcal{S}$-separators $S_1, S_2, \ldots$ all disjoint from $S_0 := S$ as shown in Lemma 6.2. As by Lemma 2.9 no two of them are separated by less than $\operatorname{ord}(\mathcal{S})$ many vertices, Menger's Theorem implies that there are $\operatorname{ord}(\mathcal{S})$ many pairwise disjoint induced S_0-S_i paths for all $0 < i$. Let $\mathcal{P}_i$ be the subgraph of G consisting of these paths. Since $\mathcal{P}_{i-1}$ covers S_{i-1} on R, we may choose $\mathcal{P}_i$ such that $\mathcal{P}_{i-1} \subseteq \mathcal{P}_i$. The union $\bigcup_{i \in \mathbb{N}} \mathcal{P}_i$ is a subgraph of $\operatorname{ord}(\mathcal{S})$ many pairwise disjoint induced rays each starting at S_0. Clearly, each of those rays intersects with every $\mathcal{S}$-separator on R. □

For a connected k-CS-transitive graph G with $k \geq 3$ and at least two ends and a basic cut system $\mathcal{S}$ of G, there are two profoundly different cases. In the first case the graph is covered with $\mathcal{S}$-separators while in the second case there are vertices in G that do not belong to any $\mathcal{S}$-separator.

For an $\mathcal{S}$-block X we define the *open ($\mathcal{S}$-)block*

$$\mathring{X} := X - \bigcup \{A \cap B \mid (A, B) \in \mathcal{S}\}.$$

Further down the line it turns out that the two cases above correspond to whether there exist non-empty open blocks or not. In Lemma 6.9 we get rid of any vertices that lie neither in an $\mathcal{S}$-separator nor in an open $\mathcal{S}$-block. In the proof of Theorem 1.1 we got this property for free as the graphs considered there are vertex transitive; here it turns out to require some effort.

Lemma 6.4. *Let $k \geq 3$, let G be a connected k-CS-transitive graph with at least two ends, and let $\mathcal{S}$ be a basic cut system of G. If the $\mathcal{S}$-separators do not cover G, then k is even.*

Proof. Let k be odd. We show that every vertex lies in some $\mathcal{S}$-separator. By Lemma 6.3 there is an induced ray R meeting infinitely many vertices that lie in $\mathcal{S}$-separators. As k is odd, there is an induced path $P \subseteq R$ of length $k-1$ whose middle vertex v belongs to some $\mathcal{S}$-separator. We may map the path anywhere into the ray and thus know that there are k succeeding vertices on the ray that belong to $\mathcal{S}$-separators. Thus every induced path of length $k-1$ in G has all its vertices in $\mathcal{S}$-separators. As the diameter of G is not finite according to Lemma 6.2, every vertex lies on an induced path of length $k-1$. Therefore every vertex lies in some $\mathcal{S}$-separator. □

Lemma 6.5. *Let $k \geq 3$, let G be a connected k-CS-transitive graph with at least two ends, and let $\mathcal{S}$ be a basic cut system of G. If G has a vertex not in any $\mathcal{S}$-separator, then every edge on every induced path of length $k-1$ in G has precisely one of its incident vertices in an $\mathcal{S}$-separator.*

In particular, if there is a vertex not in any $\mathcal{S}$-separator, then every induced path $P \subseteq G$ of length at least $k-1$ alternates between vertices in $\mathcal{S}$-separators and vertices outside every $\mathcal{S}$-separator.

Proof. As there is a vertex outside every $\mathcal{S}$-separator, k is even by Lemma 6.4. Since the structure tree $\mathcal{T}$ of G and $\mathcal{S}$ has no leaves by Lemma 6.1, every $\mathcal{S}$-separator S lies on a double ray in $\mathcal{T}$. By Lemma 6.3 there are two induced rays R_1, R_2 starting at $s \in S$ with all their other vertices in two distinct components of $G - S$. Let $R = R_1 \cup R_2$, that means R is an induced double ray in G. As G is k-CS-transitive, there are automorphisms of G mapping a path $P \subseteq R$ of length $k-1$ with its middle edge incident with s to any other path of length $k-1$ on R. Thus every edge on R is incident with a vertex in an $\mathcal{S}$-separator.

If some edge on R has both its incident vertices in $\mathcal{S}$-separators, this implies by the same argument that every edge on R has both its incident vertices in $\mathcal{S}$-separators. As there is a vertex $v \in V(G)$ not contained in any $\mathcal{S}$-separator and as the diameter of G is not bounded, there is an induced path P of length $k-1$ starting at v. Thus, as by k-CS-transitivity there is an automorphism of G mapping P to R, there is no edge on R with both its incident vertices in $\mathcal{S}$-separators. Hence every edge on R has precisely one of its incident vertices in an $\mathcal{S}$-separator. By k-CS-transitivity, every edge on an induced path of length $k-1$ in G has precisely one of its incident vertices in an $\mathcal{S}$-separator. Thus any induced path of length at least $k-1$ is such an *alternating* path. □

As a corollary of the proof of the previous lemma, we obtain the following result.

Corollary 6.6. *Let $k \geq 3$, let G be a connected k-CS-transitive graph with at least two ends, and let $\mathcal{S}$ be a basic cut system of G. Then, any two vertices on an induced path P of length $k-1$ that have a vertex from an $\mathcal{S}$-separator between them on P are separated by some $\mathcal{S}$-separator in G.*

In particular, if two vertices on P have distance at least 3 on P, then they are separated by an $\mathcal{S}$-separator in G.

Proof. We recall the definitions from the proof of Lemma 6.5: We have a double ray R in G such that for every vertex $s \in V(R)$ that lies in an $\mathcal{S}$-separator, there is one $\mathcal{S}$-separator S with $s \in S$ such that the two components of $R - s$ lie in distinct components of $G - S$. Remark that we obtain such a double ray also in the situation that G is covered by $\mathcal{S}$-separators.

The two components of $R - s$ do not meet any common $\mathcal{S}$-block for any vertex s on R that lies in an $\mathcal{S}$-separator. Thus for any path P of length $k-1$ and any vertex s' in the interior of P, that belongs to some $\mathcal{S}$-separator, the two components of $P - s'$ are separated properly by an $\mathcal{S}$-separator. This implies the first assertion, the second one follows immediately since any two vertices on an induced path of length $k-1$ with distance at least 3 have—by Lemma 6.5 or as every vertex lies in an $\mathcal{S}$-separator—a vertex from some $\mathcal{S}$-separator between them on P. □

Corollary 6.7. *Let $k \geq 3$, let G be a connected k-CS-transitive graph with at least two ends, and let $\mathcal{S}$ be a basic cut system of G such that some vertex of G is not contained in any $\mathcal{S}$-separator. If two vertices belong to different sets of $\mathcal{S}$-separators then they are not adjacent.*

Proof. Suppose that there is an $\mathcal{S}$-separator S, and vertices $s \in S$ and $s' \notin S$ but in a different $\mathcal{S}$-separator such that s and s' are adjacent. Then there is an induced path of length $k-1$ that contains this edge and lies otherwise in a component of $G - S$ that does not contain s'. But no such path exists according to Lemma 6.5. □

Lemma 6.8. *Let $k \geq 3$, let G be a connected k-CS-transitive graph with at least two ends, let $\mathcal{S}$ be a basic cut system of G, and let S be an $\mathcal{S}$-separator. Every component C of $G - S$ contains an end and the separation $(V(C) \cup S, \sim)$ lies in $\mathcal{S}$.*

In particular, S separates any two vertices in distinct components of $G - S$ properly.

Proof. The claim is true if the $\mathcal{S}$-separators cover G, because then, for every component C of $G-S$, there is an $\mathcal{S}$-separator S' that meets C. Let $(A', B') \in \mathcal{S}$ with separator S'. According to Remark 2.8 the separator S' lies in $C \cup S$ and hence one wing of (A', B') lies in C. Thus, C contains an end and according to Lemma 2.6 is the wing of an $\mathcal{S}$-separation. Since $\mathcal{S}$ is a cut system and C is a component of $G - S$, we know that $(V(C) \cup S, \sim)$ is an $\mathcal{S}$-separation.

Hence, we may assume that there is a vertex outside every $\mathcal{S}$-separator. By Lemma 6.4 this implies that k is even. Let (A, B) be an $\mathcal{S}$-separation with $A \cap B = S$. As $\mathcal{S}$ is a cut system, there is an $\mathcal{S}$-separation (A', B') such that $A' \subseteq B$ and $S \subseteq A' \cap B'$. Since $\mathcal{S}$ is minimal, it holds that $S = A' \cap B'$.

First, let us assume that $G - S$ consists of precisely two components. Both components contain an end as every $\mathcal{S}$-separator separates two ends. Thus, we have $(A, B) = (B', A')$ and the assertion holds.

Therefore, we may assume that $G - S$ contains at least three components. Then there exists a component C of $G - S$ that lies neither in A nor in A'. Let $c \in V(C)$ be a vertex that is adjacent to $s \in S$, and let $d \in A' \setminus B'$ be adjacent to s. By Lemma 6.3 there is an induced path P in G of length $k-1$ that starts at a vertex $x \in A \setminus B$ and ends in c while meeting S only in s. This implies that P has only one vertex in C which is c. By the choice of P and d, the graph Psd is also an induced path of length $k-1$. As G is k-CS-transitive, there is an automorphism $\alpha \in \mathrm{Aut}(G)$ mapping P to Psd. By Lemma 6.5 both these paths alternate between vertices in $\mathcal{S}$-separators and other vertices. Thus, precisely one endvertex of P and one of Psd is contained in an $\mathcal{S}$-separator as k is even and we have $c^\alpha = d$. The component C^α of $G - S^\alpha$ contains $d = c^\alpha$. Both S and S^α

separate x from d. Suppose $S \neq S^\alpha$, then either S separates S^α from x properly or S^α separates S from x properly according to Remark 2.8. If S separates S^α from x let $S_1 = S$, $S_2 = S^\alpha$, and $\beta = \alpha$. Otherwise, if S^α separates S from x let $S_1 = S^\alpha$, $S_2 = S$, and $\beta = \alpha^{-1}$. Let S_0 be the $\mathcal{S}$-separator that contains x and—among all those—is closest to S_1 in the structure tree $\mathcal{T}$ of G and $\mathcal{S}$.

The separator S_0^β contains x as $x^\alpha = x$, and thus lies in A. By the choice of S_0 and as S_1 separates S_0 and S_2, it holds that

$$d_\mathcal{T}(S_0, S_1) = d_\mathcal{T}(S_0^\beta, S_2) > d_\mathcal{T}(S_0^\beta, S_1),$$

since the path between S_2 and S_0^β in $\mathcal{T}$ has to contain S_1. Thus, S_0^β is closer to S_1 in $\mathcal{T}$ than S_0 contradicting the choice of S_0. This implies that $S = S^\alpha$ and $(V(C) \cup S, \sim) = (A', B')^{\alpha^{-1}}$. As every component of $G - S$ is a wing of an end separating cut the component contains an end and thus any two components of $G - S$ are separated properly by S. □

Lemma 6.9. *Let $k \geq 3$, let G be a k-CS-transitive graph with at least two ends, and let $\mathcal{S}$ be a basic cut system of G. Then every vertex of G lies in an $\mathcal{S}$-block.*

Proof. Let v be a vertex of G. If v belongs to some $\mathcal{S}$-separator, it lies in an $\mathcal{S}$-block. So we may assume that v lies outside every $\mathcal{S}$-separator. Let S and S' be two distinct $\mathcal{S}$-separators such that S' separates S and v. By Lemma 6.8 and as $\mathcal{S}$ is nested, S' separates S and v properly. There are only finitely many $\mathcal{S}$-separators separating an $\mathcal{S}$-separator and v properly according to Lemma 2.10 and thus there are only finitely many $\mathcal{S}$-separators separating S and v, as v lies outside every $\mathcal{S}$-separator. As $\mathcal{S}$ is basic and by Remark 2.8 there is one $\mathcal{S}$-separator S_0 that separates S and v such that no other $\mathcal{S}$-separator separates S_0 and v. We show that v and S_0 lie in a common $\mathcal{S}$-block. Let C_0 be the component of $G - S_0$ that contains v. Then $(V(C_0) \cup S_0, \sim)$ lies in $\mathcal{S}$ according to Lemma 6.8. There is an $\mathcal{S}$-block X adjacent to S_0 in the structure tree of G and $\mathcal{S}$ whose vertices lie in $V(C_0) \cup S_0$. This block contains S_0 and, as there is no $\mathcal{S}$-separator separating v from $S_0 \subseteq X$, there is no $\mathcal{S}$-separator separating v from X. Thus v lies in X. □

Corollary 6.10. *Let $k \geq 3$, let G be a k-CS-transitive graph with at least two ends, and let $\mathcal{S}$ be a basic cut system of G. There is a non-empty open $\mathcal{S}$-block if and only if there is a vertex outside every $\mathcal{S}$-separator.* □

This corollary shows that the distinction between '$\mathcal{S}$-separators cover G' and 'there is a vertex outside every $\mathcal{S}$-separator' is in fact a distinction between whether all open $\mathcal{S}$-blocks are empty or not. For this reason we characterize the cases by stating if there is a non-empty open $\mathcal{S}$-block or not from now on. In addition we use the fact that for all cut systems we investigate every vertex lies in a block without refering to Lemma 6.9.

In the construction of $X_{\kappa,\lambda}(H)$ and $Z_{\kappa,\lambda}(H_1, H_2)$ the appropriate copies of H and H_1, H_2, respectively, are completely adjacent. The next lemma provides the corresponding property for k-CS-transitive graphs.

Lemma 6.11. *Let $k \geq 3$, let G be a connected k-CS-transitive graph with at least two ends, and let $\mathcal{S}$ be a basic cut system of G. Let X be an $\mathcal{S}$-block, let S be an $\mathcal{S}$-separator with $S \subseteq X$, and let $s \in S$. If $\mathring{X} = \emptyset$, let $x \in V(X - S)$ and if $\mathring{X} \neq \emptyset$, let $x \in V(\mathring{X})$. Then s and x are adjacent.*

Proof. Let $\mathcal{T}$ be the structure tree of G and $\mathcal{S}$. Suppose s and x are not adjacent in G. Let P be a shortest s-x path whose inner vertices lie in the component of $G - S$ that contains x. As P is a shortest path it is induced. Let C be a component of $G - S$ not containing x. As there is an induced ray R starting at s with all its other vertices in C, there is an induced path P' of length at least $k - 1$ starting at x and containing P.

Suppose that the distance of s and x on P is 2. The common neighbour y of s and x on P does not lie in any $\mathcal{S}$-separator, because of Corollary 6.6 and as s and x do lie in a common $\mathcal{S}$-block. As every vertex lies in an $\mathcal{S}$-block according to Lemma 6.9, this implies that y lies in an open $\mathcal{S}$-block $\mathring{Y}$. By Lemma 6.4 k is even and hence at least 4. As $y \in V(\mathring{Y})$ its neighbours s and x have to lie in Y. Suppose $Y \neq X$, then let $S' \subseteq V(Y)$ be an $\mathcal{S}$-separator separating X and Y and thus containing s and x. As every component of $G - S'$ contains a ray and does not have finite diameter (according to Lemma 6.1, 6.3, and 6.8) there is an induced ray starting at y and avoiding S'. Furthermore we require R to have precisely one other vertex in Y not adjacent to s, which is possible as Lemma 6.5 implies that the neighbour x' of y on R lies in an $\mathcal{S}$-separator and Corollary 6.7 implies that there is no edge between s and x' while all the other vertices on R are separated properly from S' by an $\mathcal{S}$-separator that contains x'. Let P_1 be the subpath of P' that contains x and has length $k-1$. Let v be the other endvertex of P_1 and let $P_2 = x'yP_1v$. Then there is an automorphism α of G that maps P_1 onto P_2. By Lemma 6.5 the automorphism α has to map x to x' and fix the remainder of P_1. By the same lemma, y does not lie in the same $\mathcal{S}$-block as v. So there is an $\mathcal{S}$-separator separating y and v properly. Every such separator lies together with X in the same component of $\mathcal{T} - Y$. As $\mathcal{S}$ is minimal and nested, every $\mathcal{S}$-separator that separates x and v properly also separates any other vertex in S' and v. Thus, according to Remark 2.8, every $\mathcal{S}$-separator that separates x and v properly separates v and S' properly. Since S' separates v and x' properly, Remark 2.8 also implies that every $\mathcal{S}$-separator that separates v and S' properly also separates v and x' properly. By k-CS-transitivity and according to Lemma 2.10, the same finite number of $\mathcal{S}$-separators separates v from x as from x' properly. This yields a contradiction as S' separates x' and v properly but not x and v. This contradiction shows that $X = Y$ and that $\mathring{X}$ is not empty. Thus, we have $x \in \mathring{X}$ and with Lemma 6.5 this implies that s and x have odd distance as they lie on the alternating path P', in particular $d_P(s, x) \neq 2$.

Therefore, the distance between x and s on P is at least 3. If the length of P is at most $k - 1$, then we may choose P' as above of length precisely $k - 1$. By Corollary 6.6 the vertices x and s are properly separated by some $\mathcal{S}$-separator.

Thus, we may assume that P has length at least k. As P contains a subpath of length $k-1$ containing x, there has to be an $\mathcal{S}$-separator separating a vertex on P from x properly. Let S' be an $\mathcal{S}$-separator furthest away in $\mathcal{T}$ from X such that there is a vertex on P separated properly by S' from X. Let C be a component of $G - S'$ that meets P and avoids X. Then $(V(C) \cup S', \sim) \in \mathcal{S}$ and there is an $\mathcal{S}$-block $Y \subseteq G[C + S']$ adjacent to S' in $\mathcal{T}$. By the choice of S' all vertices of $P \cap C$ lie in Y and S' separates X and Y. In particular P has a vertex y in $Y - S'$ such that $d_P(y, s)$ is smallest possible. Let y_1 be the neighbour of y on yPx. As no induced subpath of length 3 on P lies in one $\mathcal{S}$-block by Corollary 6.6, the neighbour of y_1 on y_1Px must not lie in Y and thus not in $V(C) \cup S'$. Hence we have $y_1 \in S'$.

As above there is an induced ray starting at y and having no other vertex adjacent to S' than y. Let y_2 be the neighbour of y on that ray. Again, we may elongate—if necessary—y_1Ps in C to obtain an induced path P_1 of length $k-1$ that ends in y_1 and either contains y_1Ps or lies on it. Let v be the other endvertex of P_1 and let $P_2 = y_2yP_1$ be the same path as P_1 with y_1 substituted by y_2. As both subgraphs are induced paths of length $k-1$, there is an automorphism α of G mapping P_1 onto P_2. This automorphism has to map the endvertices of P_1 to the endvertices of P_2. By a similar argument as above, we obtain that the number of $\mathcal{S}$-separators that separate v and y_2 properly and the number of $\mathcal{S}$-separators that separate v and y_1 properly differ which is a contradiction as $P_1^\alpha = P_2$. This contradiction shows that x and s are adjacent. □

Corollary 6.12. *Let $k \geq 3$, let G be a connected k-CS-transitive graph with at least two ends, and let $\mathcal{S}$ be a basic cut system of G. Then any two distinct $\mathcal{S}$-separators are disjoint.*

Proof. Suppose that there are two distinct $\mathcal{S}$-separators S, S' that are not disjoint. Every $\mathcal{S}$-separator on the S-S' path in $\mathcal{T}$ contains $S \cap S'$. Thus we may assume that $d_{\mathcal{T}}(S, S') = 2$ and hence S and S' lie in a common $\mathcal{S}$-block X. Let $s \in S \cap S'$ and let x_1 be a neighbour of s in a component of $G - S$ avoiding S'. Let $x_2 \in \mathring{X}$ if $\mathring{X}$ is not empty, and if $\mathring{X} = \emptyset$, then let x_2 be a vertex in $S \setminus S'$. By Lemma 6.11 the vertices x_2 and s are adjacent in both cases. Let P be an induced path of length $k-2$ in G that starts at s and has its other vertices in a component of $G - S'$ avoiding S which exists according to Lemma 6.1 and 6.3. Since G is k-CS-transitive there is an automorphism of G mapping x_1P to x_2P, as both are induced paths in G of length $k-1$. Similar to the proof of the previous lemma and as the $\mathcal{S}$-blocks cover G according to Lemma 6.9, the endvertices of x_1P and those of x_2P are separated by a different finite number of $\mathcal{S}$-separators. By contradiction, this shows that $\mathcal{S}$-separators are either equal or disjoint. □

Let $k \geq 3$, let G be a connected k-CS-transitive graph with at least two ends, and let $\mathcal{S}$ be a basic cut system of G. A *k-spoon* is an induced subgraph of G that consists of a triangle and a path of length $k-2$, its *handle*, starting at one of its triangle vertices with all in all precisely k vertices. A k-spoon H *pokes* in an $\mathcal{S}$-block X, an $\mathcal{S}$-separator S, or two $\mathcal{S}$-separators S, S' if two of its degree 2 vertices[6] of the triangle are contained in $\mathring{X}$, S, or one in S and one in S', respectively. A *k-fork* is another induced subgraph of G on k vertices that consists of its *prongs*, a pair of two non-adjacent vertices, and of its *handle*, a path such that both prongs are adjacent only to the same endvertex of the handle. A k-fork H *pokes* in an $\mathcal{S}$-block X, an $\mathcal{S}$-separator S, two $\mathcal{S}$-blocks X, Y, or two $\mathcal{S}$-separators S, S' if its prongs are contained in $\mathring{X}$, in S, meet $\mathring{X}$ and $\mathring{Y}$, or meet S and S', respectively.

[6]Remark that for $k > 3$ there are precisely two such vertices, but for $k = 3$ a k-spoon is just the triangle.

6.1 Empty open blocks

In this subsection we investigate k-CS-transitive graphs that have a basic cut system all of whose open blocks are empty. Remember that by Lemma 6.4, this is the only case if k is odd.

Lemma 6.13. *Let $k \geq 3$, let G be a connected k-CS-transitive graph with at least two ends, and let $\mathcal{S}$ be a basic cut system of G. If every open $\mathcal{S}$-block is empty, then all $\mathcal{S}$-blocks lie in the same* $\mathrm{Aut}(G)$*-orbit, or k is odd and there is a cardinal $\kappa \geq 3$ such that $G \cong Y_\kappa$.*

Proof. Suppose that there are two $\mathcal{S}$-blocks X and Y that lie in distinct $\mathrm{Aut}(G)$-orbits. As every $\mathcal{S}$-block contains an $\mathcal{S}$-separator and $\mathcal{S}$ is basic, there is an automorphism φ of G with $X \cap Y^\varphi = S$ for an $\mathcal{S}$-separator S. Hence we may assume that $X \cap Y = S$. If S contains two distinct vertices, then by Lemma 6.3 there is either a k-spoon with its triangle—the subgraph isomorphic to a K^3—in X and one k-spoon with its triangle in Y or there is a k-fork with both edges incident with its prongs in X and one such k-fork for Y, such that in each case the handle does not contain any vertex from S. As G is k-CS-transitive, there is, in both cases, an automorphism α of G mapping one edge in X that does not lie in any $\mathcal{S}$-separator to one such edge in Y. Thus $X^\alpha \cap Y$ is not contained in an $\mathcal{S}$-separator and $X^\alpha = Y$.

Hence two distinct $\mathcal{S}$-blocks intersect in at most one vertex and $\mathrm{ord}(\mathcal{S}) = 1$. By Lemma 6.11 and as every open $\mathcal{S}$-block is empty, any two $\mathcal{S}$-separators in a common block are completely adjacent and thus every $\mathcal{S}$-block is complete. For any two $\mathcal{S}$-blocks each of which has more than two vertices, there is a k-spoon with its triangle in each of these $\mathcal{S}$-blocks, respectively. Thus these blocks are $\mathrm{Aut}(G)$-isomorphic as G is k-CS-transitive.

Let P be an induced double ray in G whose edges alternate between two orbits of $\mathcal{S}$-blocks. Such a double ray exists, as one may start at any vertex of G and add appropriate edges greedily, since every vertex lies in blocks of all orbits of blocks. Clearly, every induced path of length $k-1$ shares this property with the ray and thus every vertex lies in at most one block of each orbit. As otherwise, if there is a vertex that lies in more than one block of the same orbit, then one may construct an induced path of length $k-1$ without this property. With the same argument for any two kinds of orbits, there is an induced path of length $k-1$ with edges only in these orbits. Since G is k-CS-transitive, this implies that there are precisely two distinct orbits of $\mathcal{S}$-blocks: in one orbit each $\mathcal{S}$-block is isomorphic to a K^2 and in the other one each $\mathcal{S}$-block is isomorphic to a K^κ for some cardinal $\kappa \geq 2$. If $\kappa = 2$ then G is a double ray and this contradicts that there are two distinct $\mathrm{Aut}(G)$-orbits of $\mathcal{S}$-blocks. Thus $\kappa \geq 3$ and $G \cong Y_\kappa$.

Let us suppose that k is even. Then there is a path of length $k-1$ with both outermost edges in $\mathcal{S}$-blocks isomorphic to a K^2 and there is a path of length $k-1$ with both outermost edges in $\mathcal{S}$-blocks isomorphic to a K^κ with $\kappa \geq 3$. As no automorphism of G maps one of these paths to the other, this is a contradiction and hence k is odd. □

Lemma 6.14. *Let $k \geq 3$, let G be a connected k-CS-transitive graph with at least two ends, and let $\mathcal{S}$ be a basic cut system of G such that every open $\mathcal{S}$-block is empty. If any two $\mathcal{S}$-blocks lie in the same* $\mathrm{Aut}(G)$*-orbit, then $G \cong$*

$X_{2,2}(H)$ for some finite graph H that is neither complete nor the complement of a complete graph, or there are cardinals $\kappa, \lambda \geq 2$ and integers $2 \leq m < \frac{k+2}{3}$ and $2 \leq n < \frac{k}{2} + 1$ such that $G \cong X_{2,\lambda}(K^n)$ or $G \cong X_{\kappa,2}(\overline{K^m})$ or $G \cong X_{\kappa,\lambda}(K^1)$.

Proof. Let $H = G[S]$ for some $\mathcal{S}$-separator S. According to Lemma 6.11 and Corollary 6.12 it holds that $G \cong X_{\kappa,\lambda}(H)$ for some cardinals $\kappa \geq 2$ and $\lambda \geq 2$. We may assume that $G \not\cong X_{2,2}(H)$ where H is neither complete nor the complement of a complete graph. If there are edges in H and $\lambda \geq 3$ then there are two kinds of k-spoons: one with its triangle meeting three $\mathcal{S}$-separators and one meeting precisely two $\mathcal{S}$-separators. If there are two non-adjacent vertices in H and $\kappa \geq 3$ then there are two kinds of k-forks: one pokes in a single separator and one pokes in two different separators. As G is k-CS-transitive, all k-spoons as well as all k-forks lie in one $\mathrm{Aut}(G)$-orbit, respectively. Thus it holds that either $G \cong X_{\kappa,2}(\overline{K^m})$ with $m \geq 2$, or $G \cong X_{2,\lambda}(K^n)$ with $n \geq 2$, or $G \cong X_{\kappa,\lambda}(K^1)$. It remains to show that $m < \frac{k+2}{3}$ and $n < \frac{k}{2} + 1$.

Let $G \cong X_{\kappa,2}(\overline{K^m})$ and suppose that $m \geq \frac{k+2}{3}$. Let S_1, S_2 be $\mathcal{S}$-separators in different $\mathcal{S}$-blocks both (completely) adjacent to an $\mathcal{S}$-separator S_0. As $3m \geq k+2$ there are sets $A_i \subseteq S_i$ for $i = 0, 1, 2$ such that $A_1 \cup A_0 \cup A_2$ has cardinality $k+2$, is connected in G—that is $A_0 \neq \emptyset$—and such that each of A_1 and A_2 contains at least two vertices. Let $a, b \in A_1$ and $c \in A_2$. By the construction of G it holds that

$$G[(A_1 \setminus \{a,b\}) \cup A_0 \cup A_2] \cong G[(A_1 \setminus \{a\}) \cup A_0 \cup (A_2 \setminus \{c\})].$$

As there is no automorphism of G mapping the first to the second graph, this is a contradiction and thus $m < \frac{k+2}{3}$.

Let $G \cong X_{2,\lambda}(K^n)$ and suppose $n \geq \frac{k}{2} + 1$. Let S_0, S_1 be two (completely) adjacent $\mathcal{S}$-separators. Let $A_i \subseteq S_i$ with $|A_0| = \lceil \frac{k}{2} \rceil + 1$ and $|A_1| = \lfloor \frac{k}{2} \rfloor - 1$, and let $B_i \subseteq S_i$ with $|B_0| = \lceil \frac{k}{2} \rceil$ and $|B_1| = \lfloor \frac{k}{2} \rfloor$ which exist as $n \geq \frac{k}{2} + 1$ implies that $n \geq \lceil \frac{k}{2} \rceil + 1$ for any integer n. It holds that $|A_0 \cup A_1| = |B_0 \cup B_1| = k$, but there is no automorphism of G that maps the complete graph on k vertices $G[A_0 \cup A_1]$ to the complete graph on k vertices $G[B_0 \cup B_1]$. By contradiction we obtain that $n < \frac{k}{2} + 1$. □

Lemma 6.15. *Let $k \geq 3$, let G be a connected k-CS-transitive graph with at least two ends, and let $\mathcal{S}$ be a basic cut system of G such that every open $\mathcal{S}$-block is empty and all $\mathcal{S}$-blocks lie in one orbit of $\mathrm{Aut}(G)$. If $G \cong X_{2,2}(E)$ for some finite graph E that is neither complete nor the complement of a complete graph, then $2|E| - 2 < k$ and $E \in \mathcal{E}_{k,m,n}$ for $m \leq k-2$ and $n < \frac{k-|E|}{2} + 2$.*

Proof. By Corollary 4.2 if it holds that (a) the maximum degree of E is at most $k-2$, (b) E is l-S-transitive for all $l \leq k-1$, (c) any induced subgraph of order at least $\frac{k-|E|}{2} + 1$ in E is connected, and (d) no two non-adjacent vertices of E have $k-2$ common neighbours, then $E \in \mathcal{E}_{k,m,n}$ for $m \leq k-2$ and $n < \frac{k-|E|}{2} + 2$. Considering the distinct boundaries in (b) and for n, we note that a graph on at least $\frac{k-|E|}{2} + 1$ vertices has at least n vertices.

(a) Let S ($G[S] \cong E$) be an $\mathcal{S}$-separator. Suppose there is a vertex v of degree at least $k-1$ in $G[S]$. Let $A \subseteq S$ consist of v and $k-1$ of its neighbours. Let w be some vertex from an $\mathcal{S}$-separator that is adjacent to S. Then

$G[A] - v + w$ is isomorphic to $G[A]$, but there is no automorphism of G mapping one onto the other. Thus no vertex in S has degree at least $k-1$.

(b) Let $A, B \subseteq S$ induce isomorphic graphs with at most $k-1$ vertices for some $\mathcal{S}$-separator S ($G[S] \cong E$). Then there is a common neighbour v of these vertices in an adjacent $\mathcal{S}$-separator S_0. Let P be an induced path of length $k-1-|A|$ that starts at v and each of its other vertices is separated properly from $A \subseteq S$ by S_0. By construction of $X_{2,2}(E)$, the path P meets each $\mathcal{S}$-separator in at most one vertex. As G is k-CS-transitive, there is an automorphism α of G that maps $G[P + A]$ to $G[P + B]$. If $|A| \neq 1$, then α must map S onto S as it is the only $\mathcal{S}$-separator meeting more than one vertex of $G[P + A]$ and of $G[P + B]$; clearly this implies $A^\alpha = B$ and A and B lie in the same $\mathrm{Aut}(G[S])$-orbit. If $|A| = 1$, let S' be an $\mathcal{S}$-separator such that some induced path of length $k-1$ starting at A ends in S'. Let φ, φ' be the isomorphisms from E to S, S', respectively. Let $A' \subseteq S'$ be $(A^{\varphi^{-1}})^{\varphi'}$. Then we may assume that the path P ends in A'. Thus α maps A to B or A' to B and as A and A' are $\mathrm{Aut}(G)$-isomorphic so are A and B. Again A and B lie in the same $\mathrm{Aut}(G[S])$-orbit.

(c) Suppose there is an induced subgraph $X \subseteq E$ of order at least $\frac{k-|E|}{2} + 1$ that is not connected. Let S_0, S_1, S_2 be three distinct $\mathcal{S}$-separators such that S_0 is adjacent to the other two. Let $A_i \subseteq S_i$ for $i \geq 1$ be of cardinality at least $\frac{k-|E|}{2} + 1$ such that $G[A_1] \cong G[A_2]$ are not connected. Let H be a connected induced subgraph on $k+2$ vertices in $G[S_0 \cup A_1 \cup A_2]$ such that there is an isomorphism φ from $H[A_1]$ onto $H[A_2]$ and $H[A_1]$ is not connected. Such a graph exists as $|S_0 \cup A_1 \cup A_2| \geq |E| + 2(\frac{k-|E|}{2} + 1) = k+2$. Let $a, b \in A_1$ be vertices that lie in distinct components of $H[A_1]$. Then there is no automorphism of G that maps one of its two isomorphic induced and connected subgraphs $H - \{a, b\}$ and $H - \{a^\varphi, b\}$ onto the other. Thus every induced subgraph of E of order at least $\frac{k-|E|}{2} + 1$ is connected.

(d) Suppose that there are two non-adjacent vertices x, y in an $\mathcal{S}$-separator S' ($G[S'] \cong E$) with at least $k-2$ common neighbours in S' and let $N \subseteq S'$ be $k-2$ of these neighbours. Let S, S'' be distinct $\mathcal{S}$-separators adjacent to S' and let $s \in S$ and $s'' \in S''$. Then $G[N \cup \{x, y\}]$ and $G[N \cup \{s, s''\}]$ are isomorphic induced connected subgraphs of G of order k but there is no automorphism of G mapping one onto the other.

It remains to show that $2|E| - 2 < k$. As the values of k, m, n imply this inequality whenever E is not a K_r^t, we need to show that if $G \cong X_{2,2}(K_r^t)$, then $2|K_r^t| - 2 = 2tr - 2 < k$. Let X be an $\mathcal{S}$-block with $x, x', y \in V(X)$ and $xx' \in E(G)$, such that x and x' belong to the same $\mathcal{S}$-separator and y belongs to the other $\mathcal{S}$-separator in X. In this setting $G[x, x', y]$ is a K^3, and thus the subgraphs $X - \{x, x'\}$ and $X - \{x, y\}$ are isomorphic. Suppose that $2tr = |X| \geq k+2$, then there is an induced subgraph X' of X of size precisely $k+2$ containing x, x' and y such that $X' - \{x, x'\}$ and $X' - \{x, y\}$ are isomorphic but there is no automorphism of G mapping one onto the other. This shows that the inequality $2|E| - 2 < k$ holds in all cases. □

By Lemma 6.13, 6.14, and 6.15 we may finish the first case.

Theorem 6.16. *Let $k \geq 3$, let G be a connected k-CS-transitive graph with at least two ends, and let $\mathcal{S}$ be a basic cut system of G such that every open block is empty. Then there are cardinals $\kappa, \lambda \geq 2$ and integers m, n such that G is isomorphic to one of the following graphs:*

(1) $X_{\kappa,\lambda}(K^1)$*;*

(2) $X_{2,\lambda}(K^n)$ *with* $n < \frac{k}{2} + 1$*;*

(3) $X_{\kappa,2}(\overline{K^m})$ *with* $m < \frac{k+2}{3}$*;*

(4) $X_{2,2}(E)$ *with* $E \in \mathcal{E}_{k,m,n}$, $m \leq k-2$, $n < \frac{k-|E|}{2} + 2$ *and* $2|E| - 2 < k$*;*

(5) Y_κ *(if k is odd).* □

6.2 Non-empty open blocks

Let us discuss the connected k-CS-transitive graphs with at least two ends for $k \geq 3$ such that every basic cut system has non-empty open blocks. As mentioned before this case restricts k to be even by Lemma 6.4. According to Lemma 6.9 every vertex not in any separator of a basic cut system lies in an open block.

Let us show that the k-CS-transitive graphs with non-empty open blocks resemble some $Z_{\kappa,\lambda}(H_1, H_2)$ by proving that the automorphism group acts transitively on its open blocks.

Lemma 6.17. *Let $k \geq 3$, let G be a connected k-CS-transitive graph with at least two ends, and let $\mathcal{S}$ be a basic cut system of G such that some open $\mathcal{S}$-block is not empty. Then every open $\mathcal{S}$-block is non-empty and the automorphism group of G acts transitively on the $\mathcal{S}$-blocks.*

Proof. Let X be an $\mathcal{S}$-block. By Lemma 6.1, X contains two distinct $\mathcal{S}$-separators and any two such separators are disjoint according to Corollary 6.12. By Lemma 6.11 it holds that X contains an edge sx where s lies in an $\mathcal{S}$-separator $S \subseteq X$ and x lies in $X - S$. As there is an induced path P of length $k-2$ starting at s with all its other vertices in a component of $G - S$ that avoids X, the neighbour of x on the induced path xP of length $k-1$ lies in an $\mathcal{S}$-separator, and thus x is not contained in any $\mathcal{S}$-separator according to Lemma 6.5.

Let Y be a further $\mathcal{S}$-block. By the previous argument a vertex $y \in \mathring{Y}$ exists. Let P' be an induced path of length $k-1$ starting at y—such a path exists as showed above. Since G is k-CS-transitive, there is an automorphism α mapping P' to xP. As k is even by Lemma 6.4 it holds that $y^\alpha = x$ according to Lemma 6.5. Thus $\mathring{Y}^\alpha \cap \mathring{X} \neq \emptyset$ and even $Y^\alpha = X$, as the intersection of any two distinct $\mathcal{S}$-blocks lies in an $\mathcal{S}$-separator. □

Lemma 6.18. *Let $k \geq 3$, let G be a connected k-CS-transitive graph with at least two ends, and let $\mathcal{S}$ be a basic cut system of G. If some open $\mathcal{S}$-block is not empty, then there are graphs H_1, H_2 and cardinals κ, λ such that G is isomorphic to $Z_{\kappa,\lambda}(H_1, H_2)$.*

Proof. The structure tree $\mathcal{T}$ of G and $\mathcal{S}$ is an infinite tree where vertices of even distance have the same degree, as $\mathcal{S}$ is basic and as the automorphisms of G act transitively on the $\mathcal{S}$-blocks by Lemma 6.17. Let κ be the degree of any $\mathcal{S}$-separator in $\mathcal{T}$ and let λ be the degree of any $\mathcal{S}$-block in $\mathcal{T}$. Let H_1 be isomorphic to $G[S]$ for some $\mathcal{S}$-separator S, and let H_2 be isomorphic to some open $\mathcal{S}$-block. Then again as $\mathcal{S}$ is basic and by Lemma 6.17 all separators induce an isomorphic copy of H_1 in G and all open blocks are isomorphic to H_2. Since, according to Lemma 6.11, every vertex of an open block $\mathring{X}$ is adjacent to all vertices in $\mathcal{S}$-separators that lie in X, it holds that $G \cong Z_{\kappa,\lambda}(H_1, H_2)$. □

As every connected k-CS-transitive graph for $k \geq 3$ with more than one end and some non-empty open block is isomorphic to $Z_{\kappa,\lambda}(H_1, H_2)$ for some graphs H_1 and H_2, it remains to specify the building blocks and possible values for κ and λ of these graphs. In Section 2.1 we describe what a basic cut system for these graphs looks like if H_1 and H_2 are finite.

Lemma 6.19. *Let $k \geq 3$, let $G \cong Z_{\kappa,\lambda}(H_1, H_2)$ be a k-CS-transitive graph with at least two ends, and let $\mathcal{S}$ be a basic cut system of G such that some open $\mathcal{S}$-block is not empty. Then the following holds:*

(i) At least one of κ or λ is 2;

(ii) if H_i contains two non-adjacent vertices, then H_j $(j \neq i)$ is complete and $\kappa = \lambda = 2$;

(iii) if H_i contains an edge, then H_j $(i \neq j)$ contains no edge.

Proof. Either $H_1 \not\cong H_2$ or $\kappa \neq \lambda$ since the copies of H_1 and H_2 are not $\mathrm{Aut}(G)$-isomorphic. Suppose both κ and λ are at least 3, then there are two distinct orbits of k-forks. One whose members poke in two distinct open $\mathcal{S}$-blocks, and one whose members poke in two distinct $\mathcal{S}$-separators. As a k-CS-transitive graph has only one orbit of k-forks this proves (i) by contradiction.

Part (ii) follows using an analogous argument: Suppose κ or λ is greater than 2. Then there is a k-fork that pokes in just one copy of an H_i and one that pokes in two distinct copies of H_1 (if $\kappa > 2$) or in two distinct copies of H_2 (if $\lambda > 2$). Suppose on the other hand that there are two non-adjacent vertices in H_j, then there are two incompatible k-forks, too. One pokes in an open $\mathcal{S}$-block and the other one in an $\mathcal{S}$-separator.

For (iii), suppose that H_i as well as H_j contain edges. Then there are k-spoons that poke in open $\mathcal{S}$-blocks and others that poke in $\mathcal{S}$-separators. □

From the previous lemma we immediately get the following corollary.

Corollary 6.20. *Let $k \geq 3$, let $G \cong Z_{\kappa,\lambda}(H_1, H_2)$ be a k-CS-transitive graph with at least two ends, and let $\mathcal{S}$ be a basic cut system of G such that some open $\mathcal{S}$-block is not empty. If both H_1 and H_2 have at least two vertices, then one is a complete graph, the other one is the complement of a complete graph, and $\kappa = \lambda = 2$.* □

To finish the proof in the situation that both, H_1 and H_2, have at least two vertices, we will restrict the order of these graphs.

Lemma 6.21. *Let $k \geq 3$, let $G \cong Z_{2,2}(H_1, H_2)$ be a k-CS-transitive graph with at least two ends, and let $\mathcal{S}$ be a basic cut system of G such that some open $\mathcal{S}$-block is not empty. If $H_1 \cong \overline{K^m}$ and $H_2 \cong K^n$, then $2m + n \leq k + 1$.*

Proof. Suppose that $2m + n > k + 1$. Let X be a building block corresponding to the complete graph H_2 and let Y, Y' be the two building blocks corresponding to H_1 adjacent to X. If $m \geq 2$, then as k is even there are subsets Y_1, Y_2 in Y and Y_1', Y_2' in Y' with

$$\begin{aligned}
|Y_1| &= \min\{m-1, \tfrac{k}{2} - 1\}, \\
|Y_1'| &= \min\{m-1, \tfrac{k}{2} - 1\}, \\
|Y_2| &= \min\{m-2, \tfrac{k}{2} - 2\}, \text{ and} \\
|Y_2'| &= \min\{m, \tfrac{k}{2}\}.
\end{aligned}$$

If $n \geq 2$ let X' be a subset of $V(X)$ of cardinality $k - (|Y_1| + |Y_1'|) \geq 2$ which exists as

$$k - (|Y_1| + |Y_1'|) \leq k - 2(m-1) \leq n.$$

The graphs $G[Y_1 \cup X' \cup Y_1']$ and $G[Y_2 \cup X' \cup Y_2']$ are isomorphic, so by k-CS-transitivity, there is an automorphism of G mapping the first onto the second subgraph. This is a contradiction, as Y_2' is larger than Y_1 as well as Y_1' and every automorphism of G has to map a building block onto a building block corresponding to the same H_i by the construction of $Z_{2,2}(H_1, H_2)$ and our choices for H_1 and H_2.

If $n = 1$, then $2m > k$ and hence $2m \geq k + 2$ as k is even. By enlarging each of Y_1' and Y_2' by one vertex we obtain a similar contradiction in this case as for $n \geq 2$.

If $m = 1$, then we have $n \geq k$. Let X be a subset of the vertex set of a building block corresponding to the complete graph H_2 of cardinality k, Let $x \in X$, and let y be a vertex adjacent to x but not in the same building block. By the construction of G we know that y is adjacent to every vertex of X. Thus, the subgraphs $G[X]$ and $G[X] - x + y$ are both complete graphs on k vertices, so there is an automorphism of G that maps the first onto the second by k-CS-transitivity. But again as every automorphism of G maps a building block onto a building block corresponding to the same H_i, we obtain a contradiction. □

Lemma 6.22. *Let $k \geq 3$, let $G \cong Z_{\kappa,\lambda}(H_1, H_2)$ be a k-CS-transitive graph with at least two ends, and let $\mathcal{S}$ be a basic cut system of G such that some open $\mathcal{S}$-block is not empty. If one of κ and λ is not 2, then both H_1 and H_2 are complete, one of order 1 and the other of order at most $k - 1$.*

Proof. It follows directly from Lemma 6.19 (ii) that both H_1 and H_2 are complete. By Lemma 6.19 (iii) we may assume that $|H_1| = 1$. Suppose that H_2 has more than $k - 1$ vertices. Every open $\mathcal{S}$-block $\mathring{X}$ is a building block that corresponds to H_2 and thus contains an isomorphic copy of a K^k. There is a second isomorphic copy of a K^k in G with $k - 1$ vertices in $\mathring{X}$ and one vertex in some $\mathcal{S}$-separator $S \subseteq X$. Since there is no automorphism of G mapping one onto the other, H_2 has at most $k - 1$ vertices. □

The last part in this case of the proof (that there is some non-empty open block) is to determine the graphs H_2 if the graph H_1 has only one vertex and the open blocks are neither complete nor complements of complete graphs.

Lemma 6.23. *Let $k \geq 3$, let $G \cong Z_{2,2}(H_1, H_2)$ be a k-CS-transitive graph with at least two ends, and let $\mathcal{S}$ be a basic cut system of G such that some open $\mathcal{S}$-block is not empty. If H_2 is neither complete nor the complement of a complete graph, then $H_1 \cong K^1$ and $H_2 \in \mathcal{E}_{k,m,n}$ for $m \leq k-2$ and $n \leq \frac{k}{2}+1$.*

Proof. As H_2 is not complete, it contains two non-adjacent vertices. This implies by Lemma 6.19 (ii) that H_1 is complete. Since H_2 also contains an edge, H_1 does not and thus is isomorphic to K^1. By Corollary 4.2 it suffices to show that (a) the maximum degree of H_2 is at most $k-2$, (b) H_2 is l-S-transitive for all $l \leq k-1$, (c) any induced subgraph of order at least $\frac{k}{2}+1$ in H_2 is connected, and (d) no two non-adjacent vertices of H_2 have $k-2$ common neighbours.

The proofs of (a), (b) and (d) are analogous to those of Lemma 6.15 (a), (b) and (d).

(c) Following the argument of Lemma 6.15 (c), an induced subgraph of order at least $\frac{k-1}{2}+1$ in H_2 is connected. The '-1' in that term corresponds to the '$-|E|$' in Lemma 6.15. Since k is even, every induced subgraph of order at least $\frac{k}{2}+1$ is connected if and only if every induced subgraph of order at least $\frac{k-1}{2}+1$ is connected. □

These lemmas let us finish the case of non-empty open blocks.

Theorem 6.24. *Let $k \geq 3$, let G be a connected k-CS-transitive graph with at least two ends, and let $\mathcal{S}$ be a basic cut system of G such that some open $\mathcal{S}$-block is not empty. Then k is even and G is isomorphic to one of the following graphs:*

(6) $Z_{2,2}(\overline{K^m}, K^n)$ with $2m+n \leq k+1$;

(7) $Z_{\kappa,\lambda}(K^1, K^n)$ with $n \leq k-1$ and cardinals κ, λ with $\kappa = 2$ or $\lambda = 2$;

(8) $Z_{2,2}(K^1, E)$ with $E \in \mathcal{E}_{k,m,n}$, $m \leq k-2$ and $n \leq \frac{k}{2}+1$. □

With Corollary 5.4 and Corollary 6.10 the Theorems 6.16 and 6.24 imply our second main result, Theorem 1.3.

7 k-CS-homogeneous graphs

In this section we shall prove Corollary 1.4. The first part of the proof will be to exclude those k-CS-transitive graphs that do not occur in the list of Corollary 1.4 and then to prove that the remaining graphs are k-CS-homogeneous.

Proof of Corollary 1.4. As every k-CS-homogeneous graph is k-CS-transitive, the connected k-CS-homogeneous graphs with $k \geq 3$ and at least two ends belong to classes (1) to (8) of Theorem 1.3. Let us first show that for the appropriate k all graphs that occur in the list Theorem 1.3 but not in that of Corollary 1.4 are not k-CS-homogenous.

For odd $k \geq 3$, the graphs Y_κ for $\kappa \geq 3$ are not k-CS-homogeneous, since we cannot map an induced path in Y_κ of length $k-1$ onto itself by an automorphism of Y_κ without being the identity on that path, as its outermost edges lie in buildings blocks of distinct kinds. As the automorphism group of any non-trivial path consists of two elements, these graphs are not k-CS-homogenous.

For even $k \geq 3$, any graph $G \cong Z_{\kappa,\lambda}(H_1, H_2)$ for any distinct graphs H_1, H_2 or distinct cardinals κ, λ is not k-CS-homogeneous as we cannot map an induced path of odd length $k-1$ in G onto itself by an automorphism of the whole graph without being the identity on that path as its endvertices lie in distinct kinds of building blocks of G. If on the other hand $H_1 \cong H_2$ and $\kappa = \lambda$, then as it has to be k-CS-transitive Theorem 1.3 implies that $H_1 \cong K^1$ and $Z_{\kappa,\lambda}(H_1, H_2) \cong X_{2,\kappa}(K^1)$ and hence G belongs to the graphs in class (1) of Theorem 1.3.

Let us consider a graph $G \cong X_{2,\lambda}(K^n)$ with arbitrary $k \geq 3$. If $\frac{k}{2} \leq n < \frac{k}{2} + 1$, then there is an induced subgraph isomorphic to K^k in two (completely) adjacent building blocks of G. We cannot extend any automorphism of such a subgraph that does not respect the building blocks to an automorphism of the whole graph. This implies $n < \frac{k}{2}$ in this case.

Let $G \cong X_{\kappa,2}(\overline{K^m})$. If $\frac{k}{3} \leq m < \frac{k+2}{3}$, then take an arbitrary subgraph X on k vertices of three building blocks one of which is completely adjacent to the other two that are not adjacent to each other. Then there is at most one vertex of the three building blocks missing in X. Thus we might build an automorphism of X that maps two vertices of the two non-adjacent building blocks onto each other and fixes all the other vertices of X. As $m \geq 2$, this automorphism of X cannot be extended to an automorphism of G.

Let us now assume that $G \cong X_{2,2}(E)$ for an $E \in \mathcal{E}_{k,m,n}$ with $m \leq k-2$, $2|E| - 2 < k$ and $n < \frac{k-|E|}{2} + 2$. Suppose that E contains an induced subgraph on at least $\frac{k-|E|}{2}$ vertices that is not connected. Let E_1, E_2, E_3 be three building blocks of G such that E_2 is (completely) adjacent to the other two but E_1 and E_3 are not adjacent. Then there are two induced subgraphs $X \subseteq E_1$ and $Y \subseteq E_3$ each of order at least $\frac{k-|E|}{2}$, both not connected such that $G[X] \cong G[Y]$. By the cardinality of these vertex sets, there is a non-empty vertex set Z in E_2 such that $|X| + |Y| + |Z| = k$. There is an automorphism of $H := G[X \cup Y \cup Z]$ that exchanges a component of $G[X]$ with one of $G[Y]$ and fixes every other vertex in H. As every automorphism of G maps vertices in the same building block again to a common building block, the just described automorphism of H does not extend to an automorphism of G. By contradiction we get that $E \in \mathcal{E}_{k,m,n'}$ for k and m as above and $n' < \frac{k-|E|}{2} + 1$.

It remains to show that $2|E| < k$ in this situation. If E is isomorphic to t disjoint K^r, then the inequalities imply $tr < \frac{k-|E|}{2}$ and hence $3|E| = 3tr < k$. If $E \cong C_5$, then $4 < \frac{k-5}{2} + 1$ implies $11 < k$ and if $E \cong L(K_{3,3})$, then $6 < \frac{k-9}{2} + 1$ implies $19 < k$. Suppose $2|E| \geq k$, then none of these three previous cases may occur and we conclude $E \cong K_r^t$ with $2 \leq t$. Let H be a subgraph of G induced by two adjacent building blocks B_1, B_2. Then H has less than $k+2$ vertices. Let X be an induced subgraph of H on k vertices. Then either $X = H$ or there is one vertex x in H with $X = H - x$. There is a set Y_1 of r independent vertices in $X \cap B_1$ and a set Y_2 of r independent vertices in $X \cap B_2$. As each of these $2r$ vertices is adjacent to all other vertices of X, there is an automorphism of X that maps Y_1 onto Y_2 and vice versa and that fixes every other vertex in X. Such an automorphism of X cannot extend to an automorphism of G as vertices in the same building block have to be mapped into a common building block by every automorphism of G and this is not satisfied by the above described automorphism of X.

It remains to show that the graphs described in Corollary 1.4 are k-CS-homogeneous. In principle, the proof is similar to those in Section 5. Therefore, we just point out the important bits that have to be changed and give a sketch of the remaining part of the proof. Let G be a graph that occurs in the list of Corollary 1.4. For the corresponding assertion of Lemma 5.1, it suffices to see that the only graphs in the list of Corollary 1.4 that have a complete graph on k vertices as subgraph, are the graphs $X_{\kappa,\lambda}(K^1)$ and $X_{\kappa,2}(\overline{K^m})$ and in each of these cases the construction of the graphs admits the extension of every isomorphism between two complete subgraphs on k vertices.

For the proof of Lemma 5.2, remark that any induced subgraph of an induced connected subgraph X on k vertices with diameter 2 has to meet at least three building blocks by cardinality reasons. It easily follows in each case that either there exists a unique smallest separator in X which is precisely $X \cap B$ where B is a building block adjacent to all other building blocks of G that meet X, or $G \cong X_{\kappa,2}(\overline{K^m})$ and X is a subgraph of a complete multipartite graph with partition classes each of the same cardinality. Where the required extension of any isomorphism between two induced connected subgraphs follows from the homogeneity of complete multipartite graphs in the last case, the extension exists for the first case because the building blocks are homogeneous and by the construction of the graphs $X_{\kappa,\lambda}(H)$.

For the situation that the induced isomorphic subgraphs on k vertices are connected and have diameter at least 3, it suffices to see that any isomorphism between any paths of length at least 3 whose diameter in G is at least 3 in these graphs can be extended to an automorphism of the whole graph. The further construction of the automorphisms α_k in the proof of Lemma 5.3 can also be chosen so that they extend the given isomorphism between the two induced connected subgraphs of order k. This completes the sketch of this direction of the proof and hence the whole proof. □

8 Ends of k-CS-transitive graphs

Gray [6] asked whether every locally finite k-CS-transitive graph is end-transitive for $k \geq 3$. With Theorem 1.3 we may answer his question.

Theorem 8.1. *Let $k \geq 3$ and let G be a connected locally finite graph. If G is k-CS-transitive, then it is end-transitive.* □

This theorem does not extend to graphs with vertices of infinite degree. For example the graphs $X_{\kappa,\lambda}$ with $\kappa \geq \aleph_0, \lambda \geq 2$ contain fundamentally different ends. Let us make this precise: a ray is *local* if it meets a set of finite diameter infinitely often. An end is *local* if all its rays are local, and an end is *global* if none of its rays is local. Theorem 1.3 shows that in k-CS-transitive graphs with $k \geq 3$ and more than one end every end is either local or global and that the automorphism group acts transitively on those of each kind.

Theorem 8.2. *Let $k \geq 3$ and G be a connected k-CS-transitive graph with more than one end. Then every end of G is either local or global. The automorphism group of G acts transitively on the local ends, as well as on the global ends.*

Furthermore, G is end-transitive if and only if it has no local end. □

Krön and Möller [9, 10] introduced metric ends. They call rays *metric* if they are not local, that is, if any infinite subset of its vertices does not have finite diameter in G. Two metric rays R_1 and R_2 are *metrically equivalent* if there is no vertex set S of finite diameter such that R_1 and R_2 lie eventually in different components of $G - S$. This is an equivalence relation on metric rays, whose classes are the *metric ends* of the graph. In locally finite graphs the notions of being an end and being a metric end coincide. Thus for connected locally finite k-CS-transitive graphs with $k \geq 3$ and with more than one end its automorphism group acts transitively on its metric ends. In spite of the local ends this extends by inspection of the examples in Theorem 1.3 to graphs that are not necessarily locally finite.

Theorem 8.3. *If $k \geq 3$, then the automorphism group of any connected k-CS-transitive graph with more than one end acts transitively on the metric ends of the graph.* □

Acknowledgement

We gratefully acknowledge the careful reading and helpful comments of a referee.

References

[1] W. Dicks, M.J. Dunwoody, Groups Acting on Graphs, Cambridge Stud. Adv. Math., vol. 17, Cambridge Univ. Press, Cambridge, 1989

[2] R. Diestel, *Graph Theory*, 3rd ed. Springer-Verlag, 2005

[3] M.J. Dunwoody, B. Krön, Vertex cuts, arXiv:0905.0064v3

[4] H. Enomoto, Combinatorially homogeneous graphs, J. Combin. Theory Ser. B 30 (2) (1981), 215-223

[5] A. Gardiner, Homogeneous graphs, J. Combin. Theory Ser. B 20 (1) (1976), 94-102

[6] R. Gray, k-CS-transitive infinite graphs, J. Combin. Theory Ser. B 99 (2) (2009), 378-398

[7] R. Gray, D. Macpherson, Countable connected-homogeneous graphs, J. Combin. Theory Ser. B 100 (2) (2010), 97-118

[8] A.A. Ivanov, Bounding the diameter of a distance-regular graph, Dokl. Akad. Nauk SSSR 271 (4) (1983) 789-792

[9] B. Krön, End compactifications in non-locally-finite graphs, Math. Proc. Camb. Phil. Soc. 131 (3) (2001), 427-443

[10] B. Krön, R.G. Möller, Metric ends, fibers and automorphisms of graphs, Math. Nachr. 281 (1) (2008), 62-74

[11] H.D. Macpherson, Infinite distance transitive graphs of finite valency, Combinatorica 2 (1) (1982), 63-69

[12] R.G. Möller, Accessibility and ends of graphs, J. Combin. Theory Ser. B 66 (2) (1996), 303-309

[13] R.G. Möller, Distance-transitivity in infinite graphs, J. Combin. Theory Ser. B 60 (1) (1994), 36-39

[14] R.G. Möller, Groups acting on locally finite graphs—a survey of the infinitely ended case. In Groups '93 Galway/St. Andrews, Vol. 2, 426–456. Cambridge Univ. Press, Cambridge, 1995

[15] C. Ronse, On homogenous graphs, J. London Math. Soc. (2) 17 (3) (1978), 375-379

[16] C. Thomassen, W. Woess, Vertex-transitive graphs and accessibility, J. Combin. Theory Ser. B 58 (2) (1993), 248-268

Edge-disjoint double rays in infinite graphs: a Halin type result*

Nathan Bowler Johannes Carmesin Julian Pott

Abstract

We show that any graph that contains k edge-disjoint double rays for any $k \in \mathbb{N}$ contains also infinitely many edge-disjoint double rays. This was conjectured by Andreae in 1981.

1 Introduction

We say a graph G has *arbitrarily many vertex-disjoint* H if for every $k \in \mathbb{N}$ there is a family of k vertex-disjoint subgraphs of G each of which is isomorphic to H. Halin's Theorem says that every graph that has arbitrarily many vertex-disjoint rays, also has infinitely many vertex-disjoint rays [5]. In 1970 he extended this result to vertex-disjoint double rays [6]. Jung proved a strengthening of Halin's Theorem where the initial vertices of the rays are constrained to a certain vertex set [7].

We look at the same questions with 'edge-disjoint' replacing 'vertex-disjoint'. Consider first the statement corresponding to Halin's Theorem. It suffices to prove this statement in locally finite graphs, as each graph with arbitrarily many edge-disjoint rays contains a locally finite union of tails of these rays. But the statement for locally finite graphs follows from Halin's original Theorem applied to the line-graph.

This reduction to locally finite graphs does not work for Jung's Theorem or for Halin's statement about double rays. Andreae proved an analog of Jung's Theorem for edge-disjoint rays in 1981, and conjectured that a Halin-type Theorem would be true for edge-disjoint double rays [1]. Our aim in the current paper is to prove this conjecture.

More precisely, we say a graph G has *arbitrarily many edge-disjoint* H if for every $k \in \mathbb{N}$ there is a family of k edge-disjoint subgraphs of G each of which is isomorphic to H, and our main result is the following.

Theorem 1. *Any graph that has arbitrarily many edge-disjoint double rays has infinitely many edge-disjoint double rays.*

*J. Combin. Theory (Series B) Available online 31 October 2014

Even for locally finite graphs this theorem does not follow from Halin's analogous result for vertex-disjoint double rays applied to the line graph. For example a double ray in the line graph may correspond, in the original graph, to a configuration as in Figure 1.

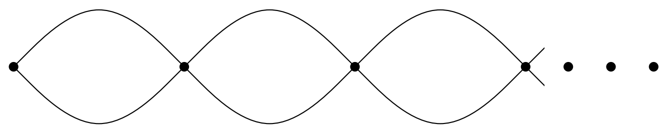

Figure 1: A graph that does not include a double ray but whose line graph does.

A related notion is that of ubiquity. A graph H is *ubiquitous* with respect to a graph relation $\leq$ if $nH \leq G$ for all $n \in \mathbb{N}$ implies $\aleph_0 H \leq G$, where nH denotes the disjoint union of n copies of H. For example, Halin's Theorem says that rays are ubiquitous with respect to the subgraph relation. It is known that not every graph is ubiquitous with respect to the minor relation [2], nor is every locally finite graph ubiquitous with respect to the subgraph relation [8, 9], or even the topological minor relation [2, 3]. However, Andreae has conjectured that every locally finite graph is ubiquitous with respect to the minor relation [2]. For more details see [3]. In Section 6 (the outlook) we introduce a notion closely related to ubiquity.

The proof is organised as follows. In Section 3 we explain how to deal with the cases that the graph has infinitely many ends, or an end with infinite vertex-degree. In Section 4 we consider the 'two ended' case: That in which there are two ends ω and ω' both of finite vertex-degree, and arbitrarily many edge-disjoint double rays from ω to ω'.

The only remaining case is the 'one ended' case: That in which there is a single end ω of finite vertex-degree and arbitrarily many edge-disjoint double rays from ω to ω. One central idea in the proof of this case is to consider 2-rays instead of double rays. Here a 2-ray is a pair of vertex-disjoint rays. For example, from each double ray one obtains a 2-ray by removing a finite path. The remainder of the proof is subdivided into two parts: In Subsection 5.3 we show that if there are arbitrarily many edge-disjoint 2-rays into ω, then there are infinitely many such 2-rays. In Subsection 5.2 we show that if there are infinitely many edge-disjoint 2-rays into ω, then there are infinitely many edge-disjoint double rays from ω to ω.

We finish by discussing the outlook and mentioning some open problems.

2 Preliminaries

All our basic notation for graphs is taken from [4]. In particular, two rays in a graph are equivalent if no finite set separates them. The equivalence classes of this relation are called the *ends* of G. We say that a ray in an end ω *converges* to ω. A double ray *converges* to all the ends of which it includes a ray.

2.1 The structure of a thin end

It follows from Halin's Theorem that if there are arbitrarily many vertex-disjoint rays in an end of G, then there are infinitely many such rays. This fact motivated the central definition of the *vertex-degree* of an end ω: the maximal cardinality of a set of vertex-disjoint rays in ω.

An end is *thin* if its vertex-degree is finite, and otherwise it is *thick*. A pair (A, B) of edge-disjoint subgraphs of G is a *separation* of G if $A \cup B = G$. The number of vertices of $A \cap B$ is called the *order* of the separation.

Definition 2. Let G be a locally finite graph and ω a thin end of G. A countable infinite sequence $((A_i, B_i))_{i \in \mathbb{N}}$ of separations of G *captures* ω if for all $i \in \mathbb{N}$

- $A_i \cap B_{i+1} = \emptyset$,
- $A_{i+1} \cap B_i$ is connected,
- $\bigcup_{i \in \mathbb{N}} A_i = G$,
- the order of (A_i, B_i) is the vertex-degree of ω, and
- each B_i contains a ray from ω.

Lemma 3. *Let G be a locally finite graph with a thin end ω. Then there is a sequence that captures ω.*

Proof. Without loss of generality G is connected, and so is countable. Let $v_1, v_2, \ldots$ be an enumeration of the vertices of G. Let k be the vertex-degree of ω. Let $\mathcal{R} = \{R_1, \ldots, R_k\}$ be a set of vertex-disjoint rays in ω and let S be the set of their start vertices. We pick a sequence $((A_i, B_i))_{i \in \mathbb{N}}$ of separations and a sequence (T_i) of connected subgraphs recursively as follows. We pick (A_i, B_i) such that S is included in A_i, such that there is a ray from ω included in B_i, and such that B_i does not meet $\bigcup_{j<i} T_j$ or $\{v_j \mid j \leq i\}$: subject to this we minimise the size of the set X_i of vertices in $A_i \cap B_i$. Because of this minimization B_i is connected and X_i is finite. We take T_i to be a finite connected subgraph of B_i including X_i. Note that any ray that meets all of the B_i must be in ω.

By Menger's Theorem [4] we get for each $i \in \mathbb{N}$ a set $\mathcal{P}_i$ of vertex-disjoint paths from X_i to X_{i+1} of size $|X_i|$. From these, for each i we get a set of $|X_i|$ vertex-disjoint rays in ω. Thus the size of X_i is at most k. On the other hand it is at least k as each ray R_j meets each set X_i.

Assume for contradiction that there is a vertex $v \in A_i \cap B_{i+1}$. Let R be a ray from v to ω inside B_{i+1}. Then R must meet X_i, contradicting the definition of B_{i+1}. Thus $A_i \cap B_{i+1}$ is empty.

Observe that $\bigcup \mathcal{P}_i \cup T_i$ is a connected subgraph of $A_{i+1} \cap B_i$ containing all vertices of X_i and X_{i+1}. For any vertex $v \in A_{i+1} \cap B_i$ there is a v–X_{i+1} path P in B_i. P meets B_{i+1} only in X_{i+1}. So P is included in $A_{i+1} \cap B_i$. Thus $A_{i+1} \cap B_i$ is connected. The remaining conditions are clear. □

Remark 4. *Every infinite subsequence of a sequence capturing ω also captures ω.* □

The following is obvious:

Remark 5. *Let G be a graph and $v, w \in V(G)$ If G contains arbitrarily many edge-disjoint v–w paths, then it contains infinitely many edge-disjoint v–w paths.* □

We will need the following special case of the theorem of Andreae mentioned in the Introduction.

Theorem 6 (Andreae [1])**.** *Let G be a graph and $v \in V(G)$. If there are arbitrarily many edge-disjoint rays all starting at v, then there are infinitely many edge-disjoint rays all starting at v.*

3 Known cases

Many special cases of Theorem 1 are already known or easy to prove. For example Halin showed the following.

Theorem 7 (Halin)**.** *Let G be a graph and ω an end of G. If ω contains arbitrarily many vertex-disjoint rays, then G has a half-grid as a minor.*

Corollary 8. *Any graph with an end of infinite vertex-degree has infinitely many edge-disjoint double rays.* □

Another simple case is the case where the graph has infinitely many ends.

Lemma 9. *A tree with infinitely many ends contains infinitely many edge-disjoint double rays.*

Proof. It suffices to show that every tree T with infinitely many ends contains a double ray such that removing its edges leaves a component containing infinitely many ends, since then one can pick those double rays recursively.

There is a vertex $v \in V(T)$ such that $T - v$ has at least 3 components C_1, C_2, C_3 that each have at least one end, as T contains more than 2 ends. Let e_i be the edge vw_i with $w_i \in C_i$ for $i \in \{1, 2, 3\}$. The graph $T \setminus \{e_1, e_2, e_3\}$ has precisely 4 components (C_1, C_2, C_3 and the one containing v), one of which, D say, has infinitely many ends. By symmetry we may assume that D is neither C_1 nor C_2. There is a double ray R all whose edges are contained in $C_1 \cup C_2 \cup$

$\{e_1, e_2\}$. Removing the edges of R leaves the component D, which has infinitely many ends. □

Corollary 10. *Any connected graph with infinitely many ends has infinitely many edge-disjoint double rays.* □

4 The 'two ended' case

Using the results of Section 3 it is enough to show that any graph with only finitely many ends, each of which is thin, has infinitely many edge-disjoint double rays as soon as it has arbitrarily many edge-disjoint double rays. Any double ray in such a graph has to join a pair of ends (not necessarily distinct), and there are only finitely many such pairs. So if there are arbitrarily many edge-disjoint double rays, then there is a pair of ends such that there are arbitrarily many edge-disjoint double rays joining those two ends. In this section we deal with the case where these two ends are different, and in Section 5 we deal with the case that they are the same. We start with two preparatory lemmas.

Lemma 11. *Let G be a graph with a thin end ω, and let $\mathcal{R} \subseteq \omega$ be an infinite set. Then there is an infinite subset of $\mathcal{R}$ such that any two of its members intersect in infinitely many vertices.*

Proof. We define an auxilliary graph H with $V(H) = \mathcal{R}$ and an edge between two rays if and only if they intersect in infinitely many vertices. By Ramsey's Theorem either H contains an infinite clique or an infinite independent set of vertices. Let us show that there cannot be an infinite independent set in H. Let k be the vertex-degree of ω: we shall show that H does not have an independent set of size $k+1$. Suppose for a contradiction that $X \subseteq \mathcal{R}$ is a set of $k+1$ rays that is independent in H. Since any two rays in X meet in only finitely many vertices, each ray in X contains a tail that is disjoint to all the other rays in X. The set of these $k+1$ vertex-disjoint tails witnesses that ω has vertex-degree at least $k+1$, a contradiction. Thus there is an infinite clique $K \subseteq H$, which is the desired infinite subset. □

Lemma 12. *Let G be a graph consisting of the union of a set $\mathcal{R}$ of infinitely many edge-disjoint rays of which any pair intersect in infinitely many vertices. Let $X \subseteq V(G)$ be an infinite set of vertices, then there are infinitely many edge-disjoint rays in G all starting in different vertices of X.*

Proof. If there are infinitely many rays in $\mathcal{R}$ each of which contains a different vertex from X, then suitable tails of these rays give the desired rays. Otherwise there is a ray $R \in \mathcal{R}$ meeting X infinitely often. In this case, we choose the desired rays recursively such that each contains a tail from some ray in $\mathcal{R} - R$. Having chosen finitely many such rays, we can always pick another: we start at some point in X on R which is beyond all the (finitely many) edges on R used so far. We follow R until we reach a vertex of some ray R' in $\mathcal{R} - R$ whose tail has not been used yet, then we follow R'. □

Lemma 13. *Let G be a graph with only finitely many ends, all of which are thin. Let ω_1, ω_2 be distinct ends of G. If G contains arbitrarily many edge-disjoint double rays each of which converges to both ω_1 and ω_2, then G contains infinitely many edge-disjoint double rays each of which converges to both ω_1 and ω_2.*

Proof. For each pair of ends, there is a finite set separating them. The finite union of these finite sets is a finite set $S \subseteq V(G)$ separating any two ends of G. For $i = 1, 2$ let C_i be the component of $G - S$ containing ω_i.

There are arbitrarily many edge-disjoint double rays from ω_1 to ω_2 that have a common last vertex v_1 in S before staying in C_1 and also a common last vertex v_2 in S before staying in C_2. Note that v_1 may be equal to v_2. There are arbitrarily many edge-disjoint rays in $C_1 + v_1$ all starting in v_1. By Theorem 6 there is a countable infinite set $\mathcal{R}_1 = \{R_1^i \mid i \in \mathbb{N}\}$ of edge-disjoint rays each included in $C_1 + v_1$ and starting in v_1. By replacing $\mathcal{R}_1$ with an infinite subset of itself, if necessary, we may assume by Lemma 11 that any two members of $\mathcal{R}_1$ intersect in infinitely many vertices. Similarly, there is a countable infinite set $\mathcal{R}_2 = \{R_2^i \mid i \in \mathbb{N}\}$ of edge-disjoint rays each included in $C_2 + v_2$ and starting in v_2 such that any two members of $\mathcal{R}_2$ intersect in infinitely many vertices.

Let us subdivide all edges in $\bigcup \mathcal{R}_1$ and call the set of subdivision vertices X_1. Similarly, we subdivide all edges in $\bigcup \mathcal{R}_2$ and call the set of subdivision vertices X_2. Below we shall find double rays in the subdivided graph, which immediately give rise to the desired double rays in G.

Suppose for a contradiction that there is a finite set F of edges separating X_1 from X_2. Then v_i has to be on the same side of that separation as X_i as there are infinitely many v_i–X_i edges. So F separates v_1 from v_2, which contradicts the fact that there are arbitrarily many edge-disjoint double rays containing both v_1 and v_2. By Remark 5 there is a set $\mathcal{P}$ of infinitely many edge-disjoint X_1–X_2 paths. As all vertices in X_1 and X_2 have degree 2, and by taking an infinite subset if necessary, we may assume that each end-vertex of a path in $\mathcal{P}$ lies on no other path in $\mathcal{P}$.

By Lemma 12 there is an infinite set Y_1 of start-vertices of paths in $\mathcal{P}$ together with an infinite set $\mathcal{R}'_1$ of edge-disjoint rays with distinct start-vertices whose set of start-vertices is precisely Y_1. Moreover, we can ensure that each ray in $\mathcal{R}'_1$ is included in $\bigcup \mathcal{R}_1$. Let Y_2 be the set of end-vertices in X_2 of those paths in $\mathcal{P}$ that start in Y_1. Applying Lemma 12 again, we obtain an infinite set $Z_2 \subseteq Y_2$ together with an infinite set $\mathcal{R}'_2$ of edge-disjoint rays included in $\bigcup \mathcal{R}_2$ with distinct start-vertices whose set of start-vertices is precisely Z_2.

For each path P in $\mathcal{P}$ ending in Z_2, there is a double ray in the union of P and the two rays from $\mathcal{R}'_1$ and $\mathcal{R}'_2$ that P meets in its end-vertices. By construction, all these infinitely many double rays are edge-disjoint. Each of those double rays converges to both ω_1 and ω_2, since each ω_i is the only end in C_i. □

Remark 14. *Instead of subdividing edges we also could have worked in the line graph of G. Indeed, there are infinitely many vertex-disjoint paths in the line* *graph from* $\bigcup \mathcal{R}_1$ *to* $\bigcup \mathcal{R}_2$.

5 The 'one ended' case

We are now going to look at graphs G that contain a thin end ω such that there are arbitrarily many edge-disjoint double rays converging only to the end ω. The aim of this section is to prove the following lemma, and to deduce Theorem 1.

Lemma 15. *Let G be a countable graph and let ω be a thin end of G. Assume there are arbitrarily many edge-disjoint double rays all of whose rays converge to ω. Then G has infinitely many edge-disjoint double rays.*

We promise that the assumption of countability will not cause problems later.

5.1 Reduction to the locally finite case

A key notion for this section is that of a 2-ray. A *2-ray* is a pair of vertex-disjoint rays. For example, from each double ray one obtains a 2-ray by removing a finite path.

In order to deduce that G has infinitely many edge-disjoint double rays, we will only need that G has arbitrarily many edge-disjoint 2-rays. In this subsection, we illustrate one advantage of 2-rays, namely that we may reduce to the case where G is locally finite.

Lemma 16. *Let G be a countable graph with a thin end ω. Assume there is a countable infinite set $\mathcal{R}$ of rays all of which converge to ω.*

Then there is a locally finite subgraph H of G with a single end which is thin such that the graph H includes a tail of any $R \in \mathcal{R}$.

Proof. Let $(R_i \mid i \in \mathbb{N})$ be an enumeration of $\mathcal{R}$. Let $(v_i \mid i \in \mathbb{N})$ be an enumeration of the vertices of G. Let U_i be the unique component of $G \setminus \{v_1, \ldots, v_i\}$ including a tail of each ray in ω.

For $i \in \mathbb{N}$, we pick a tail R'_i of R_i in U_i. Let $H_1 = \bigcup_{i \in \mathbb{N}} R'_i$. Making use of H_1, we shall construct the desired subgraph H. Before that, we shall collect some properties of H_1.

As every vertex of G lies in only finitely many of the U_i, the graph H_1 is locally finite. Each ray in H_1 converges to ω in G since $H_1 \setminus U_i$ is finite for every $i \in \mathbb{N}$. Let Ψ be the set of ends of H_1. Since ω is thin, Ψ has to be finite: $\Psi = \{\omega_1, \ldots, \omega_n\}$. For each $i \leq n$, we pick a ray $S_i \subseteq H_1$ converging to ω_i.

Now we are in a position to construct H. For any $i > 1$, the rays S_1 and S_i are joined by an infinite set $\mathcal{P}_i$ of vertex-disjoint paths in G. We obtain H from H_1 by adding all paths in the sets $\mathcal{P}_i$. Since H_1 is locally finite, H is locally finite.

It remains to show that every ray R in H is equivalent to S_1. If R contains infinitely many edges from the $\mathcal{P}_i$, then there is a single $\mathcal{P}_i$ which R meets infinitely, and thus R is equivalent to S_1. Thus we may assume that a tail of R is a ray in H_1. So it converges to some $\omega_i \in \Psi$. Since S_i and S_1 are equivalent, R and S_1 are equivalent, which completes the proof. □

Corollary 17. *Let G be a countable graph with a thin end ω and arbitrarily many edge-disjoint 2-rays of which all the constituent rays converge to ω. Then there is a locally finite subgraph H of G with a single end, which is thin, such that H has arbitrarily many edge-disjoint 2-rays.*

Proof. By Lemma 16 there is a locally finite graph $H \subseteq G$ with a single end such that a tail of each of the constituent rays of the arbitrarily many 2-rays is included in H. □

5.2 Double rays versus 2-rays

A connected subgraph of a graph G including a vertex set $S \subseteq V(G)$ is a *connector* of S in G.

Lemma 18. *Let G be a connected graph and S a finite set of vertices of G. Let $\mathcal{H}$ be a set of edge-disjoint subgraphs H of G such that each connected component of H meets S. Then there is a finite connector T of S, such that at most $2|S|-2$ graphs from $\mathcal{H}$ contain edges of T.*

Proof. By replacing $\mathcal{H}$ with the set of connected components of graphs in $\mathcal{H}$, if necessary, we may assume that each member of $\mathcal{H}$ is connected. We construct graphs T_i recursively for $0 \leq i < |S|$ such that each T_i is finite and has at most $|S| - i$ components, at most $2i$ graphs from $\mathcal{H}$ contain edges of T_i, and each component of T_i meets S. Let $T_0 = (S, \emptyset)$ be the graph with vertex set S and no edges. Assume that T_i has been defined.

If T_i is connected let $T_{i+1} = T_i$. For a component C of T_i, let C' be the graph obtained from C by adding all graphs from $\mathcal{H}$ that meet C.

As G is connected, there is a path P (possibly trivial) in G joining two of these subgraphs C_1' and C_2' say. And by taking the length of P minimal, we may assume that P does not contain any edge from any $H \in \mathcal{H}$. Then we can extend P to a C_1–C_2 path Q by adding edges from at most two subgraphs from $\mathcal{H}$ — one included in C_1' and the other in C_2'. We obtain T_{i+1} from T_i by adding Q.

$T = T_{|S|-1}$ has at most one component and thus is connected. And at most $2|S| - 2$ many graphs from $\mathcal{H}$ contain edges of T. Thus T is as desired. □

Let d, d' be 2-rays. d is a *tail* of d' if each ray of d is a tail of a ray of d'. A set D' is a *tailor* of a set D of 2-rays if each element of D' is a tail of some element of D but no 2-ray in D includes more than one 2-ray in D'.

Lemma 19. *Let G be a locally finite graph with a single end ω, which is thin. Assume that G contains an infinite set $D = \{d_1, d_2, \dots\}$ of edge-disjoint 2-rays.*

Then G contains an infinite tailor D' of D and a sequence $((A_i, B_i))_{i\in\mathbb{N}}$ capturing ω (see Definition 2) such that there is a family of vertex-disjoint connectors T_i of $A_i \cap B_i$ contained in $A_{i+1} \cap B_i$, each of which is edge-disjoint from each member of D'.

Proof. Let k be the vertex-degree of ω. By Lemma 3 there is a sequence $((A'_i, B'_i))_{i\in\mathbb{N}}$ capturing ω. By replacing each 2-ray in D with a tail of itself if necessary, we may assume that for all $(r, s) \in D$ and $i \in \mathbb{N}$ either both r and s meet A'_i or none meets A'_i. By Lemma 18 there is a finite connector T'_i of $A'_i \cap B'_i$ in the connected graph B'_i which meets in an edge at most $2k-2$ of the 2-rays of D that have a vertex in A'_i.

Thus, there are at most $2k - 2$ 2-rays in D that meet all but finitely many of the T'_i in an edge. By throwing away these finitely many 2-rays in D we may assume that each 2-ray in D is edge-disjoint from infinitely many of the T'_i. So we can recursively build a sequence $N_1, N_2, \dots$ of infinite sets of natural numbers such that $N_i \supseteq N_{i+1}$, the first i elements of N_i are all contained in N_{i+1}, and d_i only meets finitely many of the T'_j with $j \in N_i$ in an edge. Then $N = \bigcap_{i\in\mathbb{N}} N_i$ is infinite and has the property that each d_i only meets finitely many of the T'_j with $j \in N$ in an edge. Thus there is an infinite tailor D' of D such that no 2-ray from D' meets any T'_j for $j \in N$ in an edge.

We recursively define a sequence $n_1, n_2, \dots$ of natural numbers by taking $n_i \in N$ sufficiently large that B'_{n_i} does not meet T'_{n_j} for any $j < i$. Taking $(A_i, B_i) = (A'_{n_i}, B'_{n_i})$ and $T_i = T'_{n_i}$ gives the desired sequences. □

Lemma 20. *If a locally finite graph G with a single end ω which is thin contains infinitely many edge-disjoint 2-rays, then G contains infinitely many edge-disjoint double rays.*

Proof. Applying Lemma 19 we get an infinite set D of edge-disjoint 2-rays, a sequence $((A_i, B_i))_{i\in\mathbb{N}}$ capturing ω, and connectors T_i of $A_i \cap B_i$ for each $i \in \mathbb{N}$ such that the T_i are vertex-disjoint from each other and edge-disjoint from all members of D.

We shall construct the desired set of infinitely many edge-disjoint double rays as a nested union of sets D_i. We construct the D_i recursively. Assume that a set D_i of i edge-disjoint double rays has been defined such that each of its members is included in the union of a single 2-ray from D and one connector T_j. Let $d_{i+1} \in D$ be a 2-ray distinct from the finitely many 2-rays used so far. Let C_{i+1} be one of the infinitely many connectors that is different from all the finitely many connectors used so far and that meets both rays of d_{i+1}. Clearly, $d_{i+1} \cup C_{i+1}$ includes a double ray R_{i+1}. Let $D_{i+1} = D_i \cup \{R_{i+1}\}$. The union $\bigcup_{i\in\mathbb{N}} D_i$ is an infinite set of edge-disjoint double rays as desired. □

5.3 Shapes and allowed shapes

Let G be a graph and (A, B) a separation of G. A *shape* for (A, B) is a word $v_1x_1v_2x_2 \dots x_{n-1}v_n$ with $v_i \in A \cap B$ and $x_i \in \{l, r\}$ such that no vertex appears twice. We call the v_i the *vertices* of the shape. Every ray R *induces* a shape

$\sigma = \sigma_R(A,B)$ on every separation (A,B) of finite order in the following way: Let $<_R$ be the *natural order* on $V(R)$ induced by the ray, where $v <_R w$ if w lies in the unique infinite component of $R - v$. The vertices of σ are those vertices of R that lie in $A \cap B$ and they appear in σ in the order given by $<_R$. For v_i, v_{i+1} the path $v_i R v_{i+1}$ has edges only in A or only in B but not in both. In the first case we put l between v_i and v_{i+1} and in the second case we put r between v_i and v_{i+1}.

Let $(A_1,B_1),(A_2,B_2)$ be separations with $A_1 \cap B_2 = \emptyset$ and thus also $A_1 \subseteq A_2$ and $B_2 \subseteq B_1$. Let σ_i be a nonempty shape for (A_i,B_i). The word $\tau = v_1x_1v_2 \dots x_{n-1}v_n$ is an *allowed shape linking* σ_1 *to* σ_2 with *vertices* $v_1 \dots v_n$ if the following holds.

- v is a vertex of τ if and only if it is a vertex of σ_1 or σ_2,
- if v appears before w in σ_i, then v appears before w in τ,
- v_1 is the initial vertex of σ_1 and v_n is the terminal vertex of σ_2,
- $x_i \in \{l, m, r\}$,
- the subword vlw appears in τ if and only if it appears in σ_1,
- the subword vrw appears in τ if and only if it appears in σ_2,
- $v_i \neq v_j$ for $i \neq j$.

Each ray R defines a word $\tau = \tau_R[(A_1,B_1),(A_2,B_2)] = v_1x_1v_2 \dots x_{n-1}v_n$ with vertices v_i and $x_i \in \{l,m,r\}$ as follows. The vertices of τ are those vertices of R that lie in $A_1 \cap B_1$ or $A_2 \cap B_2$ and they appear in τ in the order given by $<_R$. For v_i, v_{i+1} the path $v_i R v_{i+1}$ has edges either only in A_1, only in $A_2 \cap B_1$, or only in B_2. In the first case we set $x_i = l$ and τ contains the subword $v_i l v_{i+1}$. In the second case we set $x_i = m$ and τ contains the subword $v_i m v_{i+1}$. In the third case we set $x_i = r$ and τ contains the subword $v_i r v_{i+1}$.

For a ray R to induce an allowed shape $\tau_R[(A_1,B_1),(A_2,B_2)]$ we need at least that R starts in A_2. However, each ray in ω has a tail such that whenever it meets an A_i it also starts in that A_i. Let us call such rays *lefty*. A 2-ray is *lefty* if both its rays are.

Remark 21. *Let (A_1,B_1), and (A_2,B_2) be two separations of finite order with $A_1 \subseteq A_2$, and $B_2 \subseteq B_1$. For every lefty ray R meeting A_1, the word $\tau_R[(A_1,B_1),(A_2,B_2)]$ is an allowed shape linking $\sigma_R(A_1,B_1)$ and $\sigma_R(A_2,B_2)$.* □

From now on let us fix a locally finite graph G with a thin end ω of vertex-degree k. And let $((A_i,B_i))_{i\in\mathbb{N}}$ be a sequence capturing ω such that each member has order k.

A 2-*shape* for a separation (A,B) is a pair of shapes for (A,B). Every 2-ray induces a 2-shape coordinatewise in the obvious way. Similarly, an *allowed* 2-*shape* is a pair of allowed shapes.

Clearly, there is a global constant $c_1 \in \mathbb{N}$ depending only on k such that there are at most c_1 distinct 2-shapes for each separation (A_i, B_i). Similarly, there is a global constant $c_2 \in \mathbb{N}$ depending only on k such that for all $i, j \in \mathbb{N}$ there are at most c_2 distinct allowed 2-shapes linking a 2-shape for (A_i, B_i) with a 2-shape for (A_j, B_j).

For most of the remainder of this subsection we assume that for every $i \in \mathbb{N}$ there is a set D_i consisting of at least $c_1 \cdot c_2 \cdot i$ edge-disjoint 2-rays in G. Our aim will be to show that in these circumstances there must be infinitely many edge-disjoint 2-rays.

By taking a tailor if necessary, we may assume that every 2-ray in each D_i is lefty.

Lemma 22. *There is an infinite set $J \subseteq \mathbb{N}$ and, for each $i \in \mathbb{N}$, a tailor D'_i of D_i of cardinality $c_2 \cdot i$ such that for all $i \in \mathbb{N}$ and $j \in J$ all 2-rays in D'_i induce the same 2-shape $\sigma[i, j]$ on (A_j, B_j).*

Proof. We recursively build infinite sets $J_i \subseteq \mathbb{N}$ and tailors D'_i of D_i such that for all $k \leq i$ and $j \in J_i$ all 2-rays in D'_k induce the same 2-shape on (A_j, B_j). For all $i \geq 1$, we shall ensure that J_i is an infinite subset of J_{i-1} and that the $i - 1$ smallest members of J_i and J_{i-1} are the same. We shall take J to be the intersection of all the J_i.

Let $J_0 = \mathbb{N}$ and let D'_0 be the empty set. Now, for some $i \geq 1$, assume that sets J_k and D'_k have been defined for all $k < i$. By replacing 2-rays in D_i by their tails, if necessary, we may assume that each 2-ray in D_i avoids A_ℓ, where ℓ is the $(i-1)$st smallest value of J_{i-1}. As D_i contains $c_1 \cdot c_2 \cdot i$ many 2-rays, for each $j \in J_{i-1}$ there is a set $S_j \subseteq D_i$ of size at least $c_2 \cdot i$ such that each 2-ray in S_j induces the same 2-shape on (A_j, B_j). As there are only finitely many possible choices for S_j, there is an infinite subset J_i of J_{i-1} on which S_j is constant. For D'_i we pick this value of S_j. Since each $d \in D'_i$ induces the empty 2-shape on each (A_k, B_k) with $k \leq \ell$ we may assume that the first $i - 1$ elements of J_{i-1} are also included in J_i.

It is immediate that the set $J = \bigcap_{i \in \mathbb{N}} J_i$ and the D'_i have the desired property. □

Lemma 23. *There are two strictly increasing sequences $(n_i)_{i \in \mathbb{N}}$ and $(j_i)_{i \in \mathbb{N}}$ with $n_i \in \mathbb{N}$ and $j_i \in J$ for all $i \in \mathbb{N}$ such that $\sigma[n_i, j_i] = \sigma[n_{i+1}, j_i]$ and $\sigma[n_i, j_i]$ is not empty.*

Proof. Let H be the graph on $\mathbb{N}$ with an edge $vw \in E(H)$ if and only if there are infinitely many elements $j \in J$ such that $\sigma[v, j] = \sigma[w, j]$.

As there are at most c_1 distinct 2-shapes for any separator (A_i, B_i), there is no independent set of size $c_1 + 1$ in H and thus no infinite one. Thus, by Ramsey's theorem, there is an infinite clique in H. We may assume without loss of generality that H itself is a clique by moving to a subsequence of the D'_i if necessary. With this assumption we simply pick $n_i = i$.

Now we pick the j_i recursively. Assume that j_i has been chosen. As i and $i + 1$ are adjacent in H, there are infinitely many indicies $\ell \in \mathbb{N}$ such that

$\sigma[i,\ell] = \sigma[i+1,\ell]$. In particular, there is such an $\ell > j_i$ such that $\sigma[i+1,\ell]$ is not empty. We pick j_{i+1} to be one of those ℓ.

Clearly, $(j_i)_{i\in\mathbb{N}}$ is an increasing sequence and $\sigma[i,j_i] = \sigma[i+1,j_i]$ as well as $\sigma[i,j_i]$ is non-empty for all $i \in \mathbb{N}$, which completes the proof. □

By moving to a subsequence of (D'_i) and $((A_j, B_j))$, if necessary, we may assume by Lemma 22 and Lemma 23 that for all $i, j \in \mathbb{N}$ all $d \in D'_i$ induce the same 2-shape $\sigma[i,j]$ on (A_j, B_j), and that $\sigma[i,i] = \sigma[i+1,i]$, and that $\sigma[i,i]$ is non-empty.

Lemma 24. *For all $i \in \mathbb{N}$ there is $D''_i \subseteq D'_i$ such that $|D''_i| = i$, and all $d \in D''_i$ induce the same allowed 2-shape $\tau[i]$ that links $\sigma[i,i]$ and $\sigma[i,i+1]$.*

Proof. Note that it is in this proof that we need all the 2-rays in D''_i to be lefty as they need to induce an allowed 2-shape that links $\sigma[i,i]$ and $\sigma[i,i+1]$ as soon as they contain a vertex from A_i. As $|D'_i| \geq i \cdot c_2$ and as there are at most c_2 many distinct allowed 2-shapes that link $\sigma[i,i]$ and $\sigma[i,i+1]$ there is $D''_i \subseteq D'_i$ with $|D''_i| = i$ such that all $d \in D''_i$ induce the same allowed 2-shape. □

We enumerate the elements of D''_j as follows: $d^j_1, d^j_2, \ldots, d^j_j$. Let (s^j_i, t^j_i) be a representation of d^j_i. Let $S^j_i = s^j_i \cap A_{j+1} \cap B_j$, and let $\mathcal{S}_i = \bigcup_{j\geq i} S^j_i$. Similarly, let $T^j_i = t^j_i \cap A_{j+1} \cap B_j$, and let $\mathcal{T}_i = \bigcup_{j\geq i} T^j_i$.

Clearly, $\mathcal{S}_i$ and $\mathcal{T}_i$ are vertex-disjoint and any two graphs in $\bigcup_{i\in\mathbb{N}}\{\mathcal{S}_i, \mathcal{T}_i\}$ are edge-disjoint. We shall find a ray R_i in each of the $\mathcal{S}_i$ and a ray R'_i in each of the $\mathcal{T}_i$. The infinitely many pairs (R_i, R'_i) will then be edge-disjoint 2-rays, as desired.

Lemma 25. *Each vertex v of $\mathcal{S}_i$ has degree at most 2. If v has degree 1 it is contained in $A_i \cap B_i$.*

Proof. Clearly, each vertex v of $\mathcal{S}_i$ that does not lie in any separator $A_j \cap B_j$ has degree 2, as it is contained in precisely one S^j_i, and all the leaves of S^j_i lie in $A_j \cap B_j$ and $A_{j+1} \cap B_{j+1}$ as d^j_i is lefty. Indeed, in S^j_i it is an inner vertex of a path and thus has degree 2 in there. If v lies in $A_i \cap B_i$ it has degree at most 2, as it is only a vertex of S^j_i for one value of j, namely $j = i$.

Hence, we may assume that $v \in A_j \cap B_j$ for some $j > i$. Thus, $\sigma[j,j]$ contains v and $l : \sigma[j,j] : r$ contains precisely one of the four following subwords:

$$lvl, lvr, rvl, rvr$$

(Here we use the notation $p : q$ to denote the concatenation of the word p with the word q.) In the first case $\tau[j-1]$ contains mvm as a subword and $\tau[j]$ has no m adjacent to v. Then S^{j-1}_i contains precisely 2 edges adjacent to v and S^j_i has no such edge. The fourth case is the first one with l and r and j and $j-1$ interchanged.

In the second and third cases, each of $\tau[j-1]$ and $\tau[j]$ has precisely one m adjacent to v. So both S^{j-1}_i and S^j_i contain precisely 1 edge adjacent to v.

As v appears only as a vertex of S^ℓ_i for $\ell = j$ or $\ell = j-1$, the degree of v in $\mathcal{S}_i$ is 2. □

Lemma 26. *There are an odd number of vertices in $\mathcal{S}_i$ of degree* 1.

Proof. By Lemma 25 we have that each vertex of degree 1 lies in $A_i \cap B_i$. Let v be a vertex in $A_i \cap B_i$. Then, $\sigma[i,i]$ contains v and $l : \sigma[i,i] : r$ contains precisely one of the four following subwords:

$$lvl, lvr, rvl, rvr$$

In the first and fourth case v has even degree. It has degree 1 otherwise. As $l : \sigma[i,i] : r$ starts with l and ends with r, the word lvr appear precisely once more than the word rvl. Indeed, between two occurrences of lvr there must be one of rvl and vice versa. Thus, there are an odd number of vertices with degree 1 in $\mathcal{S}_i$. □

Lemma 27. *$\mathcal{S}_i$ includes a ray.*

Proof. By Lemma 25 every vertex of $\mathcal{S}_i$ has degree at most 2 and thus every component of $\mathcal{S}_i$ has at most two vertices of degree 1. By Lemma 26 $\mathcal{S}_i$ has a component C that contains an odd number of vertices with degree 1. Thus C has precisely one vertex of degree 1 and all its other vertices have degree 2, thus C is a ray. □

Corollary 28. *G contains infinitely many edge-disjoint* 2*-rays.*

Proof. By symmetry, Lemma 27 is also true with $\mathcal{T}_i$ in place of $\mathcal{S}_i$. Thus $\mathcal{S}_i \cup \mathcal{T}_i$ includes a 2-ray X_i. The X_i are edge-disjoint by construction. □

Recall that Lemma 15 states that a countable graph with a thin end ω and arbitrarily many edge-disjoint double rays all whose subrays converge to ω, also has infinitely many edge-disjoint double rays. We are now in a position to prove this lemma.

Proof of Lemma 15. By Lemma 20 it suffices to show that G contains a subgraph H with a single end which is thin such that H has infinitely many edge-disjoint 2-rays. By Corollary 17, G has a subgraph H with a single end which is thin such that H has arbitrarily many edge-disjoint 2-rays. But then by the argument above H contains infinitely many edge-disjoint 2-rays, as required. □

With these tools at hand, the remaining proof of Theorem 1 is easy. Let us collect the results proved so far to show that each graph with arbitrarily many edge-disjoint double rays also has infinitely many edge-disjoint double rays.

Proof of Theorem 1. Let G be a graph that has a set D_i of i edge-disjoint double rays for each $i \in \mathbb{N}$. Clearly, G has infinitely many edge-disjoint double rays if its subgraph $\bigcup_{i\in\mathbb{N}} D_i$ does, and thus we may assume without loss of generality that $G = \bigcup_{i\in\mathbb{N}} D_i$. In particular, G is countable.

By Corollary 10 we may assume that each connected component of G includes only finitely many ends. As each component includes a double ray we may assume that G has only finitely many components. Thus, there is one component containing arbitrarily many edge-disjoint double rays, and thus we may assume that G is connected.

By Corollary 8 we may assume that all ends of G are thin. Thus, as mentioned at the start of Section 4, there is a pair of ends (ω, ω') of G (not necessarily distinct) such that G contains arbitrarily many edge-disjoint double rays each of which converges precisely to ω and ω'. This completes the proof as, by Lemma 13 G has infinitely many edge-disjoint double rays if ω and ω' are distinct and by Lemma 15 G has infinitely many edge-disjoint double rays if $\omega = \omega'$. □

6 Outlook and open problems

We will say that a graph H is *edge-ubiquitous* if every graph having arbitrarily many edge-disjoint H also has infinitely many edge-disjoint H.

Thus Theorem 1 can be stated as follows: the double ray is edge-ubiquitous. Andreae's Theorem implies that the ray is edge-ubiquitous. And clearly, every finite graph is edge-ubiquitous.

We could ask which other graphs are edge-ubiquitous. It follows from our result that the 2-ray is edge-ubiquitous. Let G be a graph in which there are arbitrarily many edge-disjoint 2-rays. Let $v * G$ be the graph obtained from G by adding a vertex v adjacent to all vertices of G. Then $v * G$ has arbitrarily many edge-disjoint double rays, and thus infinitely many edge-disjoint double rays. Each of these double rays uses v at most once and thus includes a 2-ray of G.

The vertex-disjoint union of k rays is called a *k-ray*. The k-ray is edge-ubiquitous. This can be proved with an argument similar to that for Theorem 1: Let G be a graph with arbitrarily many edge-disjoint k-rays. The same argument as in Corollaries 10 and 8 shows that we may assume that G has only finitely many ends, each of which is thin. By removing a finite set of vertices if necessary we may assume that each component of G has at most one end, which is thin. Now we can find numbers k_C indexed by the components C of G and summing to k such that each component C has arbitrarily many edge-disjoint k_C-rays. Hence, we may assume that G has only a single end, which is thin. By Lemma 16 we may assume that G is locally finite.

In this case, we use an argument as in Subsection 5.3. It is necessary to use k-shapes instead of 2-shapes but other than that we can use the same combinatorial principle. If C_1 and C_2 are finite sets, a *(C_1, C_2)-shaping* is a pair (c_1, c_2) where c_1 is a partial colouring of $\mathbb{N}$ with colours from C_1 which is defined at all but finitely many numbers and c_2 is a colouring of $\mathbb{N}^{(2)}$ with colours from C_2 (in our argument above, C_1 would be the set of all k-shapes and C_2 would be the set of all allowed k-shapes for all pairs of k-shapes).

Lemma 29. *Let $D_1, D_2, \dots$ be a sequence of sets of (C_1, C_2)-shapings where D_i has size i. Then there are strictly increasing sequences $i_1, i_2, \dots$ and $j_1, j_2, \dots$ and subsets $S_n \subseteq D_{i_n}$ with $|S_n| \geq n$ such that*

- *for any $n \in \mathbb{N}$ all the values of $c_1(j_n)$ for the shapings $(c_1, c_2) \in S_{n-1} \cup S_n$ are equal (in particular, they are all defined).*

- *for any $n \in N$, all the values of $c_2(j_n, j_{n+1})$ for the shapings $(c_1, c_2) \in S_n$ are equal.*

Lemma 29 can be proved by the same method with which we constructed the sets D''_i from the sets D_i. The advantage of Lemma 29 is that it can not only be applied to 2-rays but also to more complicated graphs like k-rays.

A *talon* is a tree with a single vertex of degree 3 where all the other vertices have degree 2. An argument as in Subsection 5.2 can be used to deduce that talons are edge-ubiquitous from the fact that 3-rays are. However, we do not know whether the graph in Figure 2 is edge-ubiquitous.

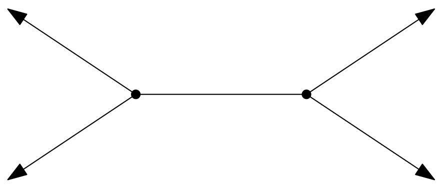

Figure 2: A graph obtained from 2 disjoint double rays, joined by a single edge. Is this graph edge-ubiquitous?

We finish with the following open problem.

Problem 30. *Is the directed analogue of Theorem 1 true? More precisely: Is it true that if a directed graph has arbitrarily many edge-disjoint directed double rays, then it has infinitely many edge-disjoint directed double rays?*

It should be noted that if true the directed analogue would be a common generalization of Theorem 1 and the fact that double rays are ubiquitous with respect to the subgraph relation.

7 Acknowledgement

We appreciate the helpful and accurate comments of a referee.

References

[1] Th. Andreae. Über maximale Systeme von kantendisjunkten unendlichen Wegen in Graphen. *Math. Nachr.*, 101:219–228, 1981.

[2] Thomas Andreae. On disjoint configurations in infinite graphs. *Journal of Graph Theory*, 39(4):222–229, 2002.

[3] Thomas Andreae. Classes of locally finite ubiquitous graphs. *Journal of Combinatorial Theory, Series B*, 103(2):274 – 290, 2013.

[4] R. Diestel. *Graph Theory.* Springer, 4th edition, 2010.

[5] R. Halin. Über unendliche Wege in Graphen. *Math. Annalen*, 157:125–137, 1964.

[6] R. Halin. Die Maximalzahl fremder zweiseitig unendlicher Wege in Graphen. *Math. Nachr.*, 44:119–127, 1970.

[7] H.A. Jung. Wurzelbäume und unendliche Wege in Graphen. *Math. Nachr.*, 41:1–22, 1969.

[8] John Lake. A problem concerning infinite graphs. *Discrete Mathematics*, 14(4):343 – 345, 1976.

[9] D.R Woodall. A note on a problem of Halin's. *Journal of Combinatorial Theory, Series B*, 21(2):132 – 134, 1976.

Linkages in Large Graphs of Bounded Tree-Width

Jan-Oliver Fröhlich* Ken-ichi Kawarabayashi†‡

Theodor Müller§ Julian Pott¶ Paul Wollan‖

Abstract

We show that all sufficiently large $(2k + 3)$-connected graphs of bounded tree-width are k-linked. Thomassen has conjectured that all sufficiently large $(2k + 2)$-connected graphs are k-linked.

1 Introduction

Given an integer $k \geq 1$, a graph G is *k-linked* if for any choice of $2k$ distinct vertices $s_1, \ldots, s_k$ and $t_1, \ldots, t_k$ of G there are disjoint paths $P_1, \ldots, P_k$ in G such that the end vertices of P_i are s_i and t_i for $i = 1, \ldots, k$. Menger's theorem implies that every k-linked graph is k-connected.

One can conversely ask how much connectivity (as a function of k) is required to conclude that a graph is k-linked. Larman and Mani [12] and Jung [8] gave the first proofs that a sufficiently highly connected graph is also k-linked. The bound was steadily improved until Bollobás and Thomason [3] gave the first linear bound on the necessary connectivity, showing that every $22k$-connected graph is k-linked. The current best bound shows that $10k$-connected graphs are also k-linked [18].

What is the best possible function $f(k)$ one could hope for which implies an $f(k)$-connected graph must also be k-linked? Thomassen [20] conjectured

*Fachbereich Mathematik, Universität Hamburg, Hamburg, Germany. jan-oliver.froehlich@math.uni-hamburg.de

†National Institute of Informatics, Tokyo, Japan.

‡JST, ERATO, Kawarabayashi Large Graph Project.

§Fachbereich Mathematik, Universität Hamburg, Hamburg, Germany.

¶Fachbereich Mathematik, Universität Hamburg, Hamburg, Germany.

‖Department of Computer Science, University of Rome "La Sapienza", Rome, Italy. wollan@di.uniroma1.it. Research supported by the European Research Council under the European Unions Seventh Framework Programme (FP7/2007 - 2013)/ERC Grant Agreement no. 279558.

that $(2k+2)$-connected graphs are k-linked. However, this was quickly proven to not be the case by Jørgensen with the following example [21]. Consider the graph obtained from K_{3k-1} obtained by deleting the edges of a matching of size k. This graph is $(3k-3)$-connected but is not k-linked. Thus, the best possible function $f(k)$ one could hope for to imply k-linked would be $3k-2$. However, all known examples of graphs which are roughly $3k$-connected but not k-linked are similarly of bounded size, and it is possible that Thomassen's conjectured bound is correct if one assumes that the graph has sufficiently many vertices.

In this paper, we show Thomassen's conjectured bound is almost correct with the additional assumption that the graph is large and has bounded tree-width. This is the main result of this article.

Theorem 1.1. *For all integers k and w there exists an integer N such that a graph G is k-linked if*

$$\kappa(G) \geq 2k+3, \qquad \text{tw}(G) < w, \quad \textit{and} \quad |G| \geq N.$$

where κ is the connectivity of the graph and tw *is the tree-width.*

The tree-width of the graph is a parameter commonly arising in the theory of graph minors; we will delay giving the definition until Section 2 where we give a more in depth discussion of how tree-width arises naturally in tackling the problem. The value $2k+2$ would be best possible; see Section 8 for examples of arbitrarily large graphs which are $(2k+1)$-connected but not k-linked.

Our work builds on the theory of graph minors in large, highly connected graphs begun by Böhme, Kawarabayashi, Maharry and Mohar [1]. Recall that a graph G contains K_t as a *minor* if K_t can be obtained from a subgraph of G by repeatedly contracting edges. Böhme et al. showed that there exists an absolute constant c such that every sufficiently large ct-connected graph contains K_t as a minor. This statement is not true without the assumption that the graph be sufficiently large, as there are examples of small graphs which are $(t\sqrt{\log t})$-connected but still have no K_t minor [11, 19]. In the case where we restrict our attention to small values of t, one is able to get an explicit characterisation of the large t-connected graphs which do not contain K_t as a minor.

Theorem 1.2 (Kawarabayashi et al. [10])**.** *There exists a constant N such that every 6-connected graph G on N vertices either contains K_6 as a minor or there exists a vertex $v \in V(G)$ such that $G-v$ is planar.*

Jorgensen [7] conjectures that Theorem 1.2 holds for all graphs without the additional restriction to graphs on a large number of vertices. In 2010, Norine and Thomas [17] announced that Theorem 1.2 could be generalised to arbitrary values of t to either find a K_t minor in a sufficiently large t-connected graph or alternatively, find a small set of vertices whose deletion leaves the graph planar. They have indicated that their methodology could be used to show a similar bound of $2k+3$ on the connectivity which ensures a large graph is k-linked.

2 Outline

In this section, we motivate our choice to restrict our attention to graphs of bounded tree-width and give an outline of the proof of Theorem 1.1.

We first introduce the basic definitions of tree-width. A *tree-decomposition* of a graph G is a pair $(T, \mathcal{X})$ where T is a tree and $\mathcal{X} = \{X_t \subseteq V(G) : t \in V(T)\}$ is a collection of subsets of $V(G)$ indexed by the vertices of T. Moreover, $\mathcal{X}$ satisfies the following properties.

1. $\bigcup_{t \in V(T)} X_t = V(G)$,
2. for all $e \in E(G)$, there exists $t \in V(T)$ such that both ends of e are contained in X_t, and
3. for all $v \in V(G)$, the subset $\{t \in V(T) : v \in X_t\}$ induces a connected subtree of T.

The sets in $\mathcal{X}$ are sometimes called the *bags* of the decomposition. The *width* of the decomposition is $\max_{t \in V(T)} |X_t| - 1$, and the *tree-width* of G is the minimum width of a tree-decomposition.

Robertson and Seymour showed that if a $2k$-connected graph contains K_{3k} as a minor, then it is k-linked [15]. Thus, when one considers $(2k+3)$-connected graphs which are not k-linked, one can further restrict attention to graphs which exclude a fixed clique minor. This allows one to apply the excluded minor structure theorem of Robertson and Seymour [16]. The structure theorem can be further strengthened if one assumes the graph has large tree-width [5]. This motivates one to analyse separately the case when the tree-width is large or bounded. The proofs of the main results in [1] and [10] similarly split the analysis into cases based on either large or bounded tree-width.

We continue with an outline of how the proof of Theorem 1.1 proceeds. Assume Theorem 1.1 is false, and let G be a $(2k+3)$-connected graph which is not k-linked. Fix a set $\{s_1, \ldots, s_k, t_1, \ldots, t_k\}$ such that there do not exist disjoint paths $P_1, \ldots, P_k$ where the ends of P_i are s_i and t_i for all i. Fix a tree-decomposition $(T, \mathcal{X})$ of G of minimal width w.

We first exclude the possibility that T has a high degree vertex. Assume t is a vertex of T of large degree. By Property 3 in the definition of a tree-decomposition, if we delete the set X_t of vertices from G, the resulting graph must have at least $\deg_T(t)$ distinct connected components. By the connectivity of G, each component contains $2k + 3$ internally disjoint paths from a vertex v to $2k+3$ distinct vertices in X_t. If the degree of t is sufficiently large, we conclude that the graph G contains a subdivision of $K_{a,2k+3}$ for some large value a. We now prove that that if a graph contains such a large complete bipartite subdivision and is $2k$-connected, then it must be k-linked (Lemma 7.1).

We conclude that the tree T does not have a high degree vertex, and consequently contains a long path. It follows that the graph G has a long *path decomposition*, that is, a tree-decomposition where the tree is a path. As the bags of the decomposition are linearly ordered by their position on the path, we simply give the path decomposition as a linearly ordered set of bags $(B_1, \ldots, B_t)$ for some large value t. At this point in the argument, the path-decomposition $(B_1, \ldots, B_t)$ may not have bounded width, but it will have the property that $|B_i \cap B_j|$ is bounded, and this will suffice for the argument to proceed. Section 3 examines this path decomposition in detail and presents a series of refinements allowing us to assume the path decomposition satisfies a set of desirable properties. For example, we are able to assume that $|B_i \cap B_{i+1}|$ is the same for all i, $1 \leq i < t$. Moreover, there exist a set $\mathcal{P}$ of $|B_1 \cap B_2|$ disjoint paths starting in B_1 and ending in B_t. We call these paths the *foundational linkage* and they play an important role in the proof. A further property of the path decomposition which we prove in Section 3 is that for each i, $1 < i < t$, if there is a bridge connecting two foundational paths in $\mathcal{P}$ in B_i, then for all j, $1 < j < t$, there exists a bridge connecting the same foundational paths in B_j. This allows us to define an auxiliary graph H with vertex set $\mathcal{P}$ and two vertices of $\mathcal{P}$ adjacent in H if there exists a bridge connecting them in some B_i $1 < i < t$.

Return to the linkage problem at hand; we have $2k$ terminals $s_1, \ldots, s_k$ and $t_1, \ldots, t_k$ which we would like to link appropriately, and $B_1, \ldots, B_t$ is our path decomposition with the foundational linkage running through it. Let the set $B_i \cap B_{i+1}$ be labeled S_i. As our path decomposition developed in the previous paragraph is very long, we can assume there exists some long subsection $B_i, B_{i+1}, \ldots, B_{i+a}$ such that no vertex of $s_1, \ldots, s_k, t_1, \ldots, t_k$ is contained in $\bigcup_i^{i+a} B_i - (S_{i-1} \cup S_{i+a})$ for some large value a. By Menger's theorem, there exist $2k$ paths linking $s_1, \ldots, s_k, t_1, \ldots, t_k$ to the set $S_{i-1} \cup S_{i+a}$. We attempt to link the terminals by continuing these paths into the subgraph induced by the vertex set $B_i \cup \cdots \cup B_{i+a}$. More specifically, we extend the paths along the foundational paths and attempt to link up the terminals with the bridges joining the various foundational paths in each of the B_j. By construction, the connections between foundational paths are the same in B_j for all j, $1 < j < t$; thus we translate the problem into a token game played on the auxiliary graph H. There each terminal has a corresponding token, and the desired linkage in G will exist if it is possible to slide the tokens around H in such a way to match up the tokens of the corresponding pairs of terminals. The token game is rigorously defined in Section 4, and we present a characterisation of what properties on H will allow us to find the desired linkage in G.

The final step in the proof of Theorem 1.1 is to derive a contradiction when H doesn't have sufficient complexity to allow us to win the token game. In order to do so, we use the high degree in G and a theorem of Robertson and Seymour on crossing paths. We give a series of technical results in preparation in Section 5 and Section 6 and present the proof of Theorem 1.1 in Section 7.

3 Stable Decompositions

In this section we present a result which, roughly speaking, ensures that a highly connected, sufficiently large graph of bounded tree-width either contains a subdivision of a large complete bipartite graph or has a long path decomposition whose bags all have similar structure.

Such a theorem was first established by Böhme, Maharry, and Mohar in [2] and extended by Kawarabayashi, Norine, Thomas, and Wollan in [9], both using techniques from [13]. We shall prove a further extension based on the result by Kawarabayashi et al. from [9] so our terminology and methods will be close to theirs.

For all basic definitions and notation we refer to Diestel's textbook [4]. We begin this section with a general Lemma about nested separations. Let G be a graph. A *separation* of G is an ordered pair (A, B) of sets $A, B \subseteq V(G)$ such that $G[A] \cup G[B] = G$. If (A, B) is a separation of G, then $A \cap B$ is called its *separator* and $|A \cap B|$ its *order*. Two separations (A, B) and (A', B') of G are called *nested* if either $A \subseteq A'$ and $B \supseteq B'$ or $A \supseteq A'$ and $B \subseteq B'$. In the former case we write $(A, B) \leq (A', B')$ and in the latter $(A, B) \geq (A', B')$. This defines a partial order $\leq$ on all separations of G. A set $\mathcal{S}$ of separations is called *nested* if the separations of $\mathcal{S}$ are pairwise nested, that is, $\leq$ is a linear order on $\mathcal{S}$. To avoid confusion about the order of the separations in $\mathcal{S}$ we do not use the usual terms like smaller, larger, maximal, and minimal when talking about this linear order but instead use *left*, *right*, *rightmost*, and *leftmost*, respectively (we still use *successor* and *predecessor* though). To distinguish $\leq$ from $<$ we say 'left' for the former and 'strictly left' for the latter (same for $\geq$ and right).

If (A, B) and (A', B') are both separations of G, then so are $(A \cap A', B \cup B')$ and $(A \cup A', B \cap B')$ and a simple calculation shows that the orders of $(A \cap A', B \cup B')$ and $(A \cup A', B \cap B')$ sum up to the same number as the orders of (A, B) and (A', B'). Clearly each of $(A \cap A', B \cup B')$ and $(A \cup A', B \cap B')$ is nested with both, (A, B) and (A', B').

For two sets $X, Y \subseteq V(G)$ we say that a separation (A, B) of G is an *X–Y separation* if $X \subseteq A$ and $Y \subseteq B$. If (A, B) and (A', B') are X–Y separations in G, then so are $(A \cap A', B \cup B')$ and $(A \cup A', B \cap B')$. Furthermore, if (A, B) and (A', B') are X–Y separations of G of minimum order, say m, then so are $(A \cap A', B \cup B')$ and $(A \cup A', B \cap B')$ as none of the latter two can have order less than m but their orders sum up to $2m$.

Lemma 3.1. *Let G be a graph and $X, Y, Z \subseteq V(G)$. If for every $z \in Z$ there is an X–Y separation of G of minimal order with z in its separator, then there is a nested set $\mathcal{S}$ of X–Y separations of minimal order such that their separators cover Z.*

Proof. Let $\mathcal{S}$ be a maximal nested set of X–Y separations of minimal order in G (as $\mathcal{S}$ is finite the existence of a leftmost and a rightmost element in any subset of $\mathcal{S}$ is trivial). Suppose for a contradiction that some $z \in Z$ is not contained in any separator of the separations of $\mathcal{S}$.

Set $\mathcal{S}_L := \{(A, B) \in \mathcal{S} \mid z \in B\}$ and $\mathcal{S}_R := \{(A, B) \in \mathcal{S} \mid z \in A\}$. Clearly $\mathcal{S}_L \cup \mathcal{S}_R = \mathcal{S}$ and $\mathcal{S}_L \cap \mathcal{S}_R = \emptyset$. Moreover, if $\mathcal{S}_L$ and $\mathcal{S}_R$ are both non-empty, then the rightmost element (A_L, B_L) of $\mathcal{S}_L$ is the predecessor of the leftmost element (A_R, B_R) of $\mathcal{S}_R$ in $\mathcal{S}$. Loosely speaking, $\mathcal{S}_L$ and $\mathcal{S}_R$ contain the separations of $\mathcal{S}$ "on the left" and "on the right" of z, respectively, and (A_L, B_L) and (A_R, B_R) are the separations of $\mathcal{S}_L$ and $\mathcal{S}_R$ whose separators are "closest" to z.

By assumption there is an X–Y separation (A, B) of minimal order in G with $z \in A \cap B$. Set

$$(A', B') := (A \cup A_L, B \cap B_L) \quad \text{and} \quad (A'', B'') := (A' \cap A_R, B' \cup B_R)$$

(but $(A', B') := (A, B)$ if $\mathcal{S}_L = \emptyset$ and $(A'', B'') := (A', B')$ if $\mathcal{S}_R = \emptyset$). As (A_L, B_L), (A, B), and (A_R, B_R) are all X–Y separations of minimal order in

G so must be (A', B') and (A'', B''). Moreover, we have $z \in A'' \cap B''$ and thus $(A'', B'') \notin \mathcal{S}$.

By construction we have $(A_L, B_L) \le (A', B')$ and $(A'', B'') \le (A_R, B_R)$. To verify that $(A_L, B_L) \le (A'', B'')$ we need to show $A_L \subseteq A' \cap A_R$ and $B_L \supseteq B' \cup B_R$. All required inclusions follow from $(A_L, B_L) \le (A', B')$ and $(A_L, B_L) \le (A_R, B_R)$. So by transitivity (A'', B'') is right of all elements of $\mathcal{S}_L$ and left of all elements of $\mathcal{S}_R$, in particular, it is nested with all elements of $\mathcal{S}$, contradicting the maximality of the latter. □

We assume that every path comes with a fixed linear order of its vertices. If a path arises as an X–Y path, then we assume it is ordered from X to Y and if a path Q arises as a subpath of some path P, then we assume that Q is ordered in the same direction as P unless explicitly stated otherwise.

Given a vertex v on a path P we write Pv for the initial subpath of P with last vertex v and vP for the final subpath of P with first vertex v. If v and w are both vertices of P, then by vPw or wPv we mean the subpath of P that ends in v and w and is ordered from v to w or from w to v, respectively. By P^{-1} we denote the path P with inverse order.

Let $\mathcal{P}$ be a set of disjoint paths in some graph G. We do not distinguish between $\mathcal{P}$ and the graph $\bigcup \mathcal{P}$ formed by uniting these paths; both will be denoted by $\mathcal{P}$. By a *path of* $\mathcal{P}$ we always mean an element of $\mathcal{P}$, not an arbitrary path in $\bigcup \mathcal{P}$.

Let G be a graph. For a subgraph $S \subseteq G$ an *S-bridge in G* is a connected subgraph $B \subseteq G$ such that B is edge-disjoint from S and either B is a single edge with both ends in S or there is a component C of $G - S$ such that B consists of all edges that have at least one end in C. We call a bridge *trivial* in the former case and *non-trivial* in the latter. The vertices in $V(B) \cap V(S)$ and $V(B) \setminus V(S)$ are called the *attachments* and the *inner vertices* of B, respectively. Clearly an S-bridge has an inner vertex if and only if it is non-trivial. We say that an S-bridge B *attaches* to a subgraph $S' \subseteq S$ if B has an attachment in S'. Note that S-bridges are pairwise edge-disjoint and each common vertex of two S-bridges must be an attachment of both.

A *branch vertex* of S is a vertex of degree $\neq 2$ in S and a *segment* of S is a maximal path in S such that its ends are branch vertices of S but none of its inner vertices are. An S-bridge B in G is called *unstable* if some segment of S contains all attachments of B, and *stable* otherwise. If an unstable S-bridge B has at least two attachments on a segment P of S, then we call P a *host* of B and say that B is *hosted by* P. For a subgraph $H \subseteq G$ we say that two segments of S are *S-bridge adjacent* or just *bridge adjacent* in H if H contains an S-bridge that attaches to both.

If a graph is the union of its segments and no two of its segments have the same end vertices, then it is called *unambiguous* and *ambiguous* otherwise. It is easy to see that a graph S is unambiguous if and only if all its cycles contain a least three branch vertices. In our application S will always be a union of disjoint paths so its segments are precisely these paths and S is trivially unambiguous.

Let $S \subseteq G$ be unambiguous. We say that $S' \subseteq G$ is a *rerouting* of S if there is a bijection φ from the segments of S to the segments of S' such that every segment P of S has the same end vertices as $\varphi(P)$ (and thus φ is unique by the unambiguity). If S' contains no edge of a stable S-bridge, then we call S' a *proper rerouting* of S. Clearly any rerouting of the unambiguous graph S has the same branch vertices as S and hence is again unambiguous.

The following Lemma states two observations about proper reroutings. The proofs are both easy and hence we omit them.

Lemma 3.2. *Let S' be a proper rerouting of an unambiguous graph $S \subseteq G$ and let φ be as in the definition. Both of the following statements hold.*

(i) Every hosted S-bridge has a unique host. For each segment P of S the segment $\varphi(P)$ of S' is contained in the union of P and all S-bridges hosted by P.

(ii) For every stable S-bridge B there is a stable S'-bridge B' with $B \subseteq B'$. Moreover, if B attaches to a segment P of S, then B' attaches to $\varphi(P)$.

Note that Lemma 3.2 (ii) implies that no unstable S'-bridge contains an edge of a stable S-bridge. Together with (i) this means that being a proper rerouting of an unambiguous graph is a transitive relation.

The next Lemma is attributed to Tutte; we refer to [9, Lemma 2.2] for a proof[1].

Lemma 3.3. *Let G be a graph and $S \subseteq G$ unambiguous. There exists a proper rerouting S' of S in G such that if B' is an S'-bridge hosted by some segment P' of S', then B' is non-trivial and there are vertices $v, w \in V(P')$ such that the component of $G - \{v, w\}$ that contains $B' - \{v, w\}$ is disjoint from $S' - vP'w$.*

[1] To check that Lemma 2.2 in [9] implies our Lemma 3.3 note that if S' is obtained from S by "a sequence of proper reroutings" as defined in [9], then by transitivity S' is a proper rerouting of S according to our definition. And although not explicitly included in the statement, the given proof shows that no trivial S'-bridge can be unstable.

This implies that the segments of S' are induced paths in G as trivial S'-bridges cannot be unstable and no two segments of S' have the same end vertices.

Let G be a graph. A set of disjoint paths in G is called a *linkage*. If $X, Y \subseteq V(G)$ with $k := |X| = |Y|$, then a set of k disjoint X–Y paths in G is called an *X–Y linkage* or a *linkage from X to Y*. Let $\mathcal{W} = (W_0, \ldots, W_l)$ be an ordered tuple of subsets of $V(G)$. Then l is the *length* of $\mathcal{W}$, the sets W_i with $0 \leq i \leq l$ are its *bags*, and the sets $W_{i-1} \cap W_i$ with $1 \leq i \leq l$ are its *adhesion sets*. We refer to the bags W_i with $1 \leq i \leq l-1$ as *inner* bags. When we say that a bag W of $\mathcal{W}$ *contains* some graph H, we mean $H \subseteq G[W]$. Given an inner bag W_i of $\mathcal{W}$, the sets $W_{i-1} \cap W_i$ and $W_i \cap W_{i+1}$ are called the *left* and *right* adhesion set of W_i, respectively. Whenever we introduce a tuple $\mathcal{W}$ as above without explicitly naming its elements, we shall denote them by $W_0, \ldots, W_l$ where l is the length of $\mathcal{W}$. For indices $0 \leq j \leq k \leq l$ we use the shortcut $W_{[j,k]} := \bigcup_{i=j}^{k} W_i$.

The tuple $\mathcal{W}$ with the following five properties is called a *slim decomposition* of G.

(L1) $\bigcup \mathcal{W} = V(G)$ and every edge of G is contained in some bag of $\mathcal{W}$.

(L2) If $0 \leq i \leq j \leq k \leq l$, then $W_i \cap W_k \subseteq W_j$.

(L3) All adhesion sets of $\mathcal{W}$ have the same size.

(L4) No bag of $\mathcal{W}$ contains another.

(L5) G contains a $(W_0 \cap W_1)$–$(W_{l-1} \cap W_l)$ linkage.

The unique size of the adhesion sets of a slim decomposition is called its *adhesion*. A linkage $\mathcal{P}$ as in (L5) together with an enumeration $P_1, \ldots, P_q$ of its paths is called a *foundational linkage* for $\mathcal{W}$ and its members are called *foundational paths*. Each path P_α contains a unique vertex of every adhesion set of $\mathcal{W}$ and we call this vertex the *α-vertex* of that adhesion set. For an inner bag W of $\mathcal{W}$ the α-vertex in the left and right adhesion set of W are called the *left* and *right* α-vertex of W, respectively. Note that $\mathcal{P}$ is allowed to contain trivial paths so $\bigcap \mathcal{W}$ may be non-empty.

The enumeration of a foundational linkage $\mathcal{P}$ for $\mathcal{W}$ is a formal tool to compare arbitrary linkages between adhesion sets of $\mathcal{W}$ to $\mathcal{P}$ by their 'induced permutation' as detailed below. When considering another foundational linkage $\mathcal{Q} = \{Q_1, \ldots, Q_q\}$ for $\mathcal{W}$ we shall thus always assume that it induces the same enumeration as $\mathcal{P}$ on $W_0 \cap W_1$, in other words, Q_α and P_α start on the same vertex.

Suppose that $\mathcal{W}$ is a slim decomposition of some graph G with foundational linkage $\mathcal{P}$. Then any $\mathcal{P}$-bridge B in G is contained in a bag of $\mathcal{W}$, and this bag is unique unless B is trivial and contained in one or more adhesion sets.

We say that a linkage $\mathcal{Q}$ in a graph H is *p-attached* if each path of $\mathcal{Q}$ is induced in H and if some non-trivial $\mathcal{Q}$-bridge B attaches to a non-trivial path P of $\mathcal{Q}$, then either B attaches to another non-trivial path of $\mathcal{Q}$ or there are at least $p-2$ trivial paths Q of $\mathcal{Q}$ such that H contains a $\mathcal{Q}$-bridge (which may be different from B) attaching to P and Q.

We call a pair $(\mathcal{W}, \mathcal{P})$ of a slim decomposition $\mathcal{W}$ of G and a foundational linkage $\mathcal{P}$ for $\mathcal{W}$ a *regular decomposition of attachedness p of G* if there is an integer p such that the axioms (L6), (L7), and (L8) hold.

(L6) $\mathcal{P}[W]$ is p-attached in $G[W]$ for all inner bags W of $\mathcal{W}$.

(L7) A path $P \in \mathcal{P}$ is trivial if $P[W]$ is trivial for some inner bag W of $\mathcal{W}$.

(L8) For every $P, Q \in \mathcal{P}$, if some inner bag of $\mathcal{W}$ contains a $\mathcal{P}$-bridge attaching to P and Q, then every inner bag of $\mathcal{W}$ contains such a $\mathcal{P}$-bridge.

The integer p is not unique: A regular decomposition of attachedness p has attachedness p' for all integers $p' \leq p$. Note that $\mathcal{P}$ satisfies (L7) if and only if every vertex of G either lies in at most two bags of $\mathcal{W}$ or in all bags. This means that either all foundational linkages for $\mathcal{W}$ satisfy (L7) or none.

The next Theorem follows[2] from the Lemmas 3.1, 3.2, and 3.5 in [9].

Theorem 3.4 (Kawarabayashi et al. [9])**.** *For all integers $a, l, p, w \geq 0$ there exists an integer N with the following property. If G is a p-connected graph of tree-width less than w with at least N vertices, then either G contains a subdivision of $K_{a,p}$, or G has a regular decomposition of length at least l, adhesion at most w, and attachedness p.*

Note that [9] features a stronger version of Theorem 3.4, namely Theorem 3.8, which includes an additional axiom (L9). We omit that axiom since our arguments do not rely on it.

Let $(\mathcal{W}, \mathcal{P})$ be a slim decomposition of adhesion q and length l for a graph G. Suppose that $\mathcal{Q}$ is a linkage from the left adhesion set of W_i to the right adhesion set of W_j for two indices i and j with $1 \leq i \leq j < l$. The enumeration $P_1, \dots, P_q$ of $\mathcal{P}$ induces an enumeration $Q_1, \dots, Q_q$ of $\mathcal{Q}$

[2] The statement of Lemma 3.1 in [9] only asserts the existence of a minor isomorphic to $K_{a,p}$ rather than a subdivision of $K_{a,p}$ like we do. But its proof refers to an argument in the proof of [13, Theorem 3.1] which actually gives a subdivision.

where Q_α is the path of $\mathcal{Q}$ starting in the left α-vertex of W_i. The map $\pi : \{1, \dots, q\} \to \{1, \dots, q\}$ such that Q_α ends in the right $\pi(\alpha)$-vertex of W_j for $\alpha = 1, \dots, q$ is a permutation because $\mathcal{Q}$ is a linkage. We call it the *induced permutation of* $\mathcal{Q}$. Clearly the induced permutation of $\mathcal{Q}$ is the composition of the induced permutations of $\mathcal{Q}[W_i]$, $\mathcal{Q}[W_{i+1}]$, ..., $\mathcal{Q}[W_j]$. For any permutation π of $\{1, \dots, q\}$ and any graph Γ on $\{1, \dots, q\}$ we write $\pi\Gamma$ to denote the graph $(\{\pi(\alpha) \mid \alpha \in V(\Gamma)\}, \{\pi(\alpha)\pi(\beta) \mid \alpha\beta \in E(\Gamma)\})$. For a subset $X \subseteq \{1, \dots, q\}$ we set $\mathcal{Q}_X := \{Q_\alpha \mid \alpha \in X\}$.

Keep in mind that the enumerations $\mathcal{P}$ induces on linkages $\mathcal{Q}$ as above always depend on the adhesion set where the considered linkage starts. For example let $\mathcal{Q}$ be as above and for some index i' with $i < i' \leq j$ set $\mathcal{Q}' := \mathcal{Q}[W_{[i',j]}]$. Then $Q_\alpha[W_{[i',j]}]$ need not be the same as Q'_α. More precisely, we have $Q_\alpha[W_{[i',j]}] = Q'_{\tau(\alpha)}$ where τ denotes the induced permutation of $\mathcal{Q}[W_{[i,i'-1]}]$.

For some subgraph H of G the *bridge graph of* $\mathcal{Q}$ *in* H, denoted $B(H, \mathcal{Q})$, is the graph with vertex set $\{1, \dots, q\}$ in which $\alpha\beta$ is an edge if and only if Q_α and Q_β are $\mathcal{Q}$-bridge adjacent in H. Any $\mathcal{Q}$-bridge B in H that attaches to Q_α and Q_β is said to *realise* the edge $\alpha\beta$. We shall sometimes think of induced permutations as maps between bridge graphs.

For a slim decomposition $\mathcal{W}$ of length l of G with foundational linkage $\mathcal{P}$ we define the *auxiliary graph* $\Gamma(\mathcal{W}, \mathcal{P}) := B(G[W_{[1,\ l-1]}], \mathcal{P})$. Clearly $B(G[W], \mathcal{P}[W]) \subseteq \Gamma(\mathcal{W}, \mathcal{P})$ for each inner bag W of $\mathcal{W}$ and if $(\mathcal{W}, \mathcal{P})$ is regular, then by (L8) we have equality.

Set $\lambda := \{\alpha \mid P_\alpha$ is non-tivial$\}$ and $\theta := \{\alpha \mid P_\alpha$ is trivial$\}$. Given a subgraph $\Gamma \subseteq \Gamma(\mathcal{W}, \mathcal{P})$ and some foundational linkage $\mathcal{Q}$ for $\mathcal{W}$, we write $G^{\mathcal{Q}}_{\Gamma}$ for the graph obtained by deleting $\mathcal{Q} \setminus \mathcal{Q}_{V(\Gamma)}$ from the union of $\mathcal{Q}$ and those $\mathcal{Q}$-bridges in inner bags of $\mathcal{W}$ that realise an edge of Γ or attach to $\mathcal{Q}_{V(\Gamma)\cap\lambda}$ but to no path of $\mathcal{Q}_{\lambda \setminus V(\Gamma)}$. For a subset $V \subseteq \{1, \dots, q\}$ we write $G^{\mathcal{Q}}_V$ instead of $G^{\mathcal{Q}}_{\Gamma(\mathcal{W},\mathcal{P})[V]}$. Note that $\mathcal{Q}_\theta = \mathcal{P}_\theta$. Hence $G^{\mathcal{P}}_\lambda$ and $G^{\mathcal{Q}}_\lambda$ are the same graph and we denote it by G_λ.

A regular decomposition $(\mathcal{W}, \mathcal{P})$ of a graph G is called *stable* if it satisfies the following two axioms where $\lambda := \{\alpha \mid P_\alpha$ is non-trivial$\}$.

(L10) If $\mathcal{Q}$ is a linkage from the left to the right adhesion set of some inner bag of $\mathcal{W}$, then its induced permutation is an automorphism of $\Gamma(\mathcal{W}, \mathcal{P})$.

(L11) If $\mathcal{Q}$ is a linkage from the left to the right adhesion set of some inner bag W of $\mathcal{W}$, then every edge of $B(G[W], \mathcal{Q})$ with one end in λ is also an edge of $\Gamma(\mathcal{W}, \mathcal{P})$.

Given these definitions we can further expound our strategy to prove the main theorem: We will reduce the given linkage problem to a linkage problem with start and end vertices in $W_0 \cup W_l$ for some stable regular decomposition $(\mathcal{W}, \mathcal{P})$ of length l. The stability ensures that we maximised the number of edges of $\Gamma(\mathcal{W}, \mathcal{P})$, i.e. no rerouting of $\mathcal{P}$ will give rise to new bridge adjacencies. We will focus on a subset $\lambda_0 \subseteq \lambda$ and show that the minimum degree of G forces a high edge density in $G^{\mathcal{P}}_{\lambda_0}$, leading to a high number of edges in $\Gamma(\mathcal{W}, \mathcal{P})[\lambda_0]$. Using combinatoric arguments, which we elaborate in Section 4, we show that we can find linkages using segments of $\mathcal{P}$ and $\mathcal{P}$-bridges in $G^{\mathcal{P}}_{\lambda_0}$ to realise any matching of start and end vertices in $W_0 \cup W_l$, showing that G is in fact k-linked.

We strengthen Theorem 3.4 by the assertion that the regular decomposition can be chosen to be stable. We like to point out that, even with the left out axiom (L9) included in the definition of a regular decomposition, Theorem 3.5 would hold. By almost the same proof as in [9] one could also obtain a stronger version of (L8) stating that for every subset $\mathcal{R}$ of $\mathcal{P}$ if some inner bag of $\mathcal{W}$ contains a $\mathcal{P}$-bridge attaching to every path of $\mathcal{R}$ but to no path of $\mathcal{P} \setminus \mathcal{R}$, then every inner bag does.

Theorem 3.5. *For all integers $a, l, p, w \geq 0$ there exists an integer N with the following property. If G is a p-connected graph of tree-width less than w with at least N vertices, then either G contains a subdivision of $K_{a,p}$, or G has a stable regular decomposition of length at least l, adhesion at most w, and attachedness p.*

Before we start with the formal proof let us introduce its central concepts: disturbances and contractions. Let $(\mathcal{W}, \mathcal{P})$ be a regular decomposition of a graph G. A linkage $\mathcal{Q}$ is called a *twisting* $(\mathcal{W}, \mathcal{P})$*-disturbance* if it violates (L10) and it is called a *bridging* $(\mathcal{W}, \mathcal{P})$*-disturbance* if it violates (L11). By a $(\mathcal{W}, \mathcal{P})$*-disturbance* we mean either of these two and a disturbance may be twisting and bridging at the same time. If the referred regular decomposition is clear from the context, then we shall not include it in the notation and just speak of a disturbance. Note that a disturbance is always a linkage from the left to the right adhesion set of an inner bag of $\mathcal{W}$.

Given a disturbance $\mathcal{Q}$ in some inner bag W of $\mathcal{W}$ which is neither the first nor the last inner bag of $\mathcal{W}$, it is not hard to see that replacing $\mathcal{P}[W]$ with $\mathcal{Q}$ yields a foundational linkage $\mathcal{P}'$ for $\mathcal{W}$ such that $\Gamma(\mathcal{W}, \mathcal{P}')$ properly contains $\Gamma(\mathcal{W}, \mathcal{P})$ and we shall make this precise in the proof. As the auxiliary graph can have at most $\binom{w}{2}$ edges, we can repeat this step until no disturbances (with respect to the current decomposition) are left and we should end up with a stable regular decomposition, given that we can somehow preserve the regularity.

This is done by "contracting" the decomposition in a certain way. The technique is the same as in [2] or [9]. Given a regular decomposition $(\mathcal{W}, \mathcal{P})$ of length l of some graph G and a subsequence $i_1, \ldots, i_n$ of $1, \ldots, l$, the *contraction of* $(\mathcal{W}, \mathcal{P})$ *along* $i_1, \ldots, i_n$ is the pair $(\mathcal{W}', \mathcal{P}')$ defined as follows. We let $\mathcal{W}' := (W'_0, W'_1, \ldots, W'_n)$ with $W'_0 := W_{[0,\ i_1 - 1]}$,

$$W'_j := W_{[i_j,\ i_{j+1}-1]} \quad \text{for} \quad j = 1, \ldots, n-1,$$

$W_n := W_{[i_n, l]}$, and $\mathcal{P}' = \mathcal{P}[W'_{[1,\ n-1]}]$ (with the induced enumeration).

Lemma 3.6. *Let* $(\mathcal{W}', \mathcal{P}')$ *be the contraction of a regular decomposition* $(\mathcal{W}, \mathcal{P})$ *of some graph* G *of adhesion* q *and attachedness* p *along the sequence* $i_1, \ldots, i_n$. *Then the following two statements hold.*

(i) $(\mathcal{W}', \mathcal{P}')$ *is a regular decomposition of length* n *of* G *of adhesion* q *and attachedness* p, *and* $\Gamma(\mathcal{W}', \mathcal{P}') = \Gamma(\mathcal{W}, \mathcal{P})$.

(ii) The decomposition $(\mathcal{W}', \mathcal{P}')$ *is stable if and only if none of the inner bags* $W_{i_1}, W_{i_1+1}, \ldots, W_{i_n-1}$ *of* $\mathcal{W}$ *contains a* $(\mathcal{W}, \mathcal{P})$*-disturbance.*

Proof. The first statement is Lemma 3.3 of [9]. The second statement follows from the fact that an inner bag W'_j of $\mathcal{W}'$ contains a $(\mathcal{W}', \mathcal{P}')$-disturbance if and only if one of the bags W_i of $\mathcal{W}$ with $i_j \le i < i_{j+1}$ contains a $(\mathcal{W}, \mathcal{P})$-disturbance (unless $\mathcal{W}'$ has no inner bag, that is, $n = 1$). The "if" direction is obvious and for the "only if" direction recall that the induced permutation of $\mathcal{P}'[W'_j]$ is the composition of the induced permutations of the $\mathcal{P}[W_i]$ with $i_j \le i < i_{j+1}$ and every $\mathcal{P}'$-bridge in W'_j is also a $\mathcal{P}$-bridge and hence must be contained in some bag W_i with $i_j \le i < i_{j+1}$. □

Let $\mathcal{Q}$ be a linkage in a graph H and denote the trivial paths of $\mathcal{Q}$ by Θ. Let $\mathcal{Q}'$ be the union of Θ with a proper rerouting of $\mathcal{Q} \setminus \Theta$ obtained from applying Lemma 3.3 to $\mathcal{Q} \setminus \Theta$ in $H - \Theta$. We call $\mathcal{Q}'$ a *bridge stabilisation of* $\mathcal{Q}$ *in* H. The next Lemma tailors Lemma 3.2 and Lemma 3.3 to our application.

Lemma 3.7. *Let* $\mathcal{Q}$ *be a linkage in a graph* H. *Denote by* Θ *the trivial paths of* $\mathcal{Q}$ *and let* $\mathcal{Q}'$ *be a bridge stabilisation of* $\mathcal{Q}$ *in* H. *Let* P *and* Q *be paths of* $\mathcal{Q}$ *and let* P' *and* Q' *be the unique paths of* $\mathcal{Q}'$ *with the same end vertices as* P *and* Q, *respectively. Then the following statements hold.*

(i) P' *is contained in the union of* P *with all* $\mathcal{Q}$*-bridges in* H *that attach to* P *but to no other path of* $\mathcal{Q} \setminus \Theta$.

(ii) If P and Q are $\mathcal{Q}$-bridge adjacent in H and one of them is non-trivial, then P' and Q' are $\mathcal{Q}'$-bridge adjacent in H.

(iii) Let Z be the set of end vertices of the paths of $\mathcal{Q}$. If p is an integer such that for every vertex x of $H - Z$ there is an x–Z fan of size p, then $\mathcal{Q}'$ is p-attached.

Proof.

(i) This is trivial if $P \in \Theta$ and follows easily from Lemma 3.2 (i) otherwise.

(ii) The statement follows directly from Lemma 3.2 (ii) if P and Q are both non-trivial so we may assume that $P = P' \in \Theta$ and Q is non-trivial. By assumption there is a P–Q path R in H. Clearly $R \cup Q$ contains the end vertices of Q'. On the other hand, by (i) it is clear that $Q \cap \mathcal{Q}' \subseteq Q'$. We claim that $R \cap \mathcal{Q}' \subseteq Q'$. Since R is internally disjoint from $\mathcal{Q}$ all its inner vertices are inner vertices of some $(\mathcal{Q} \setminus \Theta)$-bridge B. If B is stable or unstable but not hosted by any path of $\mathcal{Q}$ (that is, it has at most one attachment), then Lemma 3.2 implies that no path of $\mathcal{Q}'$ contains an inner vertex of B and that our claim follows. If B is hosted by a path of $\mathcal{Q}$, then this path must clearly be Q and thus by Lemma 3.2 (i) $R \cap \mathcal{Q}' \subseteq Q'$ as claimed. Hence $R \cup Q$ contains a P–Q' path that is internally disjoint from $\mathcal{Q}'$ as desired.

(iii) Clearly all paths of $\mathcal{Q}'$ are induced in H, either because they are trivial or by Lemma 3.3. Let B be a non-trivial hosted $\mathcal{Q}'$-bridge and let Q' be the non-trivial path of $\mathcal{Q}'$ to which it attaches. Then by Lemma 3.3 there are vertices v and w on Q' and a separation (X, Y) of H such that $V(B) \subseteq X$, $X \cap Y \subseteq \{v, w\} \cup V(\Theta)$, and apart from the inner vertices of $vQ'w$ all vertices of $\mathcal{Q}'$ are in Y, in particular, $Z \subseteq Y$. But B has an inner vertex x which must be in $X \setminus Y$. So by assumption there is an x–$\{v, w\} \cup V(\Theta)$ fan of size p in $G[X]$ and thus also an x–Θ fan of size $p - 2$. It is easy to see that this can gives rise to the desired $\mathcal{Q}'$-bridge adjacencies in H.

□

Proof of Theorem 3.5. We will trade off some length of a regular decomposition to gain edges in its auxiliary graph. To quantify this we define the function $f : \mathbb{N}_0 \to \mathbb{N}_0$ by $f(m) := (zlw!)^m l$ where $z := 2^{\binom{w}{2}}$ and call a regular decomposition $(\mathcal{W}, \mathcal{P})$ of a graph G *valid* if it has adhesion at most w, attachedness p, and length at least $f(m)$ where m is the number of edges in the complement of $\Gamma(\mathcal{W}, \mathcal{P})$ that are incident with at least one non-trivial path of $\mathcal{P}$.

Set $\lambda := f\left(\binom{w}{2}\right)$ and let N be the integer returned by Theorem 3.4 when invoked with parameters a, λ, p, and w. We claim that the assertion of Theorem 3.5 is true for this choice of N. Let G be a p-connected graph of tree-width less than w with at least N vertices and suppose that G does not contain a subdivision of $K_{a,p}$. Then by the choices of N and λ the graph G has a valid decomposition (the foundational linkage has at most w paths so there can be at most $\binom{w}{2}$ non-edges in the auxiliary graph). Among all valid decompositions of G pick $(\mathcal{W}, \mathcal{P})$ such that the number of edges of $\Gamma(\mathcal{W}, \mathcal{P})$ is maximal and denote the length of $(\mathcal{W}, \mathcal{P})$ by n.

We may assume that for any integer k with $0 \leq k \leq n - l$ one of the $l - 1$ consecutive inner bags $W_{k+1}, \ldots, W_{k+l-1}$ of $\mathcal{W}$ contains a disturbance. If not, then by Lemma 3.6, the contraction of $(\mathcal{W}, \mathcal{P})$ along the sequence $k + 1, k + 2, \ldots, k + l$ is a stable regular decomposition of G of length l, adhesion at most w, and attachedness p as desired.

Claim 3.5.1. *Let $1 \leq k \leq k' \leq n - 1$ with $k' - k \geq lw! - 1$. Then the graph $H := G[W_{[k,k']}]$ contains a linkage $\mathcal{Q}$ from the left adhesion set of W_k to the right adhesion set of $W_{k'}$ such that $B(H, \mathcal{Q})$ is a proper supergraph of $\Gamma(\mathcal{W}, \mathcal{P})$, the induced permutation π of $\mathcal{Q}$ is the identity, and $\mathcal{Q}$ is p-attached in H.*

Proof. There are indices $k_0 := k, k_1, \ldots, k_{w!} := k' + 1$, such that we have $k_j - k_{j-1} \geq l$ for $j \in \{1, \ldots, w!\}$. For each $j \in \{0, \ldots, w!-1\}$ one of the at least $l-1$ consecutive inner bags $W_{k_j+1}, W_{k_j+2}, \ldots, W_{k_{j+1}-1}$ contains a disturbance $\mathcal{Q}_j$ by our assumption. Let W_{i_j} be the bag of $\mathcal{W}$ that contains $\mathcal{Q}_j$ and let $\mathcal{Q}'_j$ be the bridge stabilisation of $\mathcal{Q}_j$ in $G[W_{i_j}]$.

If $\mathcal{Q}_j$ is a twisting $(\mathcal{W}, \mathcal{P})$-disturbance, then so is $\mathcal{Q}'_j$ as they have the same induced permutation. If $\mathcal{Q}_j$ is a bridging $(\mathcal{W}, \mathcal{P})$-disturbance, then so is $\mathcal{Q}'_j$ by Lemma 3.7 (ii). The set Z of end vertices of $\mathcal{Q}_j$ is the union of both adhesion sets of W_{i_j} and clearly for every vertex $x \in W_{i_j} \setminus Z$ there is an x–Z fan of size p in $G[W_{i_j}]$ as G is p-connected. So by Lemma 3.7 (iii) the linkage $\mathcal{Q}'_j$ is p-attached in $G[W_{i_j}]$.

For every $j \in \{0, \ldots, w!-1\}$ denote the induced permutation of $\mathcal{Q}'_j$ by π_j. Since the symmetric group S_q has order at most $q! \leq w!$ we can pick[3] indices j_0 and j_1 with $0 \leq j_0 \leq j_1 \leq w! - 1$ such that $\pi_{j_1} \circ \pi_{j_1-1} \ldots \circ \pi_{j_0} = \mathrm{id}$.

[3] Let $(G, \cdot)$ be a group of order n and $g_1, \ldots, g_n \in G$. Then of the $n + 1$ products $h_k := \prod_{i=1}^{k} g_i$ for $0 \leq k \leq n$, two must be equal by the pigeon hole principle, say $h_k = h_l$ with $k < l$. This means $\prod_{i=k+1}^{l} g_i = e$, where e is the neutral element of G.

Let $\mathcal{Q}$ be the linkage from the left adhesion set of W_k to the right adhesion set of $W_{k'}$ in H obtained from $\mathcal{P}[W_{[k,k']}]$ by replacing $\mathcal{P}[W_{i_j}]$ with $\mathcal{Q}'_j$ for all $j \in \{j_0, \ldots, j_1\}$. Of all the restrictions of $\mathcal{Q}$ to the bags $W_k, \ldots, W_{k'}$ only $\mathcal{Q}[W_{i_j}] = \mathcal{Q}_j$ with $j_0 \leq j \leq j_1$ need not induce the identity permutation. However, the composition of their induced permutations is the identity by construction and therefore the induced permutation of $\mathcal{Q}$ is the identity.

To see that $B(H, \mathcal{Q})$ is a supergraph of $\Gamma(\mathcal{W}, \mathcal{P})$ note that $k < i_{j_0}$ so $\mathcal{Q}$ and $\mathcal{P}$ coincide on W_k and hence by (L8) we have

$$\Gamma(\mathcal{W}, \mathcal{P}) = B(G[W_k], \mathcal{P}[W_k]) \subseteq B(H, \mathcal{Q}).$$

It remains to show that $B(H, \mathcal{Q})$ contains an edge that is not in $\Gamma(\mathcal{W}, \mathcal{P})$. Set $W := W_{i_{j_0}}$, $W' := W_{i_{j_0}+1}$, and $\pi := \pi_{j_0}$. If $\mathcal{Q}'_{j_0}$ is a bridging disturbance, then $B_0 := B(G[W], \mathcal{Q}[W])$ contains an edge that is not in $\Gamma(\mathcal{W}, \mathcal{P})$. Since $\mathcal{Q}$ and $\mathcal{P}$ coincide on all bags prior to W (down to W_k) we must have $B_0 \subseteq B(H, \mathcal{Q})$.

If $\mathcal{Q}'_{j_0}$ is a twisting disturbance, then $j_1 > j_0$, in particular, W' comes before $W_{i_{j_0+1}}$ (there is at least one bag between $W_{i_{j_0}}$ and $W_{i_{j_0+1}}$, namely $W_{k_{j_0+1}}$). This means $\mathcal{Q}[W'] = \mathcal{P}[W']$ and hence we have

$$B_1 := B(G[W'], \mathcal{Q}[W']) = B(G[W'], \mathcal{P}[W']) = \Gamma(\mathcal{W}, \mathcal{P}).$$

On the other hand, the induced permutation of the restriction of $\mathcal{Q}$ to all bags prior to W' is π and thus $\pi^{-1}B_1 \subseteq B(H, \mathcal{Q})$. But π is not an automorphism of $\Gamma(\mathcal{W}, \mathcal{P})$ and therefore $\pi^{-1}B_1 = \pi^{-1}\Gamma(\mathcal{W}, \mathcal{P})$ contains an edge that is not in $\Gamma(\mathcal{W}, \mathcal{P})$ as desired. This concludes the proof of Claim 3.5.1 □

To exploit Claim 3.5.1 we now contract subsegments of $lw!$ consecutive inner bags of $\mathcal{W}$ into single bags. We assumed earlier that $(\mathcal{W}, \mathcal{P})$ is not stable so the number m of non-edges of $\Gamma(\mathcal{W}, \mathcal{P})$ is at least 1 (if $\Gamma(\mathcal{W}, \mathcal{P})$ is complete there can be no disturbances). Set $n' := zf(m-1)$. As $(\mathcal{W}, \mathcal{P})$ is valid, its length n is at least $f(m) = zlw!f(m-1) = n'lw!$. Let $(\mathcal{W}', \mathcal{P}')$ be the contraction of $(\mathcal{W}, \mathcal{P})$ along the sequence $i_1, \ldots, i_{n'}$ defined by $i_j := (j-1)lw! + 1$ for $j = 1, \ldots, n'$. Then by Lemma 3.6 the pair $(\mathcal{W}', \mathcal{P}')$ is a regular decomposition of G of length n', adhesion at most w, it is p-attached, and $\Gamma(\mathcal{W}', \mathcal{P}') = \Gamma(\mathcal{W}, \mathcal{P})$.

By construction every inner bag W'_i of $\mathcal{W}'$ consists of $lw!$ consecutive inner bags of $\mathcal{W}$ and hence by Claim 3.5.1 it contains a bridging disturbance $\mathcal{Q}'_i$ such $\mathcal{Q}'_i$ is p-attached in $G[W'_i]$, its induced permutation is the identity, and $B(G[W'_i], \mathcal{Q}'_i)$ is a proper supergraph of $\Gamma(\mathcal{W}', \mathcal{P}')$.

Clearly $\Gamma(\mathcal{W}', \mathcal{P}')$ has at most $z-1$ proper supergraphs on the same vertex set. On the other hand, $\mathcal{W}'$ has at least $n' - 1 = zf(m-1) - 1$

inner bags. By the pigeonhole principle there must be $f(m-1)$ indices $0 < i_1 < \ldots < i_{f(m-1)} < n'$ such that $B(G[W'_{i_j}], \mathcal{Q}'_{i_j})$ is the same graph Γ for $j = 1, \ldots, f(m-1)$.

Let $(\mathcal{W}'', \mathcal{P}'')$ be the contraction of $(\mathcal{W}', \mathcal{P}')$ along $i_1, \ldots, i_{f(m-1)}$. Obtain the foundational linkage $\mathcal{Q}''$ for $\mathcal{W}''$ from $\mathcal{P}''$ by replacing $\mathcal{P}'[W_{i_j}]$ with $\mathcal{Q}_{i_j}$ for $1 \leq j \leq f(m-1)$. By construction $\mathcal{W}''$ is a slim decomposition of G of length $f(m-1)$ and of adhesion at most w. $\mathcal{Q}''$ is a foundational linkage for $\mathcal{W}''$ that satisfies (L7) because $\mathcal{P}''$ does. By construction $\mathcal{Q}''$ is p-attached and $B(G[W''], \mathcal{Q}''[W'']) = \Gamma$ for all inner bags W'' of $\mathcal{W}''$. Hence $(\mathcal{W}'', \mathcal{P}'')$ is regular decomposition of G. But it is valid and its auxiliary graph Γ has more edges than $\Gamma(\mathcal{W}, \mathcal{P})$, contradicting our initial choice of $(\mathcal{W}, \mathcal{P})$. □

4 Token Movements

Consider the following token game. We place distinguishable tokens on the vertices of a graph H, at most one per vertex. A move consists of sliding a token along the edges of H to a new vertex without passing through vertices which are occupied by other tokens. Which placements of tokens can be obtained from each other by a sequence of moves?

A rather well-known instance of this problem is the 15-puzzle where tokens $1, \ldots, 15$ are placed on the 4-by-4 grid. It has been observed as early as 1879 by Johnson [6] that in this case there are two placements of the tokens which cannot be obtained from each other by any number of moves.

Clearly the problem gets easier the more "unoccupied" vertices there are. The hardest case with $|H| - 1$ tokens was tackled comprehensively by Wilson [22] in 1974 but before we turn to his solution we present a formal account of the token game and show how it helps with the linkage problem.

Throughout this section let H be a graph and let $\mathcal{X}$ always denote a sequence $\mathcal{X} = X_0, \ldots, X_n$ of vertex sets of H and $\mathcal{M}$ a non-empty sequence $\mathcal{M} = M_1, \ldots, M_n$ of non-trivial paths in H. In our model the sets X_i are "occupied vertices", the paths M_i are paths along which the tokens are moved, and i is the "move count".

Formally, a pair $(\mathcal{X}, \mathcal{M})$ is called a *movement on H* if for $i = 1, \ldots, n$

(M1) the set $X_{i-1} \triangle X_i$ contains precisely the two end vertices of M_i, and

(M2) M_i is disjoint from $X_{i-1} \cap X_i$.

Then n is the *length* of $(\mathcal{X}, \mathcal{M})$, the sets in $\mathcal{X}$ are its *intermediate configurations*, in particular, X_0 and X_n are its *first* and *last configuration*, respectively.

The paths in $\mathcal{M}$ are the *moves* of $(\mathcal{X}, \mathcal{M})$. A movement with first configuration X and last configuration Y is called an X–Y *movement*. Note that our formal notion of token movements allows a move M_i to have both ends in X_{i-1} or both in X_i. In our intuitive account of the token game this corresponds to "destroying" or "creating" a pair of tokens on the end vertices of M_i.

Let us state some obvious facts about movements. If $\mathcal{M}$ is a non-empty sequence of non-trivial paths in H and one intermediate configuration X_i is given, then there is a unique sequence $\mathcal{X}$ such that $(\mathcal{X}, \mathcal{M})$ satisfies (M1). A pair $(\mathcal{X}, \mathcal{M})$ is a movement if and only if $((X_{i-1}, X_i), (M_i))$ is a movement for $i = 1, \ldots, n$. This easily implies the following Lemma so we spare the proof.

Lemma 4.1. *Let $(\mathcal{X}, \mathcal{M}) = ((X_0, \ldots, X_n), (M_1, \ldots, M_n))$ and $(\mathcal{Y}, \mathcal{N}) = ((Y_0, \ldots, Y_m), (N_1, \ldots, N_m))$ be movements on H and let $Z \subseteq V(H)$.*

(i) If $X_n = Y_0$, then the pair

$$\big((X_0, \ldots, X_n = Y_0, \ldots, Y_m), (M_1, \ldots M_n, N_1, \ldots, N_m)\big)$$

is a movement. We denote it by $(\mathcal{X}, \mathcal{M}) \oplus (\mathcal{Y}, \mathcal{N})$ and call it the concatenation *of $(\mathcal{X}, \mathcal{M})$ and $(\mathcal{Y}, \mathcal{N})$.*

(ii) If every move of $\mathcal{M}$ is disjoint from Z, then the pair

$$\big((X_0 \cup Z, \ldots, X_n \cup Z), (M_1, \ldots M_n)\big)$$

is a movement and we denote it by $(\mathcal{X} \cup Z, \mathcal{M})$.

Let $(\mathcal{X}, \mathcal{M})$ be a movement. For $i = 1, \ldots, n$ let R_i be the graph with vertex set $(X_{i-1} \times \{i-1\}) \cup (X_i \times \{i\})$ and the following edges:

1. $(x, i-1)(x, i)$ for each $x \in X_{i-1} \cap X_i$, and
2. $(x, j)(y, k)$ where x, y are the end vertices of M_i and j, k the unique indices such that $(x, j), (y, k) \in V(R_i)$.

Define a multigraph $\mathcal{R}$ with vertex set $\bigcup_{i=0}^{n}(X_i \times \{i\})$ where the multiplicity of an edge is the number of graphs R_i containing it. Observe that two graphs R_i and R_j with $i < j$ are edge-disjoint unless $j = i + 1$ and M_i and M_j both end in the same two vertices x, y of X_j, in which case they share one edge, namely $(x, j)(y, j)$. Our reason to prefer the above definition of $\mathcal{R}$ over just taking the simple graph $\bigcup_{i=1}^{n} R_i$ is to avoid a special case in the following argument.

Every graph R_i is 1-regular. Hence in $\mathcal{R}$ every vertex (x, i) with $0 < i < n$ has degree 2 as (x, i) is a vertex of R_j if an only if $j = i$ or $j = i + 1$. Every vertex (x, i) with $i = 0$ or $i = n$ has degree 1 as it only lies in R_1 or in R_n. This implies that a component of $\mathcal{R}$ is either a cycle (possibly of length 2) avoiding $(X_0 \times \{0\}) \cup (X_n \times \{n\})$ or a non-trivial path with both end vertices in $(X_0 \times \{0\}) \cup (X_n \times \{n\})$. We denote the subgraph of $\mathcal{R}$ consisting of these paths by $\mathcal{R}(\mathcal{X}, \mathcal{M})$. Intuitively, each path of $\mathcal{R}(\mathcal{X}, \mathcal{M})$ traces the position of one token over the course of the token movement or of one pair of tokens which is destroyed or created during the movement.

For vertex sets X and Y we call any 1-regular graph on $(X \times \{0\}) \cup (Y \times \{\infty\})$ an (X, Y)-*pairing*. An (X, Y)-pairing is said to be *balanced* if its edges form a perfect matching from $X \times \{0\}$ to $Y \times \{\infty\}$, that is, each edge has one end vertex in $X \times \{0\}$ and the other in $Y \times \{\infty\}$.

The components of $\mathcal{R}(\mathcal{X}, \mathcal{M})$ induce a 1-regular graph on $(X_0 \times \{0\}) \cup (X_n \times \{n\})$ where two vertices form an edge if and only if they are in the same component of $\mathcal{R}(\mathcal{X}, \mathcal{M})$. To make this formally independent of the index n, we replace each vertex (x, n) by (x, ∞). The obtained graph L is an (X_0, X_n)-pairing and we call it the *induced pairing* of the movement $(\mathcal{X}, \mathcal{M})$. A movement $(\mathcal{X}, \mathcal{M})$ with induced pairing L is called an L-*movement*. If a movement induces a balanced pairing, then we call the movement *balanced* as well.

Given two sets X and Y and a bijection $\varphi : X \to Y$ we denote by $L(\varphi)$ the balanced X–Y pairing where $(x, 0)(y, \infty)$ is an edge of $L(\varphi)$ if and only if $y = \varphi(x)$. Clearly an X–Y pairing L is balanced if and only if there is a bijection $\varphi : X \to Y$ with $L = L(\varphi)$.

Given sets X, Y, and Z let L_X be an X–Y pairing and L_Z a Y–Z pairing. Denote by $L_X \oplus L_Z$ the graph on $(X \times \{0\}) \cup (Z \times \{\infty\})$ where two vertices are connected by an edge if and only if they lie in the same component of $L_X \cup L(\mathrm{id}_Y) \cup L_Z$. The components of $L_X \cup L(\mathrm{id}_Y) \cup L_Z$ are either paths with both ends in $(X \times \{0\}) \cup (Z \times \{\infty\})$ or cycles avoiding that set. So $L_X \oplus L_Z$ is an X–Z pairing end we call it the *concatenation* of L_X and L_Z. The next Lemma is an obvious consequence of this construction (and Lemma 4.1 (i)).

Lemma 4.2. *The induced pairing of the concatenation of two movements is the concatenation of their induced pairings.*

Let $(\mathcal{X}, \mathcal{M})$ be a movement on H. A vertex x of H is called $(\mathcal{X}, \mathcal{M})$-*singular* if no move of $\mathcal{M}$ contains x as an inner vertex and $I_x := \{i \mid x \in X_i\}$ is an *integer interval*, that is, a possibly empty sequence of consecutive integers. Furthermore, x is called *strongly* $(\mathcal{X}, \mathcal{M})$-*singular* if it is $(\mathcal{X}, \mathcal{M})$-singular and I_x is empty or contains one of 0 and n where n denotes the length of $(\mathcal{X}, \mathcal{M})$.

We say that a set $W \subseteq V(H)$ is $(\mathcal{X}, \mathcal{M})$*-singular* or *strongly* $(\mathcal{X}, \mathcal{M})$*-singular* if all its vertices are. If the referred movement is clear from the context, then we shall drop it from the notation and just write singular or strongly singular.

Note that any vertex v of H that is contained in at most one move of $\mathcal{M}$ is strongly $(\mathcal{X}, \mathcal{M})$-singular. Furthermore, v is singular but not strongly singular if it is contained in precisely two moves but neither in the first nor in the last configuration.

The following Lemma shows how to obtain linkages in a graph G from movements on the auxiliary graph of a regular decomposition of G. It enables us to apply the results about token movements from this section to our linkage problem.

Lemma 4.3. *Let* $(\mathcal{W}, \mathcal{P})$ *be a stable regular decomposition of some graph* G *and set* $\lambda := \{\alpha \mid P_\alpha$ *is non-trivial*$\}$ *and* $\theta := \{\alpha \mid P_\alpha$ *is trivial*$\}$*. Let* $(\mathcal{X}, \mathcal{M})$ *be a movement of length* n *on a subgraph* $\Gamma \subseteq \Gamma(\mathcal{W}, \mathcal{P})$ *and denote its induced pairing by* L*. If* θ *is* $(\mathcal{X}, \mathcal{M})$*-singular and* W_a *and* W_b *are inner bags of* $\mathcal{W}$ *with* $b - a = 2n - 1$*, then there is a linkage* $\mathcal{Q} \subseteq G^{\mathcal{P}}_{\Gamma}[W_{[a,b]}]$ *and a bijection* $\varphi : E(L) \to \mathcal{Q}$ *such that for each* $e \in E(L)$ *the path* $\varphi(e)$ *ends in the left* α*-vertex of* W_a *if and only if* $(\alpha, 0) \in e$ *and* $\varphi(e)$ *ends in the right* α*-vertex of* W_b *if and only if* $(\alpha, \infty) \in e$*.*

Proof. Let us start with the general observation that for every connected subgraph $\Gamma_0 \subseteq \Gamma(\mathcal{W}, \mathcal{P})$ and every inner bag W of $\mathcal{W}$ the graph $G^{\mathcal{P}}_{\Gamma_0}[W]$ is connected: If $\alpha\beta$ is an edge of Γ_0, then some inner bag of $\mathcal{W}$ contains a $\mathcal{P}$-bridge realising $\alpha\beta$ and so does W by (L8). In particular, $G^{\mathcal{P}}_{\Gamma_0}[W]$, contains a P_α–P_β path. So $\mathcal{P}_{V(\Gamma_0)}[W]$ must be contained in one component of $G^{\mathcal{P}}_{\Gamma_0}[W]$ as Γ_0 is connected. But any vertex of $G^{\mathcal{P}}_{\Gamma_0}[W]$ is in $\mathcal{P}_{V(\Gamma_0)}$ or in a $\mathcal{P}$-bridge attaching to it. Therefore $G^{\mathcal{P}}_{\Gamma_0}[W]$ is connected.

The proof is by induction on n. Denote the end vertices of M_1 by α and β, that is, $X_0 \triangle X_1 = \{\alpha, \beta\}$. By definition the induced pairing L_1 of $((X_0, X_1), (M_1))$ contains the edges $(\gamma, 0)(\gamma, \infty)$ with $\gamma \in X_0 \cap X_1$ and w.l.o.g. precisely one of $(\alpha, 0)(\beta, 0)$, $(\alpha, 0)(\beta, \infty)$, and $(\alpha, \infty)(\beta, \infty)$. The above observation implies that $G^{\mathcal{P}}_{M_1}[W_a]$ is connected. Hence $P_\alpha[W_{[a,a+1]}] \cup G^{\mathcal{P}}_{M_1}[W_a] \cup P_\beta[W_{[a,a+1]}]$ is connected and thus contains a path Q such that $\mathcal{Q}_1 := \{Q\} \cup \mathcal{P}_{X_0 \cap X_1}[W_{[a,a+1]}]$ satisfies the following. There is a bijection $\varphi_1 : E(L_1) \to \mathcal{Q}_1$ such that for each $e \in E(L_1)$ the path $\varphi_1(e)$ ends in the left γ-vertex of W_a if and only if $(\gamma, 0) \in e$ and $\varphi_1(e)$ ends in the right γ-vertex of W_{a+1} if and only if $(\gamma, \infty) \in e$. Moreover, the paths of $\mathcal{Q}_1$ are internally disjoint from $W_{a+1} \cap W_{a+2}$.

In the base case $n = 1$ the linkage $\mathcal{Q} := \mathcal{Q}_1$ is as desired. Suppose that $n \geq 2$. Then $((X_1, \ldots, X_n), (M_2, \ldots, M_n))$ is a movement and we denote its induced permutation by L_2. Lemma 4.2 implies $L = L_1 \oplus L_2$. By induction there is a linkage $\mathcal{Q}_2 \subseteq G^{\mathcal{P}}_{\Gamma}[W_{[a+2,b]}]$ and a bijection $\varphi_2 : E(L_2) \to \mathcal{Q}_2$ such that for any $e \in E(L_2)$ the path $\varphi_2(e)$ ends in the left α-vertex of W_{a+2} (which is the right α-vertex of W_{a+1}) if and only if $(\alpha, 0) \in e$ and in the right α-vertex of W_b if and only if $(\alpha, \infty) \in e$.

Clearly for every $\gamma \in X_1$ the γ-vertex of $W_{a+1} \cap W_{a+2}$ has degree at most 1 in $\mathcal{Q}_1$ and in $\mathcal{Q}_2$. If a path of $\mathcal{Q}_1$ contains the γ-vertex of $W_{a+1} \cap W_{a+2}$ and $\gamma \notin X_1$, then $\gamma \in \theta$ so by assumption $I_\gamma = \{i \mid \gamma \in X_i\}$ is an integer interval which contains 0 but not 1. This means that no path of $\mathcal{Q}_2$ contains the unique vertex of P_γ. If the union $\mathcal{Q}_1 \cup \mathcal{Q}_2$ of the two graphs $\mathcal{Q}_1$ and $\mathcal{Q}_2$ contains no cycle, then it is a linkage $\mathcal{Q}$ as desired. Otherwise it only contains such a linkage. □

In the rest of this section we shall construct suitable movements as input for Lemma 4.3. Our first tool to this end is the following powerful theorem[4] of Wilson.

Theorem 4.4 (Wilson 1974)**.** *Let k be a postive integer and let H be a graph on $n \geq k+1$ vertices. If H is 2-connected and contains a triangle, then for every bijection $\varphi : X \to Y$ of sets $X, Y \subseteq V(H)$ with $|X| = k = |Y|$ there is a $L(\varphi)$-movement of length $m \leq n!/(n-k)!$ on H.*

The given bound on m is not included in the original statement but not too hard to check: Suppose that $(\mathcal{X}, \mathcal{M})$ is a shortest $L(\varphi)$-movement and m is its length. Since L is balanced we may assume that no tokes are "created" or "destroyed" during the movement, that is, all intermediate configurations have the same size and for every i with $1 \leq i \leq m$ there is an injection $\varphi_i : X \to V(H)$ such that the induced pairing of $((X_0, \ldots, X_i), (M_1, \ldots, M_i))$ is $L(\varphi_i)$. If there were $i < j$ with $\varphi_i = \varphi_j$, then

$$\big((X_0, \ldots, X_i = X_j, X_{j+1}, \ldots, X_m), (M_1, \ldots, M_i, M_{j+1}, \ldots, M_m)\big)$$

was an $L(\varphi)$-movement of length $m - j + i < m$ contradicting our choice of $(\mathcal{X}, \mathcal{M})$. But there are at most $n!/(n-k)!$ injections from X to $V(H)$ so we must have $m \leq n!/(n-k)!$.

For our application we need to generate L-movements where L is not necessarily balanced. Furthermore, Lemma 4.3 requires the vertices of θ to be singular with respect to the generated movement. Lemma 4.8 and Lemma 4.9

[4]Wilson stated his theorem for graphs which are neither bipartite, nor a cycle, nor a certain graph θ_0. If H properly contains a triangle, then it satisfies all these conditions and if H itself is a triangle, then our theorem is obviously true.

give a direct construction of movements if some subgraph of H is a large star. Lemma 4.10 provides an interface to Theorem 4.4 that incorporates the above requirements. The proofs of these three Lemmas require a few tools: Lemma 4.5 simply states that for sets X and Y of equal size there is a short balanced X–Y movement. Lemma 4.6 exploits this to show that instead of generating movements for every choice of $X, Y \subseteq V(H)$ and any (X, Y)-pairing L it suffices to consider just one choice of X and Y. Lemma 4.7 allows us to move strongly singular vertices from X to Y and vice versa without spoiling the existence of the desired X–Y movement.

We call a set A of vertices in a graph H *marginal* if $H - A$ is connected and every vertex of A has a neighbour in $H - A$.

Lemma 4.5. *For any two distinct vertex sets X and Y of some size k in a connected graph H and any marginal set $A \subseteq V(H)$ there is a balanced X–Y movement of length at most k on H such that A is strongly singular.*

Proof. We may assume that H is a tree and that all vertices of A are leaves of this tree. This already implies that vertices of A cannot be inner vertices of moves. Moreover, we may assume that $X \cap Y \cap A = \emptyset$.

We apply induction on $|H|$. The base case $|H| = 1$ is trivial. For $|H| > 1$ let e be an edge of H. If the two components H_1 and H_2 of $H - e$ each contain the same number of vertices from X as from Y, then for $i = 1, 2$ we set $X_i := X \cap V(H_i)$ and $Y_i := Y \cap V(H_i)$. By induction there is a balanced X_i–Y_i movement $(\mathcal{X}_i, \mathcal{M}_i)$ of length at most $|X_i|$ on H_i such that each vertex of A is strongly $(\mathcal{X}_i, \mathcal{M}_i)$-singular where $i = 1, 2$. By Lemma 4.1 $(\mathcal{X}, \mathcal{M}) := (\mathcal{X}_1 \cup X_2, \mathcal{M}_1) \oplus (\mathcal{X}_2 \cup Y_1, \mathcal{M}_2)$ is an X–Y movement of length at most $|X_1| + |X_2| = |X| = k$ as desired. Clearly $(\mathcal{X}, \mathcal{M})$ is balanced and A is strongly $(\mathcal{X}, \mathcal{M})$-singular as H_1 and H_2 are disjoint.

So we may assume that for every edge e of H one component of $H - e$ contains more vertices from Y than from X and direct e towards its end vertex lying in this component. As every directed tree has a sink, there is a vertex y of H such that every incident edge e is incoming, that is, the component of $H - e$ not containing y contains more vertices of X than of Y. As $|X| = |Y|$, this can only be if y is a leaf in H and $y \in Y \setminus X$.

Let M be any X–y path and denote its first vertex by x. At most one of $x \in Y$ and $x \in A$ can be true by assumption. Clearly $((\{x\}, \{y\}), (M))$ is an $\{x\}$–$\{y\}$ movement and since $H - y$ is connected, by induction there is a balanced $(X \setminus \{x\})$–$(Y \setminus \{y\})$ movement $(\mathcal{X}', \mathcal{M}')$ of length at most $k - 1$ on $H - y$ such that A is strongly singular w.r.t. both movements. As before, Lemma 4.1 implies that

$$(\mathcal{X}, \mathcal{M}) := \big((X, (X \setminus \{x\}) \cup \{y\}), (M)\big) \oplus (\mathcal{X}' \cup \{y\}, \mathcal{M}')$$

is an X–Y movement of length at most k. Clearly $(\mathcal{X}, \mathcal{M})$ is balanced and by construction A is strongly $(\mathcal{X}, \mathcal{M})$-singular. □

Lemma 4.6. *Let k be a positive integer and H a connected graph with a marginal set A. Suppose that $X, X', Y', Y \subseteq V(H)$ are sets with $|X| + |Y| = 2k$, $|X'| = |X|$, and $|Y'| = |Y|$ such that $(X \cup X') \cap (Y' \cup Y)$ does not intersect A. If for each (X', Y')-pairing L' there is an L'-movement $(\mathcal{X}', \mathcal{M}')$ of length at most n' on H such that A is strongly $(\mathcal{X}', \mathcal{M}')$-singular, then for each (X, Y)-pairing L there is an L-movement $(\mathcal{X}, \mathcal{M})$ of length at most $n' + 2k$ such that A is $(\mathcal{X}, \mathcal{M})$-singular and all vertices of A that are not strongly $(\mathcal{X}, \mathcal{M})$-singular are in $(X' \cup Y') \setminus (X \cup Y)$.*

Proof. Let $(\mathcal{X}_X, \mathcal{M}_X)$ be a balanced X–X' movement of length at most $|X|$ and let $(\mathcal{X}_Y, \mathcal{M}_Y)$ be a balanced Y'–Y movement of length at most $|Y|$ such that A is strongly singular w.r.t. both movements. These exist by Lemma 4.5. For any X'–Y' movement $(\mathcal{X}', \mathcal{M}')$ such that A is strongly $(\mathcal{X}', \mathcal{M}')$-singular,

$$(\mathcal{X}, \mathcal{M}) := (\mathcal{X}_X, \mathcal{M}_X) \oplus (\mathcal{X}', \mathcal{M}') \oplus (\mathcal{X}_Y, \mathcal{M}_Y)$$

is a movement of length at most $|X| + n' + |Y| = n' + 2k$ by Lemma 4.1.

In a slight abuse of the notation we shall write $a \in \mathcal{M}_X$, $a \in \mathcal{M}'$, and $a \in \mathcal{M}_Y$ for a vertex $a \in A$ if there is a move of $\mathcal{M}_X$, $\mathcal{M}'$, and $\mathcal{M}_Y$, respectively, that contains a. Consequently, we write $a \notin M_X$, etc. if there is no such move. The set A is strongly singular w.r.t. each of $(\mathcal{X}_X, \mathcal{M}_X)$, $(\mathcal{X}', \mathcal{M}')$, and $(\mathcal{X}_Y, \mathcal{M}_Y)$. Therefore all moves of $\mathcal{M}$ are internally disjoint from A and each $a \in A$ is contained in at most one move from each of $\mathcal{M}_X$, $\mathcal{M}'$, and $\mathcal{M}_Y$. Moreover, for each $a \in A$

1. $a \in \mathcal{M}_X$ if and only if precisely one of $a \in X$ and $a \in X'$ is true,
2. $a \in \mathcal{M}'$ if and only if precisely one of $a \in X'$ and $a \in Y'$ is true, and
3. $a \in \mathcal{M}_Y$ if and only if precisely one of $a \in Y'$ and $a \in Y$ is true.

Clearly $A \setminus (X \cup X' \cup Y' \cup Y)$ is strongly $(\mathcal{X}, \mathcal{M})$-singular as none of its vertices is contained in a path of $\mathcal{M}$.

Let $a \in X \cap A$. Then by assumption $a \notin Y \cup Y'$ and thus $a \notin \mathcal{M}_Y$. If $a \in X'$, then $a \in \mathcal{M}'$ and $a \notin \mathcal{M}_X$. Otherwise $a \notin X'$ and therefore $a \in \mathcal{M}_X$ and $a \notin \mathcal{M}'$. In either case a is in at most one move of $\mathcal{M}$ and hence $X \cap A$ is strongly $(\mathcal{X}, \mathcal{M})$-singular. A symmetric argument shows that $Y \cap A$ is strongly $(\mathcal{X}, \mathcal{M})$-singular.

Let $a \in (X' \cup Y') \cap A$ with $a \notin X \cup Y$. Then $a \in X' \triangle Y'$ so $a \in \mathcal{M}'$ and precisely one of $a \in \mathcal{M}_X$ and $a \in \mathcal{M}_Y$ is true.

We conclude that every vertex of $a \in A$ is $(\mathcal{X}, \mathcal{M})$-singular and it is even strongly $(\mathcal{X}, \mathcal{M})$-singular if and only if $a \notin (X' \cup Y') \setminus (X \cup Y)$.

The induced pairings L_X of $(\mathcal{X}_X, \mathcal{M}_X)$ and L_Y of $(\mathcal{X}_Y, \mathcal{M}_Y)$ are both balanced and it is not hard to see that for a suitable choice of L' the induced pairing $L_X \oplus L' \oplus L_Y$ of $(\mathcal{X}, \mathcal{M})$ equals L. □

Lemma 4.7. *Let H be a connected graph and let $X, Y \subseteq V(H)$. Suppose that L is an (X, Y)-pairing and $(\mathcal{X}, \mathcal{M})$ an L-movement of length n. If $x \in X \cup Y$ is strongly $(\mathcal{X}, \mathcal{M})$-singular, then the following statements hold.*

(i) *$(\mathcal{X} \triangle x, \mathcal{M})$ is an $(L \triangle x)$-movement of length n where $\mathcal{X} \triangle x := (X_0 \triangle \{x\}, \dots, X_n \triangle \{x\})$ and $L \triangle x$ denotes the graph obtained from L by replacing $(x, 0)$ with (x, ∞) or vice versa (at most one of these can be a vertex of L).*

(ii) *A vertex $y \in V(H)$ is (strongly) $(\mathcal{X}, \mathcal{M})$-singular if and only if it is (strongly) $(\mathcal{X} \triangle x, \mathcal{M})$-singular.*

Proof. Clearly $(\mathcal{X} \triangle x, \mathcal{M})$ is an $(L \triangle x)$-movement of length n. As its intermediate configurations differ from those of $\mathcal{X}$ only in x, the last assertion is trivial for $y \neq x$. For $y = x$ note that $\{i \mid x \notin X_i\}$ is an integer interval containing precisely one of 0 and n because $\{i \mid x \in X_i\}$ is. □

In the final three Lemmas of this section we put our tools to use and construct movements under certain assumptions about the graph. Note that it is not hard to improve on the upper bounds given for the lengths of the generated movements with more complex proofs. However, in our main proof we have an arbitrarily long stable regular decomposition at our disposal, so the input movements for Lemma 4.3 can be arbitrarily long as well.

Lemma 4.8. *Let k be a positive integer and H a connected graph with a marginal set A. If one of*

a) $|A| \geq 2k - 1$ *and*

b) $|N_H(v) \cap N_H(w) \cap A| \geq 2k - 3$ *for some edge vw of $H - A$*

holds, then for any X–Y pairing L such that $X, Y \subseteq V(H)$ with $|X| + |Y| = 2k$ and $X \cap Y \cap A = \emptyset$ there is an L-movement $(\mathcal{X}, \mathcal{M})$ of length at most $3k$ on H such that A is $(\mathcal{X}, \mathcal{M})$-singular.

The basic argument of the proof is that that if we place tokens on the leaves of a star but not on its centre, then we can clearly "destroy" any given pair of tokens by moving one on top of the other through the centre of the star.

Proof. Suppose that a) holds. Let $N_A \subseteq A$ with $|N_A| = 2k - 1$. There are sets $X', Y' \subseteq V(H)$ such that

1. $|X'| = |X|$ and $|Y'| = |Y|$,
2. $X \cap N_A \subseteq X'$,
3. $Y \cap N_A \subseteq Y'$, and
4. $N_A \subseteq X' \cup Y'$ and $X' \cap Y' \cap A = \emptyset$.

By Lemma 4.6 it suffices to show that for each X'–Y' pairing L' there is an L'-movement $(\mathcal{X}', \mathcal{M}')$ of length at most k on H such that A is strongly $(\mathcal{X}', \mathcal{M}')$-singular. Assume w.l.o.g. that the unique vertex of $(X' \cup Y') \setminus N_A$ is in X'. Repeated application of Lemma 4.7 implies that the desired L'-movement $(\mathcal{X}', \mathcal{M}')$ exists if and only if for every $(X' \cup Y')$–$\emptyset$ pairing L'' there is an L''-movement $(\mathcal{X}'', \mathcal{M}'')$ of length at most k on H such that A is strongly $(\mathcal{X}'', \mathcal{M}'')$-singular.

Let L'' be any $(X' \cup Y')$–$\emptyset$ pairing. Then $E(L'') = \{(x_i, 0)(y_i, 0) \mid i = 1, \ldots, k\}$ where $(X' \cup Y') \cap N_A = \{x_1, \ldots, x_k, y_2, \ldots, y_k\}$ and $(X' \cup Y') \setminus N_A = \{y_1\}$. For $i = 0, \ldots, k$ set $X_i := \{x_j, y_j \mid j > i\}$. For $i = 1, \ldots, k$ let M_i be an x_i–y_i path in H that is internally disjoint from A. Then $(\mathcal{X}'', \mathcal{M}'') := ((X_0, \ldots, X_k), (M_1, \ldots, M_k))$ is an L''-movement of length k and obviously A is strongly $(\mathcal{X}'', \mathcal{M}'')$-singular.

Suppose that b) holds and let $N_A \subseteq N_H(v) \cap N_H(w) \cap A$ with $|N_A| = 2k-3$ and set $N_B := \{v, w\}$. There are sets $X', Y' \subseteq V(H)$ such that

1. $|X'| = |X|$ and $|Y'| = |Y|$,
2. $X \cap N_A \subseteq X'$ and $X' \cap A \subseteq X \cup N_A$,
3. $Y \cap N_A \subseteq Y'$ and $Y' \cap A \subseteq Y \cup N_A$,
4. $N_A \subseteq X' \cup Y'$ and $X' \cap Y' \cap A = \emptyset$,
5. $N_B \subseteq X'$ or $X' \subset N_A \cup N_B$, and
6. $N_B \subseteq Y'$ or $Y' \subset N_A \cup N_B$.

By Lemma 4.6 (see case a) for the details) it suffices to find an L'-movement $(\mathcal{X}', \mathcal{M}')$ of length at most k on H such that A is strongly $(\mathcal{X}', \mathcal{M}')$-singular where L' is any X'–Y' pairing. Since $|(X' \cup Y') \setminus N_A| = 3$ we may asssume w.l.o.g. that $N_B \subseteq X'$ and $Y' \subseteq N_A \cup \{v\}$. So either there is $z \in X' \setminus (N_A \cup N_B)$ or $v \in Y'$. By repeated application of Lemma 4.7 we may assume that $N_A \subseteq X'$. This means that L' has the vertices $\bar{N}_A := N_A \times \{0\}$, $\bar{v} := (v,0)$, $\bar{w} := (w,0)$, and $\bar{z} := (z,0)$ in the first case or $\bar{z} := (v,\infty)$ in the second case. So L' must satisfy one of the following.

1. No edge of L' has both ends in $\{\bar{v}, \bar{w}, \bar{z}\}$.
2. $\bar{v}\bar{w} \in E(L')$.
3. $\bar{v}\bar{z} \in E(L')$.
4. $\bar{w}\bar{z} \in E(L')$.

This leaves us with eight cases in total. Since construction is almost the same for all cases we provide the details for only one of them: We assume that $v \in Y'$ and $\bar{w}\bar{z} \in E(L')$. Then L' has edges $(w,0)(v,\infty)$ and $\{(x_i,0)(y_i,0) \mid i = 1,\ldots,k-1\}$ where $x_1 := v$ and $X' \cap N_A = \{x_2,\ldots,x_{k-1},y_1,\ldots,y_{k-1}\}$. For $i = 0,\ldots,k-1$ set $X_i := \{w\} \cup \bigcup_{j>i}\{x_j,y_j\}$ and let $X_k := \{v\}$. Set $M_1 := vy_1$ and $M_i := x_i v y_i$ for $i = 2,\ldots,k-1$ and let M_k be a w–z path in H that is internally disjoint from A. Then $(\mathcal{X}',\mathcal{M}') := ((X_0,\ldots,X_k),(M_1,\ldots,M_k))$ is an L'-movement and A is strongly $(\mathcal{X}',\mathcal{M}')$-singular. □

Lemma 4.9. *Let k be a positive integer and H a connected graph with a marginal set A. Let $X, Y \subseteq V(H)$ with $|X| + |Y| = 2k$ and $X \cap Y \cap A = \emptyset$. Suppose that there is a vertex v of $H - (X \cup Y \cup A)$ such that*

$$2|N_H(v) \setminus A| + |N_H(v) \cap A| \geq 2k + 1.$$

Then for any (X,Y)-pairing L there is an L-movement of length at most $k(k+2)$ on H such that A is singular.

Although the basic idea is still the same as in Lemma 4.8 it gets a little more complicated here as our star might not have enough leaves to hold all tokens at the same time. Hence we prefer an inductive argument over an explicit construction.

Proof. Set $N_A := N_H(v) \cap A$ and $N_B := N_H(v) \setminus A$. If $|N_A| \geq 2k-1$, then we are done by Lemma 4.8 as $3k \leq k(k+2)$. So we may assume that $|N_A| \leq 2k-2$. Under this additional assumption we prove a slightly stronger statement than that of Lemma 4.9 by induction on k: We not only require that A is singular but also that all vertices of A that are not strongly singular are in $N_A \setminus (X \cup Y)$.

The base case $k = 1$ is trivial. Suppose that $k \geq 2$. There are sets $X', Y' \subseteq V(H)$ such that

1. $|X'| = |X|$ and $|Y'| = |Y|$,
2. $X \cap N_A \subseteq X'$ and $X' \cap A \subseteq X \cup N_A$,
3. $Y \cap N_A \subseteq Y'$ and $Y' \cap A \subseteq Y \cup N_A$,
4. $N_A \subseteq X' \cup Y'$ and $X' \cap Y' \cap A = \emptyset$,
5. $N_B \subseteq X'$ or $X' \subset N_A \cup N_B$,
6. $N_B \subseteq Y'$ or $Y' \subset N_A \cup N_B$, and
7. $v \notin X'$ and $v \notin Y'$.

By Lemma 4.6 it suffices to find an L'-movement $(\mathcal{X}', \mathcal{M}')$ of length at most k^2 such that A is strongly $(\mathcal{X}', \mathcal{M}')$-singular where L' is any X'–Y' pairing.

If there are $x, y \in X' \cap N_H(v)$ such that $(x, 0)(y, 0) \in E(L')$, then set $X'' := X' \setminus \{x, y\}$, $Y'' := Y'$, $H'' := H - (A \setminus (X'' \cup Y''))$, and $L'' := L' - \{(x, 0), (y, 0)\}$. We have $N_{H''}(v) \setminus A = N_B$ and $N_{H''}(v) \cap A = N_A \setminus \{x, y\}$ as $N_A \subseteq X' \cup Y'$. This means

$$2|N_{H''}(v) \setminus A| + |N_{H''}(v) \cap A| \geq 2|N_B| + |N_A| - 2 \geq 2k - 1.$$

Hence by induction there is an L''-movement $(\mathcal{X}'', \mathcal{M}'')$ of length at most $(k + 1)(k - 1)$ on H'' such that A is singular and all vertices of A that are not strongly singular are in $N_A \setminus (X'' \cup Y'')$. Since $N_A \cap V(H'') \subseteq X'' \cup Y''$ the set A is strongly $(\mathcal{X}'', \mathcal{M}'')$-singular. Then by construction $(\mathcal{X}', \mathcal{M}') := ((X', X''), (xvy)) \oplus (\mathcal{X}'', \mathcal{M}'')$ is an L'-movement of length at most k^2 and A is strongly $(\mathcal{X}', \mathcal{M}')$-singular.

The case $x, y \in Y' \cap N_H(v)$ with $(x, \infty)(y, \infty) \in E(L')$ is symmetric. If there are $x \in X' \cap N_H(v)$ and $y \in Y' \cap N_H(v)$ such that $(x, 0)(y, \infty) \in E(L')$ and at least one of x and y is in N_A, then the desired movement exists by Lemma 4.7 and one of the previous cases.

By assumption

$$2|N_H(v)| \geq 2|N_B| + |N_A| \geq 2k + 1$$

and thus $|N_H(v)| \geq k+1$. If $N_B \subseteq X'$, then $N_H(v) \subseteq X' \cup (Y' \cap A)$ and there is a pair as above by the pigeon hole principle. Hence we may assume that $X' \subset N_B$ and by symmetry also that $Y' \subset N_B$. This implies that $N_A = \emptyset$ and that L' is balanced.

So we have $|X'| = k = |Y'|$, $X', Y' \subseteq N_B$ and $|N_B| \geq k + 1$. It is easy to see that there is an L'-movement $(\mathcal{X}', \mathcal{M}')$ of length at most $2k \leq k^2$ on $H[\{v\} \cup N_B]$ such that A is strongly $(\mathcal{X}', \mathcal{M}')$-singular. □

Lemma 4.10. *Let $n \in \mathbb{N}$ and let $f : \mathbb{N}_0 \to \mathbb{N}_0$ be the map that is recursively defined by setting $f(0) := 0$ and $f(k) := 2k + 2n! + 4 + f(k-1)$ for $k > 0$. Let k be a positive integer and let H be a connected graph on at most n vertices with a marginal set A. Let $X, Y \subseteq V(H)$ with $|X| + |Y| = 2k$ and $X \cap Y \cap A = \emptyset$ such that neither X nor Y contains all vertices of $H - A$. Suppose that there is a block D of $H - A$ such that D contains a triangle and $2|D| + |N(D)| \geq 2k + 3$. Then for any (X, Y)-pairing L there is an L-movement of length at most $f(k)$ on H such that A is singular.*

Proof. Set $N_A := N(D) \cap A$ and $N_B := N(D) \setminus A$. If $|N_A| \geq 2k - 1$, then we are done by Lemma 4.8 as $3k \leq f(k)$. So we may assume that $|N_A| \leq 2k - 2$. Under this additional assumption we prove a slightly stronger statement than that of Lemma 4.10 by induction on k: We not only require that A is singular but also that all vertices of A that are not strongly singular are in $N_A \setminus (X \cup Y)$. As always the base case $k = 1$ is trivial. Suppose that $k \geq 2$.

Claim 4.10.1. *Suppose that $|V(D) \setminus X| \geq 1$ and that there is an edge $(x, 0)(y, 0) \in E(L)$ with $x \in V(D)$ and $y \in V(D) \cup N(D)$. Let $A' \subseteq A \setminus (X \cup Y)$ with $|A'| \leq 1$. Then there is an L-movement of length at most $|D|! + 1 + f(k-1)$ on $H - A'$ such that A is singular and every vertex of A that is not strongly singular is in $N_A \setminus (X \cup Y)$.*

Proof. Let y' be a neighbour of y in D. Here is a sketch of the idea: Move the token from x to y' by a movement on D which we can generate with Wilsons's Theorem 4.4 and then add the move yy'. This "destroys" one pair of tokens and allows us to invoke induction.

We assume $x \neq y'$ (in the case $x = y'$ we can skip the construction of $(\mathcal{X}_\varphi, \mathcal{M}_\varphi)$ in this paragraph). Set $X' := (X \setminus \{x\}) \cup \{y'\}$ if $y' \notin X$ and $X' := X$ otherwise. The vertices x and y' are both in the 2-connected graph D which contains a triangle. By definition $|X \cap V(D)| = |X' \cap V(D)|$ and by assumption both sets are smaller than $|D|$. Let $\varphi : X \to X'$ any bijection with $\varphi|_{X \setminus V(D)} = \mathrm{id}|_{X \setminus V(D)}$ and $\varphi(x) = y'$. By Theorem 4.4 there is a balanced $L(\varphi|_{V(D)})$-movement of length at most $|D|!$ on D so by Lemma 4.1 (ii) there is a balanced $L(\varphi)$-movement $(\mathcal{X}_\varphi, \mathcal{M}_\varphi)$ of length at most $|D|!$ on H such that all its moves are contained in D.

Set $X'' := X' \setminus \{y, y'\}$ and let L' be the X'–X'' pairing with edge set $\{(z, 0)(z, \infty) \mid z \in X''\} \cup \{(y, 0)(y', 0)\}$. Clearly $((X', X''), (yy'))$ is an L'-movement. Let L'' be the X''–Y pairing obtained from L by deleting the edge $(x, 0)(y, 0)$ and substituting every vertex $(z, 0)$ with $(\varphi(z), 0)$. By definition we have $L = L(\varphi) \oplus L' \oplus L''$.

The set $A'' := A' \cup (A \cap \{y\})$ has at most 2 elements and thus $2|D| + |N(D) \setminus A''| \geq 2k+1$. So by induction there is an L''-movement $(\mathcal{X}'', \mathcal{M}'')$ of length at most $f(k-1)$ on $H - A''$ such that A is $(\mathcal{X}'', \mathcal{M}'')$-singular and every vertex of A that is not strongly $(\mathcal{X}'', \mathcal{M}'')$-singular is in $N_A \setminus (X \cup Y)$. Hence the movement

$$(\mathcal{X}, \mathcal{M}) := (\mathcal{X}_\varphi, \mathcal{M}_\varphi) \oplus \big((X', X''), (yy')\big) \oplus (\mathcal{X}'', \mathcal{M}'')$$

on $H - A'$ has induced pairing L by Lemma 4.2 and length at most $|D|! + 1 + f(k-1)$. Every move of $\mathcal{M}$ that contains a vertex of $A \setminus \{y\}$ is in $\mathcal{M}''$. Hence $A \setminus \{y\}$ is $(\mathcal{X}, \mathcal{M})$-singular and every vertex of $A \setminus \{y\}$ that is not strongly $(\mathcal{X}, \mathcal{M})$-singular is in $N_A \setminus (X \cup Y)$. If $y \notin A$, then we are done. But if $y \in A$, then our construction of $(\mathcal{X}'', \mathcal{M}'')$ ensures that no move of $\mathcal{M}''$ contains y. Therefore y is strongly $(\mathcal{X}, \mathcal{M})$-singular. □

Claim 4.10.2. *Suppose that* $|V(D) \setminus X| \geq 2$ *and that* L *has an edge* $(x, 0)(y, 0)$ *with* $x, y \in N(D)$. *Then there is an* L*-movement of length at most* $2|D|! + 2 + f(k-1)$ *on* H *such that* A *is singular and every vertex of* A *that is not strongly singular is in* $N_A \setminus (X \cup Y)$.

Proof. The proof is very similar to that of Claim 4.10.1. Let y' be a neighbour of y in D. We assume $y' \in X$ (in the case $y' \notin X$ we can skip the construction of $(\mathcal{X}_\varphi, \mathcal{M}_\varphi)$ in this paragraph). Let $z \in V(D) \setminus X$ and let $X' := (X \setminus \{y'\}) \cup \{z\}$. Let $\varphi : X \to X'$ be any bijection with $\varphi|_{X \setminus V(D)} = \mathrm{id}_{X \setminus V(D)}$ and $\varphi(y') = z$. Applying Theorem 4.4 and Lemma 4.1 as in the proof of Claim 4.10.1 we obtain a balanced $L(\varphi)$-movement $(\mathcal{X}_\varphi, \mathcal{M}_\varphi)$ of length at most $|D|!$ such that its moves are contained in D (in fact, we could "free" the vertex y' with only $|D|$ moves by shifting each token on a y'–z path in D by one position towards z, but we stick with the proof of Claim 4.10.1 here for simplicity).

Set $X'' := (X' \setminus \{y\}) \cup \{y'\}$ and let $\varphi' : X' \to X''$ be the bijection that maps y to y' and every other element to itself. Clearly $((X', X''), (yy'))$ is an $L(\varphi')$-movement. Let L'' be the X''–Y pairing obtained from L by substituting every vertex $(z, 0)$ with $(\varphi' \circ \varphi(z), 0)$. It is not hard to see that this construction implies $L = L(\varphi) \oplus L(\varphi') \oplus L''$. Since $(0, x)(0, y')$ is an edge of L'' with $x \in V(D) \cup N(D)$ and $y' \in V(D)$ we can apply Claim 4.10.1 to obtain an L''-movement $(\mathcal{X}'', \mathcal{M}'')$ of length at most $|D|! + 1 + f(k-1)$ on $H - (\{y\} \cap A)$ (note that $y \in A \cap X$ implies $y \notin Y$ by assumption) such that $A \setminus \{y\}$ is $(\mathcal{X}'', \mathcal{M}'')$-singular and every vertex of $A \setminus \{y\}$ that is not strongly $(\mathcal{X}'', \mathcal{M}'')$-singular is in $N_A \setminus (X \cup Y)$. Hence the movement

$$(\mathcal{X}, \mathcal{M}) := (\mathcal{X}_\varphi, \mathcal{M}_\varphi) \oplus \big((X', X''), (yy')\big) \oplus (\mathcal{X}'', \mathcal{M}'')$$

on H has induced pairing L by Lemma 4.2 and length at most $2|D|! + 2 + f(k-1)$. The argument that A is $(\mathcal{X}, \mathcal{M})$-singular and the only vertices of A that are not strongly $(\mathcal{X}, \mathcal{M})$-singular are in $N_A \setminus (X \cup Y)$ is the same as in the proof of Claim 4.10.1. $\square$

Pick any vertex $v \in V(D)$. There are sets $X', Y' \subseteq V(H)$ such that

1. $|X'| = |X|$ and $|Y'| = |Y|$,
2. $X \cap N_A \subseteq X'$ and $X' \cap A \subseteq X \cup N_A$,
3. $Y \cap N_A \subseteq Y'$ and $Y' \cap A \subseteq Y \cup N_A$,
4. $N_A \subseteq X' \cup Y'$ and $X' \cap Y' \cap A = \emptyset$,
5. $N_B \subseteq X'$ or $X' \subset N_A \cup N_B$,
6. $N_B \subseteq Y'$ or $Y' \subset N_A \cup N_B$,
7. $v \notin X'$ and $v \notin Y'$,
8. $V(D) \cup N_B \subseteq X' \cup \{v\}$ or $X' \subset V(D) \cup N(D)$, and
9. $V(D) \cup N_B \subseteq Y' \cup \{v\}$ or $Y' \subset V(D) \cup N(D)$.

By Lemma 4.6 it suffices to find an L'-movement $(\mathcal{X}', \mathcal{M}')$ of length at most $f(k) - 2k$ on H such that A is strongly $(\mathcal{X}', \mathcal{M}')$-singular where L' is any X'–Y' pairing.

Since $n \geq |D|$ we have $f(k) - 2k \geq 2|D|! + 2 + f(k-1)$ and by assumption $v \in V(D) \setminus X'$. If L' has an edge $(0, x)(0, y)$ with $x \in V(D)$ and $y \in V(D) \cup N(D)$, then by Claim 4.10.1 there is an L'-movement $(\mathcal{X}', \mathcal{M}')$ of length at most $f(k) - 2k$ on H such that A is strongly $(\mathcal{X}', \mathcal{M}')$-singular (recall that $N_A \setminus (X' \cup Y')$ is empty by choice of X' and Y'). So we may assume that L' contains no such edge and by Lemma 4.7 we may also assume that it has no edge $(x, 0)(y, \infty)$ with $x \in V(D)$ and $y \in N_A$.

Counting the edges of L' that are incident with a vertex of $(V(D) \cup N(D)) \times \{0\}$ we obtain the lower bound

$$\|L'\| \geq |X' \cap V(D)| + |X' \cap N_B|/2 + |(X \cup Y) \cap N_A|/2.$$

If $V(D) \cup N_B \subseteq X' \cup \{v\}$, then $|X' \cap V(D)| = |D| - 1$ and $|X' \cap N_B| = |N_B|$. Since $|(X' \cup Y') \cap N_A| = |N_A|$ this means

$$2k = 2\|L'\| \geq 2(|D| - 1) + |N_B| + |N_A| \geq 2|D| + |N(D)| - 2 \geq 2k + 1,$$

a contradiction. So we must have $X' \subset V(D) \cup N(D)$ and $|V(D) \setminus X'| \geq 2$. Applying Claim 4.10.2 in the same way as Claim 4.10.1 above we deduce that no edge of L' has both ends in $X \times \{0\}$ or one end in $X \times \{0\}$ and the other in $N_A \times \{\infty\}$. By symmetry we can obtain statements like Claim 4.10.1 and Claim 4.10.2 for Y instead of X thus by the same argument as above we may also assume that $Y' \subset V(D) \cup N(D)$ and that no edge of L' has both ends in $Y \times \{\infty\}$ or one end in $N_A \times \{0\}$ and the other in $Y \times \{\infty\}$. Hence L' is balanced and $N_A = \emptyset$. Let $\varphi' : X' \to Y'$ be the bijection with $L' = L(\varphi')$.

In the rest of the proof we apply the same techniques that we have already used in the proof of Claim 4.10.1 and again in that of Claim 4.10.2 so from now on we only sketch how to construct the desired movements. Furthermore, all constructed movements use only vertices of $V(D) \cup N_B$ for their moves so A is trivially strongly singular w.r.t. them.

If $N_B \setminus X' \neq \emptyset$, then we have $X' \subset N_B$ by assumption, so $|N_B| \geq k+1$ and thus also $Y' \subset N_B$. This is basically the same situation as at the end of the proof for Lemma 4.9 so we find an L'-movement of length at most $2k \leq f(k)$. We may therefore assume that $N_B \subseteq X' \cap Y'$.

Claim 4.10.3. *Suppose that L' has an edge $(x,0)(y,\infty)$ with $x \in V(D)$ and $y \in N_B$. Then there is an L'-movement $(\mathcal{X}', \mathcal{M}')$ of length at most $|D|! + 2 + f(k-1)$ such that the moves of $\mathcal{M}'$ are disjoint from A.*

Proof. Let y' be a neighbour of y in $V(D)$. Since D is 2-connected, y' has two distinct neighbours y_l and y_r in D. Using Theorem 4.4 we generate a balanced movement of length at most $|D|!$ on H such that all its moves are in D and its induced pairing has the edge $(x,0)(y_l,\infty)$ and its final configuration does not contain y' or y_r. Adding the two moves $yy'y_r$ and $y_l y' y$ then results in a movement $(\mathcal{X}_x, \mathcal{M}_x)$ of length at most $|D|! + 2$ whose induced pairing L_x contains the edge $(x,0)(y,0)$.

It is not hard to see that there is a pairing L'' such that $L_x \oplus L'' = L$ and this pairing must have the edge $(y,0)(y,\infty)$. By induction there is an L''-movement $(\mathcal{X}'', \mathcal{M}'')$ of length at most $f(k-1)$ such that none if its moves contains y. So $(\mathcal{X}', \mathcal{M}') := (\mathcal{X}_x, \mathcal{M}_x) \oplus (\mathcal{X}'', \mathcal{M}'')$ is an L' movement of length at most $|D|! + 2 + f(k-1)$ as desired. □

Claim 4.10.4. *Suppose that L' has an edge $(x,0)(y,\infty)$ with $x, y \in N_B$. Then there is an L'-movement $(\mathcal{X}', \mathcal{M}')$ of length at most $2|D|! + 4 + f(k-1)$ such that the moves of $\mathcal{M}'$ are disjoint from A.*

Proof. Let x' be a neighbour of x in $V(D)$. Since D is 2-connected, x' has two distinct neighbours x_l and x_r in D. With the same construction as in Claim 4.10.3 we can generate a movement $(\mathcal{X}_x, \mathcal{M}_x)$ of length at most $|D|! + 2$ such that its induced pairing L_x contains the edge $(x, 0)(x_r, \infty)$ and $(x'', 0)(x, \infty)$ for some vertex $x'' \in V(D) \cap X'$. There is a pairing L'' such that $L = L_x \oplus L''$ and L'' contains the edge $(x_r, 0)(y, \infty)$.

By Claim 4.10.3 there is an L''-movement $(\mathcal{X}'', \mathcal{M}'')$ of length at most $|D|! + 2 + f(k-1)$. So $(\mathcal{X}', \mathcal{M}') := (\mathcal{X}_x, \mathcal{M}_x) \oplus (\mathcal{X}'', \mathcal{M}'')$ is an L' movement of length at most $2|D|! + 4 + f(k-1)$ as desired. □

Since $f(k) - 2k \geq 2|D|! + 4 + f(k-1)$ we may assume that $N_B \cap Y' = \emptyset$ and thus $N_B = \emptyset$ by Claim 4.10.3 and Claim 4.10.4. This means $X', Y' \subseteq V(D)$ and therefore by Theorem 4.4 there is an L'-movement $(\mathcal{X}', \mathcal{M}')$ of length at most $|D|! \leq n! \leq f(k) - 2k$. This concludes the induction and thus also the proof of Lemma 4.10. □

5 Relinkages

This section collects several Lemma that compare different foundational linkages for the same stable regular decomposition of a graph. To avoid tedious repetitions we use the following convention throughout the section.

Convention. Let $(\mathcal{W}, \mathcal{P})$ be a stable regular decomposition of length $l \geq 3$ and attachedness p of a p-connected graph G. Set $\lambda := \{\alpha \mid P_\alpha \text{ is non-trivial}\}$ and $\theta := \{\alpha \mid P_\alpha \text{ is trivial}\}$. Let D be a block of $\Gamma(\mathcal{W}, \mathcal{P})[\lambda]$ and let κ be the set of all cut-vertices of $\Gamma(\mathcal{W}, \mathcal{P})[\lambda]$ that are in D.

Lemma 5.1. *Let $\mathcal{Q}$ be a foundational linkage. If $\alpha\beta$ is an edge of $\Gamma(\mathcal{W}, \mathcal{Q})$ with $\alpha \in \lambda$ or $\beta \in \lambda$, then $\alpha\beta$ is an edge of $\Gamma(\mathcal{W}, \mathcal{P})$.*

Proof. Some inner bag W_k of $\mathcal{W}$ contains a $\mathcal{Q}$-bridge B realising $\alpha\beta$, that is, B attaches to Q_α and Q_β. For $i = 1, \ldots, k-1$ the induced permutation π_i of $\mathcal{Q}[W_i]$ is an automorphism of $\Gamma(\mathcal{W}, \mathcal{P})$ by (L10) and hence so is the induced permutation $\pi = \prod_{i=1}^{k-1} \pi_i$ of $\mathcal{Q}[W_{[1,k-1]}]$.

Clearly the restriction of any induced permutation to θ is always the identity, so $\pi(\alpha) \in \lambda$ or $\pi(\beta) \in \lambda$. Therefore $\pi(\alpha)\pi(\beta)$ must be an edge of $\Gamma(\mathcal{W}, \mathcal{P})$ by (L11) as B attaches to $\mathcal{Q}[W]_{\pi(\alpha)}$ and $\mathcal{Q}[W]_{\pi(\beta)}$. Since π is an automorphism this means that $\alpha\beta$ is an edge of $\Gamma(\mathcal{W}, \mathcal{P})$. □

The previous Lemma allows us to make statements about any foundational linkage $\mathcal{Q}$ just by looking at $\Gamma(\mathcal{W}, \mathcal{P})$, in particular, for every $\alpha \in \lambda$ the neighbourhood $N(\alpha)$ of α in $\Gamma(\mathcal{W}, \mathcal{P})$ contains all neighbours of α in $\Gamma(\mathcal{W}, \mathcal{Q})$. The following Lemma applies this argument.

Lemma 5.2. *Let $\mathcal{Q}$ be a foundational linkage such that $\mathcal{Q}[W]$ is p-attached in $G[W]$ for each inner bag W of $\mathcal{W}$. If λ_0 is a subset of λ such that $|N(\alpha) \cap \theta| \leq p - 3$ for each $\alpha \in \lambda_0$, then every non-trivial $\mathcal{Q}$-bridge in an inner bag of $\mathcal{W}$ that attaches to a path of $\mathcal{Q}_{\lambda_0}$ must attach to at least one other path of $\mathcal{Q}_\lambda$.*

Proof. Suppose for a contradiction that some inner bag W of $\mathcal{W}$ contains a $\mathcal{Q}$-bridge B that attaches to some path $Q_\alpha[W]$ with $\alpha \in \lambda_0$ but to no other path of $\mathcal{Q}_\lambda[W]$. Recall that either all foundational linkages for $\mathcal{W}$ satisfy (L7) or none does and $\mathcal{P}$ witnesses the former. Hence by (L7) a path of $\mathcal{Q}[W]$ is non-trivial if and only if it is in $\mathcal{Q}_\lambda[W]$. So by p-attachedness $Q_\alpha[W]$ is bridge adjacent to at least $p-2$ paths of $\mathcal{Q}_\theta$ in $G[W]$. Therefore in $\Gamma(\mathcal{W}, \mathcal{Q})$ the vertex α is adjacent to at least $p-2$ vertices of θ and by Lemma 5.1 so it must be in $\Gamma(\mathcal{W}, \mathcal{P})$, giving the desired contradiction. □

Lemma 5.3. *Let $\mathcal{Q}$ be a foundational linkage. Every $\mathcal{Q}$-bridge B that attaches to a path of $\mathcal{Q}_{\lambda \setminus V(D)}$ has no edge or inner vertex in $G_D^{\mathcal{Q}}$, in particular, it can attach to at most one path of $\mathcal{Q}_{V(D)}$.*

Proof. By assumption B attaches to some path Q_α with $\alpha \in \lambda \setminus V(D)$. This rules out the possibility that B attaches to only one path of $\mathcal{Q}_\lambda$ that happens to be in $\mathcal{Q}_{V(D)}$. So if B has an edge or inner vertex in $G_D^{\mathcal{Q}}$, then it must realise an edge of D. Hence B attaches to paths Q_β and Q_γ with $\beta, \gamma \in V(D)$. This means that $\alpha\beta$ and $\alpha\gamma$ are both edges of $\Gamma(\mathcal{W}, \mathcal{Q})$ and thus of $\Gamma(\mathcal{W}, \mathcal{P})$ by Lemma 5.1. But D is a block of $\Gamma(\mathcal{W}, \mathcal{P})[\lambda]$ so no vertex of $\lambda \setminus V(D)$ can have two neighbours in D. □

Given two foundational linkages $\mathcal{Q}$ and $\mathcal{Q}'$ and a set $\lambda_0 \subseteq \lambda$, we say that $\mathcal{Q}'$ is a $(\mathcal{Q}, \lambda_0)$-*relinkage* or a *relinkage of $\mathcal{Q}$ on* λ_0 if $Q'_\alpha = Q_\alpha$ for $\alpha \notin \lambda_0$ and $\mathcal{Q}'_{\lambda_0} \subseteq G^{\mathcal{Q}}_{\lambda_0}$.

Lemma 5.4. *If $\mathcal{Q}$ is a $(\mathcal{P}, V(D))$-relinkage and $\mathcal{Q}'$ a $(\mathcal{Q}, V(D))$-relinkage, then $G_D^{\mathcal{Q}'} \subseteq G_D^{\mathcal{Q}}$, in particular, $G_D^{\mathcal{Q}} \subseteq G_D^{\mathcal{P}}$.*

Proof. Clearly $G_D^{\mathcal{Q}}$ and $G_D^{\mathcal{Q}'}$ are induced subgraphs of G so it suffices to show $V(G_D^{\mathcal{Q}'}) \subseteq V(G_D^{\mathcal{Q}})$. Suppose for a contradiction that there is a vertex $w \in V(G_D^{\mathcal{Q}'}) \setminus V(G_D^{\mathcal{Q}})$. We have $G_D^{\mathcal{Q}'} \cap \mathcal{Q}' = \mathcal{Q}'_{V(D)} \subseteq G_D^{\mathcal{Q}}$ so w must be an inner vertex of a $\mathcal{Q}'$-bridge B'. But w is in $G_\lambda - G_D^{\mathcal{Q}}$ and thus in a $\mathcal{Q}$-bridge attaching to a path of $\mathcal{Q}_{\lambda \setminus V(D)}$, in particular, there is a w–$\mathcal{Q}_{\lambda \setminus V(D)}$ path R that avoids $G_D^{\mathcal{Q}} \supseteq \mathcal{Q}'_{V(D)}$. This means $R \subseteq B'$ and thus B' attaches to a path of $\mathcal{Q}'_{\lambda \setminus V(D)} = \mathcal{Q}_{\lambda \setminus V(D)}$, a contradiction to Lemma 5.3. Clearly $\mathcal{P}$ itself is a $(\mathcal{P}, V(D))$-relinkage so $G_D^{\mathcal{Q}} \subseteq G_D^{\mathcal{P}}$ follows from a special case of the statement we just proved. □

Lemma 5.5. *Let $\mathcal{Q}$ be a $(\mathcal{P}, V(D))$-relinkage. If in $\Gamma(\mathcal{W}, \mathcal{P})$ we have $|N(\alpha) \cap \theta| \le p-3$ for all $\alpha \in \lambda \setminus V(D)$, then there is a $(\mathcal{Q}, V(D))$-relinkage $\mathcal{Q}'$ such that for every inner bag W of $\mathcal{W}$ the linkage $\mathcal{Q}'[W]$ is p-attached in $G[W]$ and has the same induced permutation as $\mathcal{Q}[W]$. Moreover, $\Gamma(\mathcal{W}, \mathcal{Q}')$ contains all edges of $\Gamma(\mathcal{W}, \mathcal{Q})$ that have at least one end in λ.*

Proof. Suppose that some non-trivial $\mathcal{Q}$-bridge B in an inner bag W of $\mathcal{W}$ attaches to a path $Q_\alpha = P_\alpha$ with $\alpha \in \lambda \setminus V(D)$ but to no other path of $\mathcal{Q}_\lambda$. Then B is also a $\mathcal{P}$-bridge and $\mathcal{P}[W]$ is p-attached in $G[W]$ by (L6) so $P_\alpha[W]$ must be bridge adjacent to at least $p-2$ paths of $\mathcal{P}_\theta$ in $G[W]$ and thus α has at least $p-2$ neighbours in θ, a contradiction. Hence every non-trivial $\mathcal{Q}$-bridge that attaches to a path of $\mathcal{Q}_{\lambda \setminus V(D)}$ must attach to at least one other path of $\mathcal{Q}_\lambda$.

For every inner bag W_i of $\mathcal{W}$ let $\mathcal{Q}'_i$ be the bridge stabilisation of $\mathcal{Q}[W_i]$ in $G[W_i]$. Then $\mathcal{Q}'_i$ has the same induced permutation as $\mathcal{Q}[W_i]$. Note that the set Z of all end vertices of the paths of $\mathcal{Q}[W_i]$ is the union of the left and right adhesion set of W_i. So by the p-connectivity of G for every vertex x of $G[W_i] - Z$ there is an x–Z fan of size p in $G[W_i]$. This means that $\mathcal{Q}'_i$ is p-attached in $G[W_i]$ by Lemma 3.7 (iii).

Hence $\mathcal{Q}' := \bigcup_{i=1}^{l-1} \mathcal{Q}'_i$ is a foundational linkage with $\mathcal{Q}'[W_i] = \mathcal{Q}'_i$ for $i = 1, \dots, l-1$. Therefore $\mathcal{Q}'[W]$ is p-attached in $G[W]$ and $\mathcal{Q}'[W]$ has the same induced permutations as $\mathcal{Q}[W]$ for every inner bag W of $\mathcal{W}$. There is no $\mathcal{Q}$-bridge that attaches to precisely one path of $\mathcal{Q}_{\lambda \setminus V(D)}$ but to no other path of $\mathcal{Q}_\lambda$ so we have $\mathcal{Q}'_{\lambda \setminus V(D)} = \mathcal{Q}_{\lambda \setminus V(D)}$ by Lemma 3.7 (i). The same result implies $\mathcal{Q}'_{V(D)} \subseteq G_D^{\mathcal{Q}}$ so $\mathcal{Q}'$ is indeed a relinkage of $\mathcal{Q}$ on $V(D)$.

Finally, Lemma 3.7 (ii) states that $\Gamma(\mathcal{W}, \mathcal{Q}')$ contains all those edges of $\Gamma(\mathcal{W}, \mathcal{Q})$ that have at least one end in λ. □

The "compressed" linkages presented next will allow us to fulfil the size requirement that Lemma 4.10 imposes on our block D as detailed in Lemma 5.7. Given a subset $\lambda_0 \subseteq \lambda$ and a foundational linkage $\mathcal{Q}$, we say that $\mathcal{Q}$ is *compressed to* λ_0 or λ_0*-compressed* if there is no vertex v of $G_{\lambda_0}^{\mathcal{Q}}$ such that $G_{\lambda_0}^{\mathcal{Q}} - v$ contains $|\lambda_0|$ disjoint paths from the first to the last adhesion set of $\mathcal{W}$ and v has a neighbour in $G_\lambda - G_{\lambda_0}^{\mathcal{Q}}$.

Lemma 5.6. *Suppose that in $\Gamma(\mathcal{W}, \mathcal{P})$ we have $|N(\alpha) \cap \theta| \le p-3$ for all $\alpha \in \lambda \setminus V(D)$ and let $\mathcal{Q}$ be a $(\mathcal{P}, V(D))$-relinkage. Then there is a $V(D)$-compressed $(\mathcal{Q}, V(D))$-relinkage $\mathcal{Q}'$ such that for every inner bag W of $\mathcal{W}$ the linkage $\mathcal{Q}'[W]$ is p-attached in $G[W]$.*

Proof. Clearly $\mathcal{Q}$ itself is a $(\mathcal{Q}, V(D))$-relinkage. Pick $\mathcal{Q}'$ from all $(\mathcal{Q}, V(D))$-relinkages such that $G_D^{\mathcal{Q}'}$ is minimal. By Lemma 5.4 and Lemma 5.5 we may assume that we picked $\mathcal{Q}'$ such that for every inner bag of W of $\mathcal{W}$ the linkage $\mathcal{Q}'[W]$ is p-attached in $G[W]$.

It remains to show that $\mathcal{Q}'$ is $V(D)$-compressed. Suppose not, that is, there is a vertex v of $G_D^{\mathcal{Q}'}$ such that v has a neighbour in $G_\lambda - G_D^{\mathcal{Q}'}$ and $G_D^{\mathcal{Q}'} - v$ contains an X–Y linkage $\mathcal{Q}''$ where X and Y denote the intersection of $V(G_D^{\mathcal{Q}'})$ with the first and last adhesion set of $\mathcal{W}$, respectively.

By Lemma 5.4 we have $G_D^{\mathcal{Q}''} \subseteq G_D^{\mathcal{Q}'} \subseteq G_D^{\mathcal{Q}}$ and thus $\mathcal{Q}''$ is a $(\mathcal{Q}, V(D))$-relinkage as well. This implies $G_D^{\mathcal{Q}''} = G_D^{\mathcal{Q}'}$ by the minimality of $G_D^{\mathcal{Q}'}$. The vertex v does not lie on a path of $\mathcal{Q}''$ by construction so it must be in a $\mathcal{Q}''$-bridge B''. But v has a neighbour w in $G_\lambda - G_D^{\mathcal{Q}'}$ and there is a w–$\mathcal{Q}'_{\lambda \setminus V(D)}$ path R that avoids $G_D^{\mathcal{Q}'}$. This means $R \subseteq B''$ and thus B'' attaches to a path of $\mathcal{Q}''_{\lambda \setminus V(D)}$, contradicting Lemma 5.3. □

Lemma 5.7. *Let $\mathcal{Q}$ be a $V(D)$-compressed foundational linkage. Let V be the set of all inner vertices of paths of $\mathcal{Q}_\kappa$ that have degree at least 3 in $G_D^{\mathcal{Q}}$. Then the following statements are true.*

- *(i) Either $2|D| + |N(D) \cap \theta| \geq p$ or $V(G_D^{\mathcal{Q}}) = V(\mathcal{Q}_{V(D)})$ and $\kappa \neq \emptyset$.*
- *(ii) Either $2|D| + |N(D)| \geq p$ or there is $\alpha \in \kappa$ such that $|Q_\beta| \leq |V \cap V(Q_\alpha)| + 1$ for all $\beta \in V(D) \setminus \kappa$.*

Note that $V(G_D^{\mathcal{Q}}) = V(\mathcal{Q}_{V(D)})$ implies that every $\mathcal{Q}$-bridge in an inner bag of $\mathcal{W}$ that realises an edge of D must be trivial.

Proof.

(i) Denote by X and Y the intersection of $G_D^{\mathcal{Q}}$ with the first and last adhesion set of $\mathcal{W}$, respectively. Let Z be the union of X, Y, and the set of all vertices of $G_D^{\mathcal{Q}}$ that have a neighbour in $G_\lambda - G_D^{\mathcal{Q}}$. Clearly $Z \subseteq V(\mathcal{Q}_\kappa) \cup X \cup Y$. Moreover, $G_D^{\mathcal{Q}} - z$ does not contain an X–Y linkage for any $z \in Z$: For $z \in X \cup Y$ this is trivial and for the remaining vertices of Z it holds by the assumption that $\mathcal{Q}$ is $V(D)$-compressed. Therefore for every $z \in Z$ there is an X–Y separation (A_z, B_z) of $G_D^{\mathcal{Q}}$ of order at most $|D|$ with $z \in A_z \cap B_z$. On the other hand, $\mathcal{Q}_{V(D)}$ is a set of $|D|$ disjoint X–Y paths in $G_D^{\mathcal{Q}}$ so every X–Y separation has order at least $|D|$. Hence by Lemma 3.1 there is a nested set $\mathcal{S}$ of X–Y separations of $G_D^{\mathcal{Q}}$, each of order $|D|$, such that $Z \subseteq Z_0$ where Z_0 denotes the set of all vertices that lie in a separator of a separation of $\mathcal{S}$.

We may assume that $(X, V(G_D^{\mathcal{Q}})) \in \mathcal{S}$ and $(V(G_D^{\mathcal{Q}}), Y) \in \mathcal{S}$ so for any vertex v of $G_D^{\mathcal{Q}} - (X \cup Y)$ there are $(A_L, B_L) \in \mathcal{S}$ and $(A_R, B_R) \in \mathcal{S}$

such that (A_L, B_L) is rightmost with $v \in B_L \setminus A_L$ and (A_R, B_R) is leftmost with $v \in A_R \setminus B_R$. Set $S_L := A_L \cap B_L$ and $S_R := A_R \cap B_R$.

Let z be any vertex of Z_0 "between" S_L and S_R, more precisely, $z \in (B_L \setminus A_L) \cap (A_R \setminus B_R)$. There is a separation $(A_M, B_M) \in \mathcal{S}$ such that its separator $S_M := A_M \cap B_M$ contains z. Then z witnesses that $A_M \not\subseteq A_L$ and $B_M \not\subseteq B_R$ and thus (A_M, B_M) is neither left of (A_L, B_L) nor right of (A_R, B_R). But $\mathcal{S}$ is nested and therefore (A_M, B_M) is strictly right of (A_L, B_L) and strictly left of (A_R, B_R). This means $v \in S_M$ otherwise (A_M, B_M) would be a better choice for (A_L, B_L) or for (A_R, B_R). So any separator of a separation of $\mathcal{S}$ that contains a vertex of $(B_L \setminus A_L) \cap (A_R \setminus B_R)$ must also contain v.

If $v \notin Z_0$, then $(B_L \setminus A_L) \cap (A_R \setminus B_R) \cap Z_0 = \emptyset$. This means that $S_L \cup S_R$ separates v from Z in $G_D^{\mathcal{Q}}$. So $S_L \cup S_R \cup V(\mathcal{Q}_{N(D)\cap\theta})$ separates v from $G - G_D^{\mathcal{Q}}$ in G. By the connectivity of G we therefore have

$$2|D| + |N(D) \cap \theta| \geq \left|S_L \cup S_R \cup V(\mathcal{Q}_{N(D)\cap\theta})\right| \geq p.$$

So we may assume that $V(G_D^{\mathcal{Q}}) = Z_0$ Since every separator of a separation of $\mathcal{S}$ consists of one vertex from each path of $\mathcal{Q}_{V(D)}$ this means $V(\mathcal{Q}_{V(D)}) \subseteq V(G_D^{\mathcal{Q}}) = Z_0 \subseteq V(\mathcal{Q}_{V(D)})$. If $\kappa = \emptyset$, then $X \cup Y \cup V(\mathcal{Q}_{N(D)\cap\theta})$ separates $G_D^{\mathcal{Q}} - (X \cup Y)$ from $G - G_D^{\mathcal{Q}}$ in G so this is just a special case of the above argument.

(ii) We may assume $\kappa \neq \emptyset$ by (i) and $\kappa \neq V(D)$ since the statement is trivially true in the case $\kappa = V(D)$. Pick $\alpha \in \kappa$ such that $|V \cap V(Q_\alpha)|$ is maximal and let $\beta \in V(D) \setminus \kappa$. For any inner vertex v of Q_β define (A_L, B_L) and (A_R, B_R) as in the proof of (i) and set $V_v := V \cap (B_L \setminus A_L) \cap (A_R \setminus B_R)$.

By (i) we have $V_v \subseteq Z_0$ and every separator of a separation of $\mathcal{S}$ that contains a vertex of V_v must also contain v. This means that $V_v \cap V_{v'} = \emptyset$ for distinct inner vertices v and v' of Q_β since no separator of a separation of $\mathcal{S}$ contains two vertices on the same path of $\mathcal{Q}_{V(D)}$.

Furthermore, $S_L \cup S_R \cup V_v$ separates v from $V(\mathcal{Q}_\kappa) \cup X \cup Y \supseteq Z$ in $G_D^{\mathcal{Q}}$ so by the same argument as in (i) we have $2|D| + |N(D) \cap \theta| + |V_v| \geq p$. Then $|N(D) \cap \lambda| \geq |V_v|$ would imply $2|D| + |N(D)| \geq p$ so we may assume that $|N(D) \cap \lambda| < |V_v|$ for all inner vertices v of Q_β. Clearly $N(D) \cap \lambda$ is a disjoint union of the sets $(N(\gamma) \cap \lambda) \setminus V(D)$ with $\gamma \in \kappa$ and these sets are all non-empty. Hence $|\kappa| \leq |N(D) \cap \lambda|$ and thus $|\kappa| + 1 \leq |V_v|$ for all inner vertices v of Q_β.

Write V for the inner vertices of $\mathcal{Q}_\beta$. Statement (ii) easily follows from

$$|V|(|\kappa|+1) \leq \sum_{v\in V} |V_v| \leq |V| \leq |\kappa| \cdot |V \cap V(Q_\alpha)|. \qquad \square$$

6 Rural Societies

In this section we present the answer of Robertson and Seymour to the question whether or not a graph can be drawn in the plane with specified vertices on the boundary of the outer face in a prescribed order. We will apply their result to subgraphs of a graph with a stable decomposition.

A *society* is a pair (G, Ω) where G is a graph and Ω is a cyclic permutation of a subset of $V(G)$ which we denote by $\bar{\Omega}$. A society (G, Ω) is called *rural* if there is a drawing of G in a closed disc D such that $V(G) \cap \partial D = \bar{\Omega}$ and Ω coincides with a cyclic permutation of $\bar{\Omega}$ arising from traversing ∂D in one of its orientations. We say that a society (G, Ω) is *k-connected* for an integer k if there is no separation (A, B) of G with $|A \cap B| < k$ and $\bar{\Omega} \subseteq B \neq V(G)$. For any subset $X \subseteq \bar{\Omega}$ denote by $\Omega|X$ the map on X defined by $x \mapsto \Omega^k(x)$ where k is the smallest positive integer such that $\Omega^k(x) \in X$ (chosen for each x individually). Since Ω is a cyclic permutation so is $\Omega|X$.

Given two internally disjoint paths P and Q in G we write PQ for the cyclic permutation of $V(P \cup Q)$ that maps each vertex of P to its successor on P if there is one and to the first vertex of $Q - P$ otherwise and that maps each vertex of $Q - P$ to its successor on $Q - P$ if there is one and to the first vertex of P otherwise.

Let R and S be disjoint $\bar{\Omega}$-paths in a society (G, Ω), with end vertices r_1, r_2 and s_1, s_2, respectively. We say that $\{R, S\}$ is a *cross in* (G, Ω), if $\Omega|\{r_1, r_2, s_1, s_2\} = (r_1 s_1 r_2 s_2)$ or $\Omega|\{r_1, r_2, s_1, s_2\} = (s_2 r_2 s_1 r_1)$.

The following is an easy consequence of Theorems 2.3 and 2.4 in [14].

Theorem 6.1 (Robertson & Seymour 1990)**.** *Any* 4*-connected society is rural or contains a cross.*

In our application we always want to find a cross. To prevent the society from being rural we force it to violate the implication given in following Lemma which is a simple consequence of Euler's formula.

Lemma 6.2. *Let (G, Ω) be a rural society. If the vertices in $V(G) \setminus \bar{\Omega}$ have degree at least* 6 *on average, then $\sum_{v\in\bar{\Omega}} d_G(v) \leq 4|\bar{\Omega}| - 6$.*

Proof. Since (G, Ω) is rural there is a drawing of G in a closed disc D with $V(G) \cap \partial D = \bar{\Omega}$. Let H be the graph obtained by adding one extra vertex w

outside D and joining it by an edge to every vertex on ∂D. Writing $b := |\bar{\Omega}|$ and $i := |V(G) \setminus \bar{\Omega}|$, Euler's formula implies

$$\|G\| + b = \|H\| \leq 3|H| - 6 = 3(i + b) - 3$$

and thus $\|G\| \leq 3i + 2b - 3$. Our assertion then follows from

$$\sum_{v \in \bar{\Omega}} d_G(v) + 6i \leq \sum_{v \in V(G)} d_G(v) = 2\|G\| \leq 6i + 4b - 6 \qquad \square$$

In our main proof we will deal with societies where the permutation Ω is induced by paths (see Lemma 6.4 and Lemma 6.5). But every inner vertex on such a path that has degree 2 in G adds slack to the bound provided by Lemma 6.2 as it counts 2 on the left side but 4 on the right. This is remedied in the following Lemma which allows us to apply Lemma 6.2 to a "reduced" society where these vertices are suppressed.

Lemma 6.3. *Let (G, Ω) be a society and let P be a path in G such that all inner vertices of P have degree 2 in G. Denote by G' the graph obtained from G by suppressing all inner vertices of P and set $\Omega' := \Omega|V(G')$. Then (G', Ω') is rural if and only if (G, Ω) is.*

Proof. The graph G is a subdivision of G' so every drawing of G gives a drawing of G' and vice versa. Hence a drawing witnessing that (G, Ω) is rural can easily be modified to witness that (G', Ω') is rural and vice versa. $\square$

Two vertices a and b of some graph H are called *twins* if $N_H(a) \setminus \{b\} = N_H(b) \setminus \{a\}$. Clearly a and b are twins if and only if the transposition (ab) is an automorphism of H.

Lemma 6.4. *Let G be a p-connected graph and let $(\mathcal{W}, \mathcal{P})$ be a stable regular decomposition of G of length at least 3 and attachedness p. Set $\theta := \{\alpha \mid P_\alpha \text{ is trivial}\}$ and $\lambda := \{\alpha \mid P_\alpha \text{ is non-trivial}\}$. Suppose that $\alpha\beta$ is an edge of $\Gamma(\mathcal{W}, \mathcal{P})[\lambda]$ such that $|N(\alpha) \cap \theta| \leq p - 3$, $|N(\beta) \cap \theta| \leq p - 3$, and for $N_{\alpha\beta} := N(\alpha) \cap N(\beta)$ we have $N_{\alpha\beta} \subseteq \theta$ and $|N_{\alpha\beta}| \leq p - 5$. If α and β are not twins, then the society $(G^{\mathcal{P}}_{\alpha\beta}, P_\alpha P_\beta^{-1})$ is rural.*

Proof.

Claim 6.4.1. *Every $\mathcal{P}$-bridge with an edge in $G^{\mathcal{P}}_{\alpha\beta}$ must attach to P_α and P_β, in particular, $G^{\mathcal{P}}_{\alpha\beta} - P_\alpha$ and $G^{\mathcal{P}}_{\alpha\beta} - P_\beta$ are both connected.*

Proof. By Lemma 5.2 every non-trivial $\mathcal{P}$-bridge that attaches to P_α or P_β must attach to another path of $\mathcal{P}_\lambda$. Since P_α and P_β are induced this means that all $\mathcal{P}$-bridges with an edge in $G^{\mathcal{P}}_{\alpha\beta}$ must realise the edge $\alpha\beta$ and hence attach to P_α and P_β. □

Claim 6.4.2. *The set Z of all vertices of $G^{\mathcal{P}}_{\alpha\beta}$ that are end vertices of P_α or P_β or have a neighbour in $G - (G^{\mathcal{P}}_{\alpha\beta} \cup \mathcal{P}_{N_{\alpha\beta}})$ is contained in $V(P_\alpha \cup P_\beta)$.*

Proof. Any vertex v of $G^{\mathcal{P}}_{\alpha\beta} - (P_\alpha \cup P_\beta)$ is an inner vertex of some non-trivial $\mathcal{P}$-bridge B that attaches to P_α and P_β. Since $G^{\mathcal{P}}_{\alpha\beta}$ contains all inner vertices of B the neighbours of v in $G - G^{\mathcal{P}}_{\alpha\beta}$ must be attachments of B. But if B attaches to a path P_γ with $\gamma \neq \alpha, \beta$, then $\gamma \in N_{\alpha\beta}$ and therefore all neighbours of v are in $G^{\mathcal{P}}_{\alpha\beta} \cup \mathcal{P}_{N_{\alpha\beta}}$. □

Claim 6.4.3. *The society $(G^{\mathcal{P}}_{\alpha\beta}, P_\alpha P_\beta^{-1})$ is rural if and only if the society $(G^{\mathcal{P}}_{\alpha\beta}, P_\alpha P_\beta^{-1}|Z)$ is.*

Proof. Clearly $(G^{\mathcal{P}}_{\alpha\beta}, P_\alpha P_\beta^{-1}|Z)$ is rural if $(G^{\mathcal{P}}_{\alpha\beta}, P_\alpha P_\beta^{-1})$ is. For the converse suppose that $(G^{\mathcal{P}}_{\alpha\beta}, P_\alpha P_\beta^{-1}|Z)$ is rural, that is, there is a drawing of $G^{\mathcal{P}}_{\alpha\beta}$ in a closed disc D such that $G^{\mathcal{P}}_{\alpha\beta} \cap \partial D = Z$ and one orientation of ∂D induces the cyclic permutation $P_\alpha P_\beta^{-1}|Z$ on Z.

For the rurality of $(G^{\mathcal{P}}_{\alpha\beta}, P_\alpha P_\beta^{-1})$ and $(G^{\mathcal{P}}_{\alpha\beta}, P_\alpha P_\beta^{-1}|Z)$ it does not matter whether the first vertices of P_α and P_β are adjacent in $G^{\mathcal{P}}_{\alpha\beta}$ or not and the same is true for the last vertices of P_α and P_β. So we may assume that both edges exist and we denote the cycle that they form together with the paths P_α and P_β by C.

The closed disc D' bounded by C is contained in D. It is not hard to see that the interior of D' is the only region of $D - C$ that has vertices of both P_α and P_β on its boundary. But every edge of $G^{\mathcal{P}}_{\alpha\beta}$ lies on C or in a $\mathcal{P}$-bridge B with $B - (\mathcal{P} \setminus \{P_\alpha, P_\beta\}) \subseteq G^{\mathcal{P}}_{\alpha\beta}$. By Claim 6.4.1 such a bridge B must attach to P_α and P_β and in the considered drawing it must therefore be contained in D'. This means $G^{\mathcal{P}}_{\alpha\beta} \subseteq D'$ which implies that $(G^{\mathcal{P}}_{\alpha\beta}, P_\alpha P_\beta^{-1})$ is rural as desired. □

Claim 6.4.4. *The society $(G^{\mathcal{P}}_{\alpha\beta}, P_\alpha P_\beta^{-1}|Z)$ is 4-connected.*

Proof. Set $H := G^{\mathcal{P}}_{\alpha\beta}$ and $\Omega := P_\alpha P_\beta^{-1}|Z$. Note that $\bar{\Omega} = Z$ since $Z \subseteq V(P_\alpha \cup P_\beta)$ by Claim 6.4.2. Set $T := V(\mathcal{P}_{N_{\alpha\beta}})$. Clearly $Z \cup T$ separates H from $G - H$ so for every vertex v of $H - Z$ there is a v–$T \cup Z$ fan of size at least p in G as G is p-connected. Since $|T| \leq p - 5$ this fan contains a v–Z fan of size at least 4 such that all its paths are contained in H. This means that (H, Ω) is 4-connected as desired. $\square$

By the *off-road edges* of a cross $\{R, S\}$ in (H, Ω) we mean the edges in $E(R \cup S) \setminus E(P_\alpha \cup P_\beta)$. We call a component of $R \cap (P_\alpha \cup P_\beta)$ that contains an end vertex of R a *tail of* R. We define the *tails of* S similarly.

Claim 6.4.5. *If $\{R, S\}$ is a cross in (H, Ω) whose set E of off-road edges is minimal, then for every $z \in Z \setminus V(R \cup S)$ each z–$(R \cup S)$ path in $P_\alpha \cup P_\beta$ ends in a tail of R or S.*

Proof. Suppose not, that is, there is a Z–$(R \cup S)$ path T in $P_\alpha \cup P_\beta$ such that its last vertex t does not lie in a tail of R or S. W.l.o.g. we may assume that t is on R. Since t is not in a tail of R the paths Rt and tR must both contain an edge that is not in $P_\alpha \cup P_\beta$ so $E(T \cup Rt \cup S) \setminus E(P_\alpha \cup P_\beta)$ and $E(T \cup tR \cup S) \setminus E(P_\alpha \cup P_\beta)$ are both proper subsets of E. But one of $\{T \cup Rt, S\}$ and $\{T \cup tR, S\}$ is a cross in (H, Ω), a contradiction. $\square$

Suppose now that α and β are not twins.

Claim 6.4.6. (H, Ω) *does not contain a cross.*

Proof. If (H, Ω) contains a cross, then we may pick a cross $\{R, S\}$ in (H, Ω) such that its set E of off-road edges is minimal. Since $Z \subseteq V(P_\alpha \cup P_\beta)$ we may assume w.l.o.g. that $\{R, S\}$ satisfies one of the following.

1. R and S both have their ends on P_α.
2. R has both ends on P_α. S has one end on P_α and one on P_β.
3. R and S both have one end on P_α and one on P_β.

We reduce the first case to the second. As P_β contains a vertex of Z but no end of R or S it must be disjoint from $R \cup S$ by Claim 6.4.5. But R and S both contain a vertex outside P_α (recall that P_α is induced by (L6)) so $R \cup S$ meets $H - P_\alpha$ which is connected by Claim 6.4.1.

Therefore there is a P_β–$(R \cup S)$ in $H - P_\alpha$, in particular, there is a Z–$(R \cup S)$ path T with its first vertex z in $Z \cap V(P_\beta)$ and we may assume that its last vertex t is on S. Denote by v the end of S that separates the ends of R in P_α.

Then $\{R, vSt \cup T\}$ is a cross in (H, Ω) and we may pick a cross $\{R', S'\}$ in (H, Ω) such that its set E' of off-road edges is minimal and contained in the set F of off-road edges of $\{R, vSt \cup T\}$. If $R' \cup S'$ contains no edge of T, then E' is a proper subset of E as it does not contain $E(S) \setminus E(vSt)$, a contradiction to the minimality of E. Hence $R' \cup S'$ contains an edge of T and hence must meet P_β. So by Claim 6.4.5 one of its paths, say S' ends in P_β as desired.

On the other hand, all off-road edges of $\{R', S'\}$ that are incident with P_β are in T and therefore the remaining three ends of R' and S' must all be on P_α. Hence $\{R', S'\}$ is a cross as in the second case.

In the second case we reroute P_α along R, more precisely, we obtain a foundational linkage $\mathcal{Q}$ from $\mathcal{P}$ by replacing the subpath of P_α between the two end vertices of R with R.

The first vertex of $R \cup S$ encountered when following P_β from either of its ends belongs to a tail of R or S by Claim 6.4.5. Obviously a tail contains precisely one end of R or S. Since R has no end on P_β and S only one, $(R \cup S) \cap P_\beta$ is a tail of S, in particular, R is disjoint from P_β and hence the paths of $\mathcal{Q}$ are indeed disjoint.

Clearly S must end in an inner vertex z of P_α. By the definition of Z there is a $\mathcal{P}$-bridge B in some inner bag W of $\mathcal{W}$ that attaches to z and to some path P_γ with $\gamma \in N(\alpha) \setminus N(\beta)$. But $B \cup S$ is contained in a $\mathcal{Q}$-bridge in $G[W]$ and therefore $\beta\gamma$ is an edge of $B(G[W], \mathcal{Q}[W])$ and thus of $\Gamma(\mathcal{W}, \mathcal{Q})$ but not of $\Gamma(\mathcal{W}, \mathcal{P})$. This contradicts Lemma 5.1.

In the third case Claim 6.4.5 ensures that the first and last vertex of P_α and of P_β in $R \cup S$ is always in a tail and clearly these tails must all be distinct. Hence by replacing the tails of R and S with suitable initial and final segments of P_α and P_β we obtain paths P'_α and P'_β such that the foundational linkage $\mathcal{Q} := (\mathcal{P} \setminus \{P_\alpha, P_\beta\}) \cup \{P'_\alpha, P'_\beta\}$ has the induced permutation $(\alpha\beta)$. Since $P_\gamma = Q_\gamma$ for all $\gamma \notin \{\alpha, \beta\}$ it is easy to see the there must be an inner bag W of $\mathcal{W}$ such that $\mathcal{Q}[W]$ has induced permutation $(\alpha\beta)$. But clearly $(\alpha\beta)$ is an automorphism of $\Gamma(\mathcal{W}, \mathcal{P})$ if and only if α and β are twins in $\Gamma(\mathcal{W}, \mathcal{P})$. Hence $\mathcal{Q}[W]$ is a twisting disturbance by the assumption that α and β are not twins. This contradicts the stability of $(\mathcal{W}, \mathcal{P})$ and concludes the proof of Claim 6.4.6. □

By Claim 6.4.4 and Theorem 6.1 the society (H, Ω) is rural or contains a cross. But Claim 6.4.6 rules out the latter so (H, Ω) is rural and by Claim 6.4.3 so is $(G^{\mathcal{P}}_{\alpha\beta}, P_\alpha P_\beta^{-1})$. □

In the previous Lemma we have shown how certain crosses in the graph $H = G^{\mathcal{P}}_{\alpha\beta}$ "between" two bridge-adjacent paths P_α and P_β of $\mathcal{P}$ give rise to disturbances. The next Lemma has a similar flavour; here the graph H will be the subgraph of G "between" P_α and Q_α where α is a cut-vertex of $\Gamma(\mathcal{W}, \mathcal{P})[\lambda]$ and $\mathcal{Q}$ a relinkage of $\mathcal{P}$.

Lemma 6.5. *Let G be a p-connected graph with a stable regular decomposition $(\mathcal{W}, \mathcal{P})$ of attachedness p and set $\lambda := \{\alpha \mid P_\alpha$ is non-trivial$\}$ and $\theta := \{\alpha \mid P_\alpha$ is trivial$\}$. Let D be a block of $\Gamma(\mathcal{W}, \mathcal{P})[\lambda]$ and let κ be the set of cut-vertices of $\Gamma(\mathcal{W}, \mathcal{P})[\lambda]$ that are in D. If $|N(\alpha) \cap \theta| \leq p - 4$ for all $\alpha \in \lambda$, then there is a $V(D)$-compressed $(\mathcal{P}, V(D))$-relinkage $\mathcal{Q}$ such that $\mathcal{Q}[W]$ is p-attached in $G[W]$ for all inner bags W of $\mathcal{W}$ and for any $\alpha \in \kappa$ and any separation (λ_1, λ_2) of $\Gamma(\mathcal{W}, \mathcal{P})[\lambda]$ such that $\lambda_1 \cap \lambda_2 = \{\alpha\}$ and $N(\alpha) \cap \lambda_2 = N(\alpha) \cap V(D)$ the following statements hold where $H := G^{\mathcal{P}}_{\lambda_2} \cap G^{\mathcal{Q}}_{\lambda_1}$, q_1 and q_2 are the first and last vertex of Q_α, and Z_1 and Z_2 denote the vertices of $H - \{q_1, q_2\}$ that have a neighbour in $G_\lambda - G^{\mathcal{P}}_{\lambda_2}$ and $G_\lambda - G^{\mathcal{Q}}_{\lambda_1}$, respectively.*

(i) We have $Z_1 \subseteq V(P_\alpha)$ and $Z_2 \subseteq V(Q_\alpha)$. Furthermore, $Z := \{q_1, q_2\} \cup Z_1 \cup Z_2$ separates H from $G_\lambda - H$ in $G - \mathcal{P}_{N(\alpha) \cap \theta}$.

(ii) The graph H is connected and contains Q_α. The path P_α ends in q_2.

(iii) Every cut-vertex of H is an inner vertex of Q_α and is contained in precisely two blocks of H.

(iv) Every block H' of H that is not a single edge contains a vertex of $Z_1 \setminus V(Q_\alpha)$ and a vertex of $Z_2 \setminus V(P_\alpha)$ that is not a cut-vertex of H. Furthermore, $Q_\alpha[W]$ contains a vertex of Z_2 for every inner bag W of $\mathcal{W}$.

(v) There is $(\mathcal{P}, V(D))$-relinkage $\mathcal{P}'$ with $\mathcal{P}' = (\mathcal{Q} \setminus \{Q_\alpha\}) \cup \{P'_\alpha\}$ and $P'_\alpha \subseteq H$ such that $Z_1 \subseteq V(P'_\alpha)$, $V(P'_\alpha \cap Q_\alpha)$ consists of q_1, q_2, and all cut-vertices of H, and $\mathcal{P}'[W]$ is p-attached in $G[W]$ for all inner bags W of $\mathcal{W}$.

(vi) Let H' be a block of H that is not a single edge. Then $P' := H' \cap P'_\alpha$ and $Q' := H' \cap Q_\alpha$ are internally disjoint paths with common first vertex q'_1 and common last vertex q'_2 and the society $(H', P'Q'^{-1})$ is rural.

Figure 1 gives an impression of H. The upper (straight) black q_1–q_2 path is Q_α and everything above it belongs to $G^{\mathcal{Q}}_{\lambda_2}$. The lower (curvy) black path is P'_α and everything below it belongs to $G^{\mathcal{P}}_{\lambda_1}$. The grey paths are subpaths of P_α and, as shown, P_α need not be contained in H and need not contain the vertices of $P_\alpha \cap P'_\alpha$ in the same order as P'_α. The white vertices are the cut-vertices of H. The vertices with an arrow up or down symbolise vertices of Z_2 and Z_1, respectively. The blocks of H that are not single edges are bounded by cycles in $P'_\alpha \cup Q_\alpha$ and Lemma 6.5 (vi) states that the part of H "inside" such a cycle forms a rural society.

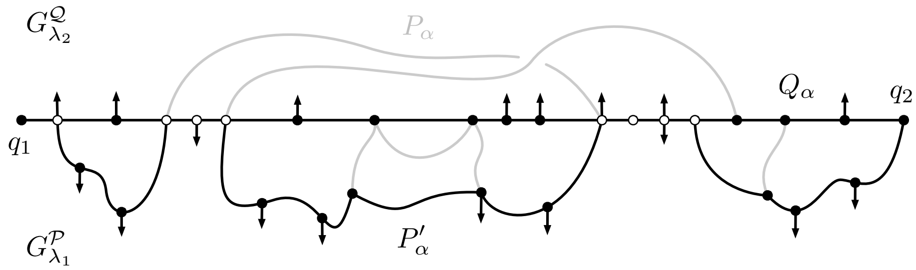

Figure 1: The graph $H = G^{\mathcal{P}}_{\lambda_2} \cap G^{\mathcal{Q}}_{\lambda_1}$.

Proof. For a $(\mathcal{P}, V(D))$-relinkage $\mathcal{Q}$ and $\beta \in \kappa$ any $G^{\mathcal{Q}}_D$-path $P \subseteq P_\beta$ such that some inner vertex of P has a neighbour in $G_\lambda - G^{\mathcal{P}}_D$ is called an *β-outlet of $\mathcal{Q}$*. By the *outlet graph of $\mathcal{Q}$* we mean the union of all components of $\mathcal{P}_\kappa - G^{\mathcal{Q}}_D$ that have a neighbour in $G_\lambda - G^{\mathcal{P}}_D$. In other words, the outlet graph of $\mathcal{Q}$ is obtained from the union of all β-outlets for all $\beta \in \kappa$ by deleting the vertices of $G^{\mathcal{Q}}_D$.

Clearly $\mathcal{P}$ itself is a $(\mathcal{P}, V(D))$-relinkage. Among all $(\mathcal{P}, V(D))$-relinkages pick $\mathcal{Q}'$ such that its outlet graph is maximal. By Lemma 5.6 there is a $V(D)$-compressed $(\mathcal{Q}', V(D))$-relinkage $\mathcal{Q}$ such that $\mathcal{Q}[W]$ is p-attached in $G[W]$ for all inner bags W of $\mathcal{W}$. Note that $G^{\mathcal{Q}}_D \subseteq G^{\mathcal{Q}'}_D$ by Lemma 5.4, so the outlet graph of $\mathcal{Q}$ is a supergraph of that of $\mathcal{Q}'$. Hence by choice of $\mathcal{Q}'$, they must be identical, in particular, the outlet graph of $\mathcal{Q}$ is maximal among the outlet graphs of all $(\mathcal{P}, V(D))$-relinkages.

Claim 6.5.1. *For any foundational linkage $\mathcal{R}$ of $\mathcal{W}$ we have $G^{\mathcal{R}}_{\lambda_1} \cup G^{\mathcal{R}}_{\lambda_2} = G_\lambda$ and $G^{\mathcal{R}}_{\lambda_1} \cap G^{\mathcal{R}}_{\lambda_2} = R_\alpha$.*

Proof. By Lemma 5.1 we have $\Gamma(\mathcal{W}, \mathcal{R})[\lambda] \subseteq \Gamma(\mathcal{W}, \mathcal{P})[\lambda]$, so (λ_1, λ_2) is also a separation of $\Gamma(\mathcal{W}, \mathcal{R})[\lambda]$. Hence each $\mathcal{R}$-bridge in an inner bag of $\mathcal{W}$ has all its attachments in $\mathcal{R}_{\lambda_1 \cup \theta}$ or all in $\mathcal{R}_{\lambda_2 \cup \theta}$ and thus $G^{\mathcal{R}}_{\lambda_1} \cup G^{\mathcal{R}}_{\lambda_2} = G_\lambda$. The induced path R_α is contained in $G^{\mathcal{R}}_{\lambda_1} \cap G^{\mathcal{R}}_{\lambda_2}$ by definition. If $G^{\mathcal{R}}_{\lambda_1} \cap G^{\mathcal{R}}_{\lambda_2}$ contains a vertex that is not on R_α, then it must be in a non-trivial $\mathcal{R}$-bridge that attaches to R_α but to no other path of $\mathcal{R}_\lambda$. Such a bridge does not exist by Lemma 5.2 (applied to $\lambda_0 := \lambda$). $\square$

Claim 6.5.2. *For every vertex v of $H - P_\alpha$ there is a v–Z_2 path in $H - P_\alpha$ and for every vertex v of $H - Q_\alpha$ there is a v–Z_1 path in $H - Q_\alpha$.*

Proof. Let v be a vertex of $H - P_\alpha \subseteq G^{\mathcal{P}}_{\lambda_2} - P_\alpha$. Then there is $\beta \in \lambda_2 \setminus \lambda_1$ such that v is on P_β or v is an inner vertex of some non-trivial $\mathcal{P}$-bridge attaching to P_β by Lemma 5.2 and the assumption that $|N(\alpha) \cap \theta| \leq p - 4$. In either case $G^{\mathcal{P}}_{\lambda_2} - P_\alpha$ contains a path R from v to the first vertex p of P_β. But p is also the first vertex of Q_β and therefore it is contained in $G_\lambda - G^{\mathcal{Q}}_{\lambda_1}$. Pick w on R such that Rw is a maximal initial subpath of R that is still contained in H. Then $w \neq p$ and the successor of w on R must be in $G_\lambda - G^{\mathcal{Q}}_{\lambda_1}$. This means $w \in Z_2$ as desired. If v is in $H - Q_\alpha$, then the argument is similar but slightly simpler as $Q_\beta = P_\beta$ for all $\beta \in \lambda_1 \setminus \lambda_2$. □

(i) Any vertex of $G^{\mathcal{P}}_{\lambda_2}$ that has a neighbour in $G^{\mathcal{P}}_{\lambda_1} - G^{\mathcal{P}}_{\lambda_2}$ must be on P_α by Claim 6.5.1. This shows $Z_1 \subseteq V(P_\alpha)$ and by a similar argument $Z_2 \subseteq V(Q_\alpha)$.

A neighbour v of H in G either is in no inner bag of $\mathcal{W}$, it is in G_λ, or it is in $\mathcal{P}_\theta$. In the first case v can only be adjacent to q_1 or q_2 as these are the only vertices of H in the first and last adhesion set of $\mathcal{W}$.

In the second case, note that $\mathcal{Q}$ is a $(\mathcal{P}, \lambda_2)$-relinkage since $V(D) \subseteq \lambda_2$ and thus Lemma 5.4 yields $G^{\mathcal{Q}}_{\lambda_2} \subseteq G^{\mathcal{P}}_{\lambda_2}$ which together with Claim 6.5.1 implies

$$\begin{aligned} G_\lambda = G^{\mathcal{P}}_{\lambda_1} \cup G^{\mathcal{P}}_{\lambda_2} &= G^{\mathcal{P}}_{\lambda_1} \cup (G^{\mathcal{P}}_{\lambda_2} \cap G^{\mathcal{Q}}_{\lambda_1}) \cup (G^{\mathcal{P}}_{\lambda_2} \cap G^{\mathcal{Q}}_{\lambda_2}) \\ &= G^{\mathcal{P}}_{\lambda_1} \cup H \cup G^{\mathcal{Q}}_{\lambda_2}. \end{aligned}$$

Hence v is in $G_\lambda - G^{\mathcal{Q}}_{\lambda_1}$ or in $G_\lambda - G^{\mathcal{P}}_{\lambda_2}$ and thus all neighbours of v in H are in Z_2 or Z_1, respectively.

In the third case v is the unique vertex of some path P_β with $\beta \in \theta$. Let w be a neighbour of v in H. Either w is on P_α or there is a w–Z_2 path in H by Claim 6.5.2 which ends on Q_α as shown above. So $\alpha\beta$ is an edge of $\Gamma(\mathcal{W}, \mathcal{P})$ or of $\Gamma(\mathcal{W}, \mathcal{Q})$. The former implies $\beta \in N(\alpha)$ directly and the latter does with the help of Lemma 5.1. Hence we have shown that $Z \cup V(\mathcal{P}_{N(\alpha)\cap\theta})$ separates H from the rest of G concluding the proof of (i).

(ii) We have $Q_\alpha \subseteq G^{\mathcal{Q}}_{\lambda_1}$ by definition and $Q_\alpha \subseteq G^{\mathcal{P}}_{\lambda_2}$ since $\mathcal{Q}$ is a $(\mathcal{P}, \lambda_2)$-relinkage. Hence $Q_\alpha \subseteq H$ and some component C of H contains Q_α. Suppose that v is a vertex of $H \cap P_\alpha$. Let w be the vertex of P_α such that $wP_\alpha v$ is a maximal subpath of P_α that is still contained in H. Since $P_\alpha \subseteq G^{\mathcal{P}}_{\lambda_2}$ we must have $w \in \{q_1\} \cup Z_2 \subseteq V(Q_\alpha)$ and hence v is in C. For any vertex v of $H - P_\alpha$ there is a v–Z_2 path in H by Claim 6.5.2

which ends on Q_α by (i). This means that v is in C and hence H is connected.

For every inner bag W of $\mathcal{W}$ the induced permutation π of $\mathcal{Q}[W]$ maps each element of $\lambda_1 \setminus \lambda_2$ to itself as $\mathcal{Q}$ is $(\mathcal{P}, \lambda_2)$-relinkage. Moreover, π is an automorphism of $\Gamma(\mathcal{W}, \mathcal{P})$ by (L10) and α is the unique vertex of λ_2 that has a neighbour in $\lambda_1 \setminus \lambda_2$. This shows $\pi(\alpha) = \alpha$. Hence Q_α and P_α must have the same end vertex, namely q_2.

(iii) Let v be a cut-vertex of H. By (ii) it suffices to show that all components of $H - v$ contain a vertex of Q_α. First note that every component of $H - v$ contains a vertex of Z: If a vertex w of $H - v$ is not in Z, then by (i) and the connectivity of G there is a w–Z fan of size at least $p - |N(\alpha) \cap \theta| \geq 2$ in H and at most one of its paths contains v. But any vertex $z \in Z \setminus V(Q_\alpha)$ is on P_α by (i) and the paths $q_1 P_\alpha z$ and $z P_\alpha q_2$ do both meet Q_α but at most one can contain v (given that $z \neq v$). So every component of $H - v$ must contain a vertex of Q_α as claimed.

Claim 6.5.3. *A Q_α-path $P \subseteq P_\alpha \cap H$ is an α-outlet if and only if some inner vertex of P is in Z_1, in particular, every vertex of $Z_1 \setminus V(Q_\alpha)$ lies in a unique α-outlet. Denoting the union of all α-outlets by U, no two components of $Q_\alpha - U$ lie in the same component of $H - U$.*

Proof. Clearly $Q_\alpha \subseteq G_D^{\mathcal{Q}} \cap H \subseteq G_{\lambda_2}^{\mathcal{Q}} \cap G_{\lambda_1}^{\mathcal{Q}} = Q_\alpha$ by Claim 6.5.1. Suppose that $P \subseteq P_\alpha \cap H$ has some inner vertex $z_1 \in Z_1$. Then P is a $G_D^{\mathcal{Q}}$-path and z_1 has a neighbour in $G_\lambda - G_{\lambda_2}^{\mathcal{P}} \subseteq G_\lambda - G_D^{\mathcal{P}}$ so P is an α-outlet.

Before we prove the converse implication let us show that $H \subseteq G_D^{\mathcal{P}}$. If some vertex v of $H \subseteq G_{\lambda_2}^{\mathcal{P}}$ is not in $G_D^{\mathcal{P}}$, then there is $\beta \in \lambda_2 \setminus V(D)$ such that v is on P_β or v is an inner vertex of a non-trivial $\mathcal{P}$-bridge attaching to P_β. But v is in $H - P_\alpha$ so by Claim 6.5.2 and (i) there is a v–Q_α path in H and hence $\alpha\beta$ is an edge of $\Gamma(\mathcal{W}, \mathcal{Q})[\lambda]$ and thus of $\Gamma(\mathcal{W}, \mathcal{P})[\lambda]$ by Lemma 5.1. But (λ_1, λ_2) is chosen such that $N(\alpha) \cap \lambda \subseteq \lambda_1 \cup V(D)$, a contradiction.

Suppose that P is an α-outlet. Then some inner vertex z of P has a neighbour in $G_\lambda - G_D^{\mathcal{P}} \subseteq G_\lambda - H$. So $z \in Z_1 \cup Z_2$ and therefore $z \in Z_1$ as $z \notin V(Q_\alpha) \supseteq Z_2$ by (i).

To conclude the proof of the claim we may assume for a contradiction that Q_α contains vertices r_1, r, and r_2 in this order such that $H - U$ contains an r_1–r_2 path R and r is the end vertex of an α-outlet. Let $\mathcal{Q}'$ be the foundational linkage obtained from $\mathcal{Q}$ by replacing the subpath $r_1 Q_\alpha r_2$ of Q_α with R. Clearly $\mathcal{Q}'$ is a $(\mathcal{P}, V(D))$-relinkage. It suffices to show that the outlet graph of $\mathcal{Q}'$ properly contains that of $\mathcal{Q}$ to derive a contradiction to our choice of $\mathcal{Q}$.

By choice of R and the construction of $\mathcal{Q}'$ each β-outlet of $\mathcal{Q}$ for any $\beta \in \kappa$ is internally disjoint from $\mathcal{Q}'$ and hence is contained in a β-outlet of $\mathcal{Q}'$. But r is not on Q'_α so it is an inner vertex of some α-outlet of $\mathcal{Q}'$ so the outlet graph of $\mathcal{Q}'$ contains that of $\mathcal{Q}$ properly as desired. □

Claim 6.5.4. *Let r_1 and r_2 be the end vertices of an α-outlet P of $\mathcal{Q}$. Then $r_1 Q_\alpha r_2$ contains a vertex of $Z_2 \setminus V(P_\alpha)$.*

Proof. We assume that r_1 and r_2 occur on Q_α in this order. Set $Q := r_1 Q_\alpha r_2$. Clearly $P \cup Q$ is a cycle. Since P_α is induced in G, some inner vertex v of Q is not on P_α. By Claim 6.5.2 there is a v–Z_2 path R in $H - P_\alpha$ and its last vertex z_2 must be on Q_α (see (i)) but not on P_α. Finally, Claim 6.5.3 implies that v and z_2 must be in the same component of $Q_\alpha - P$ so both are on Q as desired. □

(iv) Clearly H' contains a cycle. Since Q_α is induced in G there must be a vertex v in $H' - Q_\alpha$ and the v–Z_1 path in $H - Q_\alpha$ that exists by Claim 6.5.2 avoids all cut-vertices of H by (iii) and thus lies in $H' - Q_\alpha$. So H' contains a vertex of $Z_1 - V(Q_\alpha)$ which lies on P_α by (i) and thus also an α-outlet by Claim 6.5.3. So by Claim 6.5.4 we must also have a vertex of $Z_2 \setminus V(P_\alpha)$ in H' that is neither the first nor the last vertex of Q_α in H'.

For any inner bag W of $\mathcal{W}$ the end vertices of $Q_\alpha[W]$ are cut-vertices of H. By (L8) $G[W]$ contains a $\mathcal{P}$-bridge realising some edge of D that is incident with α. So some vertex of P_α has a neighbour in $G_\lambda - G^{\mathcal{P}}_{\lambda_1}$. If $Q_\alpha[W] = P_\alpha[W]$, then $G^{\mathcal{Q}}_{\lambda_1}[W] = G^{\mathcal{P}}_{\lambda_1}[W]$ so this neighbour is also in $G_\lambda - G^{\mathcal{Q}}_{\lambda_1}$ and hence $Q_\alpha[W]$ contains a vertex of Z_2. If $Q_\alpha[W] \neq P_\alpha[W]$, then some block of H in $G[W]$ is not a single edge so by the previous paragraph $Q_\alpha[W]$ contains a vertex of Z_2.

Claim 6.5.5. *Every Z_1–Z_2 path in H is a q_1–q_2 separator in H.*

Proof. Suppose not, that is, H contains a q_1–q_2 path Q'_α and a Z_1–Z_2 path R such that R and Q'_α are disjoint. Clearly $H \cap \mathcal{Q} = Q_\alpha$ so $\mathcal{Q}' := (\mathcal{Q} \setminus \{Q_\alpha\}) \cup \{Q'_\alpha\}$ is a foundational linkage. The last vertex r_2 of R is in Z_2 and hence has a neighbour in $G_\lambda - G^{\mathcal{Q}}_{\lambda_1}$. So there is an r_2–$\mathcal{Q}'_{\lambda_2 \setminus \lambda_1}$ path R_2 that meets H only in r_2. Similarly, for the first vertex r_1 of R there is an r_1–$\mathcal{Q}'_{\lambda_1 \setminus \lambda_2}$ path R_1 that meets H only in r_1. Then $R_1 \cup R \cup R_2$ witness that $\Gamma(\mathcal{W}, \mathcal{Q}')$ has an edge with one end in $\lambda_1 \setminus \lambda_2$ and the other in $\lambda_2 \setminus \lambda_1$, contradicting Lemma 5.1. □

Claim 6.5.6. *Let H' be a block of H. Then $Q := H' \cap Q_\alpha$ is a path and its first vertex q_1' equals q_1 or is a cut-vertex of H and its last vertex q_2' equals q_2 or is a cut-vertex of H. Furthermore, there is a q_1'–q_2' path $P \subseteq H'$ that is internally disjoint from Q such that $Z_1 \cap V(H') \subseteq V(P)$ and if a P-bridge B in H' has no inner vertex on Q_α, then for every $z_1 \in Z_1 \cap V(H')$ the attachments of B are either all on Pz_1 or all on z_1P.*

Proof. It follows easily from (iii) that Q is a path and q_1' and q_2' are as claimed. If H' is the single edge $q_1'q_2'$, then the statement is trivial with $P = Q$ so suppose not. Our first step is to show the existence of a q_1'–q_2' path $R \subseteq H'$ that is internally disjoint from Q.

By (iv) some inner vertex z_2 of Q is in $Z_2 \setminus V(P_\alpha)$. Since H' is 2-connected there is a Q-path $R \subseteq H' - z_2$ with first vertex r_1 on Qz_2 and last vertex r_2 on z_2Q. Pick R such that r_1Qr_2 is maximal. We claim that $r_1 = q_1'$ and $r_2 = q_2'$.

Suppose for a contradiction that $r_2 \neq q_2'$. By the same argument as before there is Q-path $S \subseteq H' - r_2$ with first vertex s_1 on Qr_2 and last vertex s_2 on r_2Q. Note that s_1 must be an inner vertex of r_1Qr_2 by choice of R. Similarly, Q separates R from S in H' otherwise there was a Q-path from r_1 to s_2 again contradicting our choice of R.

But S has an inner vertex v as Q is induced and Claim 6.5.2 asserts the existence of a v–Z_1 path S' in $H - Q_\alpha$ which must be disjoint from R as Q separates S from R. So there is a Z_1–Z_2 path in $z_2Qs_1 \cup s_1Sv \cup S'$ which is disjoint from $Q_\alpha r_1Rr_2Q_\alpha$ by construction, a contradiction to Claim 6.5.5. This shows $r_2 = q_2'$ and by symmetry also $r_1 = q_1'$.

Among all q_1'–q_2' paths in H' that are internally disjoint from Q pick P such that P contains as few edges outside P_α as possible. To show that P contains all vertices of $Z_1 \cap V(H')$ let $z_1 \in Z_1 \cap V(H')$. We may assume $z_1 \neq q_1', q_2'$. If z_1 is an inner vertex of Q, then Q contains a Z_1–Z_2 path that is disjoint from P, a contradiction to Claim 6.5.5. So there is an α-outlet R which has z_1 as an inner vertex. Then $Rz_1 \cup Q$ and $z_1R \cup Q$ both contain a Z_1–Z_2 path and by Claim 6.5.5 P must intersect both paths. But P is internally disjoint from Q so it contains a vertex t_1 of Rz_1 and a vertex t_2 of z_1R. If some edge of t_1Pt_2 is not on P_α, then $P' := q_1'Pt_1P_\alpha t_2Pq_2'$ is q_1'–q_2' path in H' that is internally disjoint from Q and has fewer edges outside P_α than P, contradicting our choice of P. This means $t_1Rt_2 \subseteq P$ and therefore z_1 is on P.

Finally, suppose that for some $z_1 \in Z_1$ there is a P-bridge B in H' with no inner vertex in Q_α and attachments $t_1, t_2 \neq z_1$ such that t_1 is on Pz_1 and t_2 is on z_1P (this implies $z_1 \neq q_1', q_2'$). Let R be the α-outlet containing z_1

and denote its end vertices by r_1 and r_2. By Claim 6.5.4 some inner vertex z_2 of r_1Qr_2 is in Z_2.

If B has an attachment in $z_1P - R$, then $z_1Rr_2 \cup z_2Qr_2$ contains a Z_1–Z_2 path that does not separate q_1' from q_2' in H' and therefore does not separate q_1 from q_2 in H, contradicting Claim 6.5.5. So B has no attachment in $z_1P - R$ and a similar argument implies that B has no attachment in $Pz_1 - R$. So all attachments of B must be in $P \cap R \subseteq P_\alpha$. As $R \cup B$ contains a cycle and P_α is induced some vertex v of B is not on P_α. But then Claim 6.5.2 implies the existence of a v–Z_2 path that avoids P_α and hence uses only inner vertices of B, in particular, some inner vertex of B is in $Z_2 \subseteq V(Q_\alpha)$, contradicting our assumption and concluding the proof of this claim. □

(v) Applying Claim 6.5.6 to every block H' of H and uniting the obtained paths P gives a q_1–q_2 path $R \subseteq H$ such that $Z_1 \subseteq V(R)$ and $V(R \cap Q_\alpha)$ consists of q_1, q_2, and all cut-vertices of H. Moreover, for every $z_1 \in Z_1$ a P-bridge B in H that has no inner vertex in Q_α has all its attachments in Rz_1 or all in z_1R.

Set $\mathcal{Q}' := (\mathcal{Q} \setminus \{Q_\alpha\}) \cup \{R\})$. Let $\mathcal{P}'$ be the foundational linkage obtained by uniting the bridge stabilisation of $\mathcal{Q}'[W]$ in $G[W]$ for all inner bags W of $\mathcal{W}$. Then $\mathcal{P}'[W]$ is p-attached in $G[W]$ for all inner bags W of $\mathcal{W}$ by Lemma 3.7.

To show $P'_\beta = Q_\beta$ for all $\beta \in \lambda \setminus \{\alpha\}$ it suffices by Lemma 3.7 to check that every non-trivial $\mathcal{Q}'$-bridge B' that attaches to Q'_β attaches to at least one other path of $\mathcal{Q}'_\lambda$. If B' is disjoint from H it is also a $\mathcal{Q}$-bridge and thus attaches to some path $Q_\gamma = Q'_\gamma$ with $\gamma \in \lambda \setminus \{\alpha, \beta\}$ by Claim 6.5.1. If B' contains a vertex of H, then it attaches to $Q'_\alpha = R$ as H is connected (see (ii)) and $\mathcal{Q}' \cap H = R$.

To verify $P'_\alpha \subseteq H$ we need to show $B' \subseteq H$ for every $\mathcal{Q}'$-bridge B' that attaches to R but to no other path of $\mathcal{Q}'_\lambda$. Clearly for every vertex v of $G^{\mathcal{P}}_{\lambda_1} - P_\alpha$ there is a v–$\mathcal{P}_{\lambda_1 \setminus \{\alpha\}}$ path in $G^{\mathcal{P}}_{\lambda_1} - P_\alpha$. Similarly, for every vertex v of $G^{\mathcal{Q}}_{\lambda_2} - Q_\alpha$ there is a v–$\mathcal{Q}_{\lambda_2 \setminus \{\alpha\}}$ path in $G^{\mathcal{Q}}_{\lambda_2} - Q_\alpha$. But $Q'_\beta = P_\beta$ for all $\beta \in \lambda_1 \setminus \{\alpha\}$ and $Q'_\beta = Q_\beta$ for all $\beta \in \lambda_2 \setminus \{\alpha\}$ and $G_\lambda - H = (G^{\mathcal{P}}_{\lambda_1} - P_\alpha) \cup (G^{\mathcal{Q}}_{\lambda_2} - Q_\alpha)$. This means that B' cannot contain a vertex of $G_\lambda - H$ and thus $B' \subseteq H$ as desired.

We have just shown that every bridge B' as above is an R-bridge in H. By construction and the properties (i) and (iv) every component of $Q_\alpha - R$ contains a vertex of Z_2 and hence lies in a $\mathcal{Q}'$-bridge attaching to some path Q'_β with $\beta \in \lambda_2 \setminus \{\alpha\}$. So B' is an R-bridge in H with no inner vertex in Q_α and therefore there must be $z_1, z_1' \in Z_1 \cup \{q_1, q_2\}$ such that z_1Rz_1' contains all attachments of B' and no inner vertex of z_1Rz_1' is in Z_1. By Lemma 3.7 this implies that P_α' contains no vertex of $Q_\alpha - R$ and $Z_1 \subseteq V(P_\alpha')$. On the other hand, P_α' must clearly contain the end vertices of R and all cut-vertices of H. This concludes the proof of (v).

(vi) We first show that (H', Ω) is rural where $\Omega := P'Q'^{-1}|Z$ where $Z' := Z \cap V(H')$. Since H is connected and $H \cap \mathcal{P}' = P_\alpha$ we must have $\beta \in N(\alpha) \cap \theta$ for each path P_β with $\beta \in \theta$ whose unique vertex has a neighbour in H. So the set T of all vertices of $\mathcal{P}_\theta$ that are adjacent to some vertex of H' has size at most $p-4$ by assumption. Clearly $Z' \cup T$ separates H' from the rest of G so for every vertex v of $H' - Z'$ there is a v–$(Z' \cup T)$ fan of size at least p and hence a v–Z fan of size at least 4. Hence (H', Ω) is 4-connected and hence it is rural or contains a cross by Theorem 6.1.

Suppose for a contradiction that (H', Ω) contains a cross. By the *off-road edges* of a cross $\{R, S\}$ in (H', Ω) we mean edge set $E(R \cup S) \setminus E(P' \cup Q')$. We call a component of $R \cap (P' \cup Q')$ that contains an end of R a *tail of R* and define the *tails of S* similarly.

Claim 6.5.7. *If $\{R, S\}$ is a cross in (H', Ω) such that its set of off-road edges is minimal, then for every $z \in Z$ that is not in $R \cup S$ the two z–$(R \cup S)$ paths in $P' \cup Q'$ both end in a tail of R or S.*

The proof is the same as for Claim 6.4.5 so we spare it.

Claim 6.5.8. *Every non-trivial $(P' \cup Q')$-bridge B in H' has an attachment in $P' - Q'$ and in $Q' - P'$.*

Proof. Let v be an inner vertex of B. Then $H - Q_\alpha$ contains a v–Z_1 path by Claim 6.5.2 so B must attach to P'. Note that v is in a non-trivial $\mathcal{P}'$-bridge B' and $B' \subseteq G^{\mathcal{P}}_{\lambda_2}$ since $Z_1 \subseteq V(P_\alpha')$. Furthermore, B' must attach to a path $P_\beta' = Q_\beta$ with $\beta \in \lambda_2 \setminus \lambda_1$: This is clear if B' does not attach to P_α' and follows from Claim 6.5.1 if it does. So B' contains a path R from v to $G^{\mathcal{Q}}_{\lambda_2} - Q_\alpha$ that avoids P'. But any such path contains a vertex of Z_2 (see (i)) and R does not contain q_1' and q_2' so some initial segment of R is a v–Z_2 path in $H' - P'$ as desired. $\square$

Claim 6.5.9. *There is a cross $\{R', S'\}$ in (H', Ω) such that its set of off-road edges is minimal and neither P' nor Q' contains all ends of R' and S'.*

Proof. Pick a cross $\{R, S\}$ in (H', Ω) such that its set E of off-road edges is minimal. We may assume that P' contains all ends of R and S. By (iv) some inner vertex z_2 of Q' is in Z_2. So if $R \cup S$ contains an inner vertex of Q', then $Q' - P'$ contains a Z_2–$(R \cup S)$ path T whose last vertex t is an inner vertex of R say. Clearly one of $\{Rt \cup T, S\}$ and $\{tR \cup T, S\}$ is a cross in (H', Ω) whose set of off-road edges is contained in that of $\{R, S\}$ and hence is minimal as well. So either we find a cross $\{R', S'\}$ as desired or $Q' - P'$ is disjoint from $R \cup S$.

But $(R \cup S) - P'$ must be non-empty as P' is induced in G. So by Claim 6.5.8 there is a Q'–$(R \cup S)$ path in $H' - P'$, in particular, there is a Z_2–$(R \cup S)$ path T in $H' - P'$ and we may assume that its last vertex t is on R. Again one of $\{Rt \cup T, S\}$ and $\{tR \cup T, S\}$ is a cross in (H', Ω) and we denote its set of off-road edges by F. Pick a cross (R', S') in (H, Ω) such that its set E' of off-road edges minimal and $E' \subseteq F$.

Since t is not on P' each of Rt and tR contains an edge that is not in $P' \cup Q'$ so $F \setminus E(T)$ is a proper subset of E. This means that E' must contain an edge of T by minimality of E and hence it must contain $F \cap E(T)$ so $R' \cup S'$ contains a vertex of $Q' - P'$ and we have already seen that we are done in this case, concluding the proof of the claim. □

Claim 6.5.10. *For $i = 1, 2$ there is a q'_i–$(R' \cup S')$ path T_i in H' such that T_1 and T_2 end on one path of $\{R', S'\}$ and the other path has its ends in Z_1 and Z_2.*

Proof. It is easy to see that by construction one path of $\{R', S'\}$, say S', has one end in $Z_1 \setminus \{q'_1, q'_2\}$ and the other in $Z_2 \setminus \{q'_1, q'_2\}$. If for some i the vertex q'_i is in $R' \cup S'$, then it must be on R' and there is a trivial q'_i–R' path T_i. We may thus assume that neither of q'_1 and q'_2 is in $R' \cup S'$.

So $P' \cup Q'$ contains two q'_1–$(R' \cup S')$ paths T_1 and T'_1 that meet only in q'_1. By Claim 6.5.7 T_1 and T'_1 must both end in a tail of R' or S'. But (R', S') is a cross and no inner vertex of $T_1 \cup T'_1$ is an end of R' or S' so we may assume that T_1 meets a tail of R'. By the same argument we find a q'_2–$(R' \cup S')$ path T_2 that end in a tail of R'. □

To conclude the proof that (H', Ω) is rural note that Claim 6.5.10 implies the existence of a Z_1–Z_2 path in H that does not separate q_1 from q_2 in H and hence contradicts Claim 6.5.5. So (H', Ω) is rural and (vi) follows from this final claim:

Claim 6.5.11. *The society* (H', Ω) *is rural if and only if the society* $(H', P'Q'^{-1})$ *is.*

This holds by a simpler version of the proof of Claim 6.4.3 where Claim 6.5.8 takes the role of Claim 6.4.1.

□

7 Constructing a Linkage

In our main theorem we want to construct the desired linkage in a long stable regular decomposition of the given graph. That decomposition is obtained by applying Theorem 3.5 which may instead give a subdivision of $K_{a,p}$. This outcome is even better for our purpose as stated by the following Lemma.

Lemma 7.1. *Every* $2k$*-connected graph containing a* $TK_{2k,2k}$ *is* k*-linked.*

Proof. Let G be a $2k$-connected graph and let $S, T \subseteq V(G)$ be disjoint and of size k each, say $S = \{s_1, \ldots, s_k\}$ and $T = \{t_1, \ldots, t_k\}$. We need to find a system of k disjoint S–T paths linking s_i to t_i for $i = 1, \ldots, k$.

By assumption G contains a subdivision of $K_{2k,2k}$, so there are disjoint sets $A, B \subseteq V(G)$ of size $2k$ each and a system $\mathcal{Q}$ of internally disjoint paths in G such that for every pair (a, b) with $a \in A$ and $b \in B$ there exists a unique a–b path in $\mathcal{Q}$ which we denote by Q_{ab}.

By the connectivity of G, there is a system $\mathcal{P}$ of $2k$ disjoint $(S \cup T)$–$(A \cup B)$ paths (with trivial members if $(S \cup T) \cap (A \cup B) \neq \emptyset$). Pick $\mathcal{P}$ such that it has as few edges outside of $\mathcal{Q}$ as possible. Our aim is to find suitable paths of $\mathcal{Q}$ to link up the paths of $\mathcal{P}$ as desired. We denote by A_1 and B_1 the vertices of A and B, respectively, in which a path of $\mathcal{P}$ ends, and let $A_0 := A \setminus A_1$ and $B_0 := B \setminus B_1$.

The paths of $\mathcal{P}$ use the system $\mathcal{Q}$ sparingly: Suppose that for some pair (a, b) with $a \in A_0$ and $b \in B$, the path Q_{ab} intersects a path of $\mathcal{P}$. Follow Q_{ab} from a to the first vertex v it shares with any path of $\mathcal{P}$, say P. Replacing P by $Pv \cup Q_{ab}v$ in $\mathcal{P}$ does not give a system with fewer edges outside $\mathcal{Q}$ by our choice of $\mathcal{P}$. In particular, the final segment vP of P must have no edges outside $\mathcal{Q}$. This means $vP = vQ_{ab}$, that is, P is the only path of $\mathcal{P}$ meeting Q_{ab} and after doing so for the first time it just follows Q_{ab} to b. Clearly the symmetric argument works if $a \in A$ and $b \in B_0$. Hence

1. Q_{ab} with $a \in A_0$ and $b \in B_0$ is disjoint from all paths of $\mathcal{P}$,

2. Q_{ab} with $a \in A_1$ and $b \in B_0$ or with $a \in A_0$ and $b \in B_1$ is met by precisely one path of $\mathcal{P}$, and

3. Q_{ab} with $a \in A_1$ and $b \in B_1$ is met by at least two paths of $\mathcal{P}$.

In order to describe precisely how we link the paths of $\mathcal{P}$, we fix some notation. Since $|A_0|+|A_1| = |A| = 2k = |\mathcal{P}| = |A_1|+|B_1|$, we have $|A_0| = |B_1|$ and similarly $|A_1| = |B_0|$. Without loss of generality we may assume that $|B_0| \geq |A_0| = |B_1|$ and therefore $|B_0| \geq k$. So we can pick k distinct vertices $b_1, \ldots, b_k \in B_0$ and an arbitrary bijection $\varphi : B_1 \to A_0$. For $x \in S \cup T$ denote by P_x the unique path of $\mathcal{P}$ starting in x and by x' its end vertex in $A \cup B$.

For each i and $x = s_i$ or $x = t_i$ set

$$R_x := \begin{cases} Q_{x'b_i} & x' \in A_1 \\ Q_{\varphi(x')x'} \cup Q_{\varphi(x')b_i} & x' \in B_1 \end{cases}.$$

By construction R_x and R_y intersect if and only if $x, y \in \{s_i, t_i\}$ for some i, i.e. they are equal or meet exactly in b_i. The paths P_x and R_y intersect if and only if P_x ends in y', that is, if $x = y$. Thus for each $i = 1, \ldots, k$ the subgraph $C_i := P_{s_i} \cup R_{s'_i} \cup R_{t'_i} \cup P_{t_i}$ of G is a tree containing s_i and t_i. Furthermore, these trees are pairwise disjoint, finishing the proof. □

We now give the proof of the main theorem, Theorem 1.1. We restate the theorem before proceeding with the proof.

Theorem 1.1. For all integers k and w there exists an integer N such that a graph G is k-linked if

$$\kappa(G) \geq 2k + 3, \qquad \mathrm{tw}(G) < w, \quad \text{and} \quad |G| \geq N.$$

Proof. Let k and w be given and let f be the function from the statement of Lemma 4.10 with $n := w$. Set

$$\begin{aligned} n_0 &:= (2k+1)(n_1 - 1) + 1 \\ n_1 &:= \max\{(2k-1)\binom{w}{2k}, 2k(k+3)+1, 12k+4, 2f(k)+1\} \end{aligned}$$

We claim that the theorem is true for the integer N returned by Theorem 3.5 for parameters $a = 2k$, $l = n_0$, $p = 2k + 3$, and w. Suppose that G is a $(2k+3)$-connected graph of tree-width less than w on at least N vertices. We want to show that G is k-linked. If G contains a subdivision of $K_{2k,2k}$, then this follows from Lemma 7.1. We may thus assume that G does not contain such a subdivision, in particular it does not contain a subdivision of $K_{a,p}$.

Let $S = (s_1, \dots, s_k)$ and $T = (t_1, \dots, t_k)$ be disjoint k-tuples of distinct vertices of G. Assume for a contradiction that G does not contain disjoint paths $P_1, \dots, P_k$ such that the end vertices of P_i are s_i and t_i for $i = 1, \dots, k$ (such paths will be called the *desired paths* in the rest of the proof).

By Theorem 3.5 there is a stable regular decomposition of G of length at least n_0, of adhesion $q \leq w$, and of attachedness at least $2k + 3$. Since this decomposition has at least $(2k + 1)(n_1 - 1)$ inner bags, there are $n_1 - 1$ consecutive inner bags which contain no vertex of $(S \cup T)$ apart from those coinciding with trivial paths. In other words, this decomposition has a contraction $(\mathcal{W}, \mathcal{P})$ of length n_1 such that $S \cup T \subseteq W_0 \cup W_{n_1}$. By Lemma 3.6 this contraction has the same attachedness and adhesion as the initial decomposition and the stability is preserved. Set $\theta := \{\alpha \mid P_\alpha \text{ is trivial}\}$ and $\lambda := \{\alpha \mid P_\alpha \text{ is non-trivial}\}$.

Claim 7.1.1. $\lambda \neq \emptyset$.

Proof. If $\lambda = \emptyset$, or equivalently, $\mathcal{P} = \mathcal{P}_\theta$, then all adhesion sets of $\mathcal{W}$ equal $V(\mathcal{P}_\theta)$. So by (L2) no vertex of $G - \mathcal{P}_\theta$ is contained in more than one bag of $\mathcal{W}$. On the other hand, (L4) implies that every bag W of $\mathcal{W}$ must contain a vertex $w \in W \setminus V(\mathcal{P}_\theta)$. Since $V(\mathcal{P}_\theta)$ separates W from the rest of G and G is $2k$-connected, there is a w–$\mathcal{P}_\theta$ fan of size $2k$ in $G[W]$. For different bags, these fans meet only in $\mathcal{P}_\theta$.

Since $\mathcal{W}$ has more than $(2k-1)\binom{q}{2k}$ bags, the pigeon hole principle implies that there are $2k$ such fans with the same $2k$ end vertices among the q vertices of $\mathcal{P}_\theta$. The union of these fans forms a $TK_{2k,2k}$ in G which may not exist by our earlier assumption. □

Claim 7.1.2. *Let Γ_0 be a component of $\Gamma(\mathcal{W}, \mathcal{P})[\lambda]$. The following all hold.*

(i) $|N(\alpha) \cap \theta| \leq 2k - 2$ *for every vertex* α *of* Γ_0.

(ii) $|N(\alpha) \cap N(\beta) \cap \theta| \leq 2k - 4$ *for every edge* $\alpha\beta$ *of* Γ_0.

(iii) $2|N(\alpha) \cap \lambda| + |N(\alpha) \cap \theta| \leq 2k$ *for every vertex* α *of* Γ_0.

(iv) $2|D| + |N(D)| \leq 2k + 2$ *for every block* D *of* Γ_0 *that contains a triangle.*

Note that (iii) implies (i) unless Γ_0 is a single vertex and (iii) implies (ii) unless Γ_0 is a single edge. We need precisely these two cases in the proof of Claim 7.1.6.

Proof. The proof is almost identical for all cases so we do it only once and point out the differences as we go. Denote by Γ_1 the union of Γ_0 with all its incident edges of $\Gamma(\mathcal{W},\mathcal{P})$. Set $L := W_0 \cap W_1 \cap V(G_{\Gamma_1})$ and $R := W_{n_1-1} \cap W_{n_1} \cap V(G_{\Gamma_1})$. In case (iv) let α be any vertex of D. Let p and q be the first and last vertex of P_α. Then $(L \cup R) \setminus \{p,q\}$ separates $G^{\mathcal{P}}_{\Gamma_1} - \{p,q\}$ from $S \cup T$ in $G - \{p,q\}$. Hence by the connectivity of G there is a set $\mathcal{Q}$ of $2k$ disjoint $(S \cup T)$–$(L \cup R)$ paths in $G - \{p,q\}$, each meeting $G^{\mathcal{P}}_{\Gamma_1}$ only in its last vertex. For $i = 1,\dots,k$ denote by s'_i the end vertex of the path of $\mathcal{Q}$ that starts in s_i and by t'_i the end vertex of the path of $\mathcal{Q}$ that starts in t_i.

Our task is to find disjoint s'_i–t'_i paths for $i = 1,\dots,k$ in $G^{\mathcal{P}}_{\Gamma_1}$ and we shall now construct sets $X, Y \subseteq V(\Gamma_1)$ and an X–Y pairing L "encoding" this by repeating the following step for each $i \in \{1,\dots,k\}$. Let $\beta,\gamma \in V(\Gamma_1)$ such that s'_i lies on P_β and t'_i lies on P_γ. If $s'_i \in L$, then add β to X and set $\bar{s}_i := (\beta, 0)$. Otherwise $s'_i \in R \setminus L$ and we add β to Y and set $\bar{s}_i := (\beta,\infty)$. Note that $s'_i \in L \cap R$ if and only if $\beta \in \theta$. In this case our decision to add β to X is arbitrary and we could also add it to Y instead (and setting $\bar{s}_i$ accordingly) without any bearing on the proof. Handle γ and t'_i similarly. Then $\{\bar{s}_i\bar{t}_i \mid i = 1,\dots,k\}$ is the edge set of an (X,Y)-pairing which we denote by L.

We claim that there is an L-movement of length at most $(n_1-1)/2 \geq f(k)$ on $H := \Gamma_1$ such that the vertices of $A := V(\Gamma_1) \cap \theta$ are singular. Clearly $H - A = \Gamma_0$ is connected and every vertex of A has a neighbour in Γ_0 so A is marginal in H. The existence of the desired L-movement follows from Lemma 4.8 if (i) or (ii) are violated, from Lemma 4.9 if (iii) is violated, and from Lemma 4.10 if (iv) is violated (note that $|H| \leq w$). But then Lemma 4.3 applied to L implies the existence of disjoint s'_i–t'_i paths in $G^{\mathcal{P}}_{\Gamma_1}$ for $i = 1,\dots,k$ contradicting our assumption that G does not contain the desired paths. This shows that all conditions must hold. □

Claim 7.1.3. *We have* $2|\Gamma_0| + |N(\Gamma_0)| \geq 2k+3$ *(and necessarily* $N(\Gamma_0) \subseteq \theta$*) for every component* Γ_0 *of* $\Gamma(\mathcal{W},\mathcal{P})[\lambda]$.

Proof. Let Γ_1 be the union of Γ_0 with all incident edges of $\Gamma(\mathcal{W},\mathcal{P})$. Set $L := W_0 \cap W_1 \cap V(G^{\mathcal{P}}_{\Gamma_1})$, $M := W_1 \cap W_2 \cap V(G^{\mathcal{P}}_{\Gamma_1})$, and $R := W_{n_1-1} \cap W_{n_1} \cap V(G^{\mathcal{P}}_{\Gamma_1})$. If $G - G^{\mathcal{P}}_{\Gamma_1}$ is non-empty, then $L \cup R$ separates it from M in G. Otherwise M separates L from R in $G = G^{\mathcal{P}}_{\Gamma_1}$. By the connectivity of G we have $2|\Gamma_0| + |N(\Gamma_0)| = |L \cup R| \geq 2k+3$ in the former case and $|M| = |\Gamma_0| + |N(\Gamma_0)| \geq 2k+3$ in the latter. □

We now want to apply Lemma 6.4 and Lemma 6.5. At the heart of both is the assertion that a certain society is rural and we already limited the number of their "ingoing" edges by Lemma 6.2. To obtain a contradiction we shall find societies exceeding this limit. Tracking these down is the purpose of the notion of "richness" which we introduce next.

Let $\Gamma \subseteq \Gamma(\mathcal{W},\mathcal{P})[\lambda]$. We say that $\alpha \in V(\Gamma)$ is *rich in* Γ if the inner vertices of P_α that have a neighbour in both $G_\lambda - G^{\mathcal{P}}_\Gamma$ and $G^{\mathcal{P}}_\Gamma - P_\alpha$ have average degree at least $2 + |N_\Gamma(\alpha)|(2+\varepsilon_\alpha)$ in $G^{\mathcal{P}}_\Gamma$ where $\varepsilon_\alpha := 1/|N(\alpha)\cap\lambda|$. A subgraph $\Gamma \subseteq \Gamma(\mathcal{W},\mathcal{P})[\lambda]$ is called *rich* if every vertex $\alpha \in V(\Gamma)$ is rich in Γ.

Claim 7.1.4. *For $\Gamma \subseteq \Gamma(\mathcal{W},\mathcal{P})[\lambda]$ and $\alpha \in V(\Gamma)$ the following is true.*

(i) If Γ contains all edges of $\Gamma(\mathcal{W},\mathcal{P})[\lambda]$ that are incident with α, then α is rich in Γ.

(ii) If α is rich in Γ, then the inner vertices of P_α that have a neighbour in $G^{\mathcal{P}}_\Gamma - P_\alpha$ have average degree at least $2 + |N_\Gamma(\alpha)|(2+\varepsilon_\alpha)$ in $G^{\mathcal{P}}_\Gamma$.

(iii) Suppose that Γ is induced in $\Gamma(\mathcal{W},\mathcal{P})[\lambda]$ and that there are subgraphs $\Gamma_1,\dots,\Gamma_m \subseteq \Gamma$ such that α separates any two of them in $\Gamma(\mathcal{W},\mathcal{P})[\lambda]$ and $\bigcup_{i=1}^m \Gamma_i$ contains all edges of Γ that are incident with α. If α is rich in Γ, then there is $j \in \{1,\dots,m\}$ such that α is rich in Γ_j.

Proof.

(i) The assumption implies that $G^{\mathcal{P}}_\Gamma$ contains every edge of G_λ that is incident with P_α so no vertex of P_α has a neighbour in $G_\lambda - G^{\mathcal{P}}_D$ and therefore the statement is trivially true.

(ii) The inner vertices of P_α that have a neighbour in $G_\lambda - G^{\mathcal{P}}_\Gamma$ and in $G^{\mathcal{P}}_\Gamma - P_\alpha$ have the desired average degree by assumption. We show that each inner vertex of P_α that has no neighbour in $G_\lambda - G^{\mathcal{P}}_\Gamma$ has at least the desired degree. Clearly we have $d_{G^{\mathcal{P}}_\Gamma}(v) = d_{G_\lambda}(v)$ for such a vertex v. Furthermore, $d_G(v) \geq 2k+3$ since G is $(2k+3)$-connected. Every neighbour of v in $\mathcal{P}_\theta$ gives rise to a neighbour of α in θ and by Claim 7.1.2 (iii) there can be at most $|N(\alpha)\cap\theta| \leq 2k - 2|N(\alpha)\cap\lambda|$ such neighbours. This means

$$d_{G^{\mathcal{P}}_\Gamma}(v) = d_{G_\lambda}(v) \geq 2k+3 - |N(\alpha)\cap\theta| \geq 2|N(\alpha)\cap\lambda| + 3$$

so (ii) clearly holds.

(iii) We may assume that α is not isolated in Γ and that each of the graphs $\Gamma_1,\dots,\Gamma_m$ contains an edge of Γ that is incident with α by simply forgetting those graphs that do not.

For $i = 0, \dots, m$ denote by Z_i the inner vertices of P_α that have a neighbour in $G_\lambda - G^{\mathcal{P}}_{\Gamma_i}$ and in $G^{\mathcal{P}}_{\Gamma_i} - P_\alpha$ where $\Gamma_0 := \Gamma$ and set $Z := \bigcup_{i=1}^m Z_i$. Clearly $P_\alpha \subseteq G^{\mathcal{P}}_{\Gamma_i}$ for all i. Each edge e of $G^{\mathcal{P}}_{\Gamma}$ that is incident with an inner vertex of P_α but does not lie in P_α is in a $\mathcal{P}$-bridge that realises an edge of Γ by (L6) and Lemma 5.2 since Claim 7.1.2 (iii) implies that $|N(\alpha) \cap \theta| \leq 2k - 2$. So at least one of the graphs $G^{\mathcal{P}}_{\Gamma_i}$ contains e. On the other hand, we have $G^{\mathcal{P}}_{\Gamma_i} \subseteq G^{\mathcal{P}}_{\Gamma}$ for $i = 1, \dots, m$. This implies $Z_0 \subseteq Z$.

By the same argument as in the proof of (ii) the vertices of Z have average degree at least $2 + |N_\Gamma(\alpha)|(2 + \varepsilon_\alpha)$ in $G^{\mathcal{P}}_{\Gamma}$. In other words, $G^{\mathcal{P}}_{\Gamma}$ contains at least $|Z||N_\Gamma(\alpha)|(2 + \varepsilon_\alpha)$ edges with one end on P_α and the other in $G^{\mathcal{P}}_{\Gamma} - P_\alpha$.

By assumption we have $|N_\Gamma(\alpha)| = \sum_{i=1}^m |N_{\Gamma_i}(\alpha)|$ and so the pigeon hole principle implies that there is $j \in \{1, \dots, m\}$ such that $G^{\mathcal{P}}_{\Gamma_j}$ contains a set E of at least $|Z||N_{\Gamma_j}(\alpha)|(2 + \varepsilon_\alpha)$ edges with one end on P_α and the other in $G^{\mathcal{P}}_{\Gamma_j} - P_\alpha$.

By assumption and Claim 6.5.1 the path P_α separates $G^{\mathcal{P}}_{\Gamma_i}$ from $G^{\mathcal{P}}_{\Gamma_j}$ in G_λ for $i \neq j$. For any vertex $z \in Z \setminus Z_j$ there is $i \neq j$ with $z \in Z_i$, so z has a neighbour in $G^{\mathcal{P}}_{\Gamma_i} - P_\alpha \subseteq G_\lambda - G^{\mathcal{P}}_{\Gamma_j}$. Then the only reason that z is not also in Z_j is that it has no neighbour in $G^{\mathcal{P}}_{\Gamma_j} - P_\alpha$, in particular, it is not incident with an edge of E. So the vertices of Z_j have average degree at least $2 + \frac{|Z|}{|Z_j|}|N_{\Gamma_j}(\alpha)|(2 + \varepsilon_\alpha)$ in $G^{\mathcal{P}}_{\Gamma_j}$ which obviously implies the claimed bound.

□

Claim 7.1.5. *Every component of* $\Gamma(\mathcal{W}, \mathcal{P})[\lambda]$ *contains a rich block.*

Proof. Let Γ_0 be a component of $\Gamma(\mathcal{W}, \mathcal{P})[\lambda]$. Suppose that α is a cut-vertex of Γ_0 and let $D_1, \dots, D_m$ be the blocks of Γ_0 that contain α. Clearly $N(\alpha) \cap \lambda \subseteq V(\bigcup_{i=1}^m D_i)$ so Claim 7.1.4 implies that α is rich in $\bigcup_{i=1}^m D_i$ by (i) and hence there is $j \in \{1, \dots, m\}$ such that α is rich in D_j by (iii).

We define an oriented tree R on the set of blocks and cut-vertices of Γ_0 as follows. Suppose that D is a block of Γ_0 and α a cut-vertex of Γ_0 with $\alpha \in V(D)$. If α is rich in D, then we let (α, D) be an edge of R. Otherwise we let (D, α) be an edge of R. Note that the underlying graph of R is the block-cut-vertex tree of Γ_0 and by the previous paragraph every cut-vertex is incident with an outgoing edge of R. But every directed tree has a sink, so there must be a block D of Γ_0 such that every $\alpha \in \kappa$ is rich in D where κ denotes the set of all cut-vertices of Γ_0 that lie in D.

But the only vertices of $G^{\mathcal{P}}_D$ that may have a neighbour in $G_\lambda - G^{\mathcal{P}}_D$ are on paths of $\mathcal{P}_{V(D)}$ by Lemma 5.3 and of these clearly only the paths of $\mathcal{P}_\kappa$ may have neighbours in $G_\lambda - G^{\mathcal{P}}_D$. So all vertices of $V(D) \setminus \kappa$ are trivially rich in D and hence D is a rich block. □

Claim 7.1.6. *Every rich block D of $\Gamma(\mathcal{W}, \mathcal{P})[\lambda]$ contains a triangle.*

Proof. Suppose that D does not contain a triangle. By Claim 7.1.3 and Claim 7.1.2 (i) we may assume D is not an isolated vertex of $\Gamma(\mathcal{W}, \mathcal{P})[\lambda]$, that is, D contains an edge. We shall obtain contradicting upper and lower bounds for the number

$$x := \sum_{v \in V(\mathcal{P}_{V(D)})} (d_{G^{\mathcal{P}}_D}(v) - d_{\mathcal{P}_{V(D)}}(v)).$$

For every $\alpha \in V(D)$ denote by V_α the subset of $V(P_\alpha)$ that consists of the ends of P_α and all inner vertices of P_α that have a neighbour in $G^{\mathcal{P}}_D - P_\alpha$. Set $V := \bigcup_{\alpha \in V(D)} V_\alpha$.

For the upper bound let $\alpha\beta$ be an edge of D. Then $N_{\alpha\beta} := N(\alpha) \cap N(\beta) \subseteq \theta$ as a common neighbour of α and β in λ would give rise to a triangle in D. Furthermore, $|N_{\alpha\beta}| \leq 2k - 4$ by Claim 7.1.2 (ii). By Lemma 6.4 the society $(G^{\mathcal{P}}_{\alpha\beta}, P_\alpha P_\beta^{-1})$ is rural if α and β are not twins. But if they are, then $N(\alpha) \cup N(\beta) = N_{\alpha\beta} \cup \{\alpha, \beta\}$. This means that D is a component of $\Gamma(\mathcal{W}, \mathcal{P})[\lambda]$ that consists only of the single edge $\alpha\beta$. So by Claim 7.1.3 we have $|N_{\alpha\beta}| = |N(D)| \geq 2k - 1$, a contradiction. Hence $(G^{\mathcal{P}}_{\alpha\beta}, P_\alpha P_\beta^{-1})$ is rural.

The graph $G - \mathcal{P}_{N_{\alpha\beta}}$ contains $G^{\mathcal{P}}_{\alpha\beta}$ and has minimum degree at least $2k + 3 - |\mathcal{P}_{N_{\alpha\beta}}| \geq 6$ by the connectivity of G. By Claim 7.1.2 (i) we have $|N(\gamma) \cap \theta| \leq 2k - 2$ for every $\gamma \in \lambda$ so Lemma 5.2 implies that every non-trivial $\mathcal{P}$-bridge in an inner bag of $\mathcal{W}$ attaches to at least two paths of $\mathcal{P}_\lambda$ or to none. A vertex v of $G^{\mathcal{P}}_{\alpha\beta} - (P_\alpha \cup P_\beta)$ is therefore an inner vertex of some non-trivial $\mathcal{P}$-bridge B that attaches to P_α and P_β and has all its inner vertices in $G^{\mathcal{P}}_{\alpha\beta}$. This means that a neighbour of v outside $G^{\mathcal{P}}_{\alpha\beta}$ must be an attachment of B on some path P_γ and hence $\gamma \in N_{\alpha\beta} \subseteq \theta$. So all vertices of $G^{\mathcal{P}}_{\alpha\beta} - (P_\alpha \cup P_\beta)$ have the same degree in $G^{\mathcal{P}}_{\alpha\beta}$ as in $G - \mathcal{P}_{N_{\alpha\beta}}$, namely at least 6.

The vertices of $G^{\mathcal{P}}_{\alpha\beta} - (P_\alpha \cup P_\beta)$ retain their degree if we suppress all inner vertices of P_α and P_β that have degree 2 in $G^{\mathcal{P}}_D$. Since the paths of $\mathcal{P}$ are induced by (L6) an inner vertex of P_α has degree 2 in $G^{\mathcal{P}}_D$ if and only if it has no neighbour in $G^{\mathcal{P}}_D - P_\alpha$. So we suppressed precisely those inner vertices of P_α and P_β that are not in V_α or V_β. By Lemma 6.3 the society obtained from $(G^{\mathcal{P}}_{\alpha\beta}, P_\alpha P_\beta^{-1})$ in this way is still rural so Lemma 6.2 implies

$$\sum_{v \in V_\alpha \cup V_\beta} d_{G^{\mathcal{P}}_{\alpha\beta}}(v) \leq 4|V_\alpha| + 4|V_\beta| - 6.$$

Clearly $G^{\mathcal{P}}_D = \bigcup_{\alpha\beta \in E(D)} G^{\mathcal{P}}_{\alpha\beta}$ and $P_\alpha \subseteq G^{\mathcal{P}}_{\alpha\beta}$ for all $\beta \in N_D(\alpha)$ and thus

$$\begin{aligned} x &= \sum_{v \in V} \Big(d_{G^{\mathcal{P}}_D}(v) - d_{\mathcal{P}_{V(D)}}(v) \Big) \\ &\leq \sum_{\alpha \in V(D)} \sum_{\beta \in N_D(\alpha)} \sum_{v \in V_\alpha} (d_{G^{\mathcal{P}}_{\alpha\beta}}(v) - d_{P_\alpha}(v)) \\ &= \sum_{\alpha\beta \in E(D)} \sum_{v \in V_\alpha \cup V_\beta} d_{G^{\mathcal{P}}_{\alpha\beta}}(v) - \sum_{\alpha \in V(D)} |N_D(\alpha)| \cdot (2|V_\alpha| - 2) \\ &\leq \sum_{\alpha\beta \in E(D)} (4|V_\alpha| + 4|V_\beta| - 6) - \sum_{\alpha \in V(D)} |N_D(\alpha)| \cdot (2|V_\alpha| - 2) \\ &= \sum_{\alpha \in V(D)} |N_D(\alpha)| \, (4|V_\alpha| - 3) - \sum_{\alpha \in V(D)} |N_D(\alpha)| \cdot (2|V_\alpha| - 2) \\ &< \sum_{\alpha \in V(D)} 2|N_D(\alpha)| \cdot |V_\alpha|. \end{aligned}$$

To obtain the lower bound for x note that Claim 7.1.4 (ii) says that for any $\alpha \in V(D)$ the vertices of V_α without the two end vertices of P_α have average degree $2 + |N_D(\alpha)|(2 + \varepsilon_\alpha)$ in $G^{\mathcal{P}}_D$ where $\varepsilon_\alpha \geq 1/k$ by Claim 7.1.2 (iii). Clearly every inner bag of $\mathcal{W}$ must contain a vertex of V_α as it contains a $\mathcal{P}$-bridge realising some edge $\alpha\beta \in E(D)$. This means $|V_\alpha| \geq n_1/2 \geq 4k + 2$ and thus

$$\begin{aligned} x &= \sum_{\alpha \in V(D)} \sum_{v \in V_\alpha} \Big(d_{G^{\mathcal{P}}_D}(v) - d_{P_\alpha}(v) \Big) \\ &\geq \sum_{\alpha \in V(D)} (|V_\alpha| - 2) \cdot |N_D(\alpha)| \cdot (2 + \varepsilon_\alpha) \\ &\geq \sum_{\alpha \in V(D)} |N_D(\alpha)| \cdot (2|V_\alpha| - 4 + 4k\varepsilon_\alpha) \\ &\geq \sum_{\alpha \in V(D)} 2|N_D(\alpha)| \cdot |V_\alpha|. \end{aligned}$$

□

Claim 7.1.7. *Every rich block D of $\Gamma(\mathcal{W}, \mathcal{P})[\lambda]$ satisfies $2|D| + |N(D)| \geq 2k + 3$.*

Proof. Suppose for a contradiction that $2|D|+|N(D)| \leq 2k+2$. By Lemma 6.5 there is a $V(D)$-compressed $(\mathcal{P}, V(D))$-relinkage $\mathcal{Q}$ with properties as listed in the statement of Lemma 6.5. Let us first show that we are done if D is *rich w.r.t. to* $\mathcal{Q}$, that is, for every $\alpha \in V(D)$ the inner vertices of Q_α that have a neighbour in $G_\lambda - G_D^{\mathcal{Q}}$ and in $G_D^{\mathcal{Q}} - Q_\alpha$ have average degree at least $2 + |N_D(\alpha)|(2 + \varepsilon_\alpha)$ in $G_D^{\mathcal{Q}}$.

Denote the cut-vertices of $\Gamma(\mathcal{W}, \mathcal{P})[\lambda]$ that lie in D by κ. For $\alpha \in \kappa$ let V_α be the set consisting of the ends of Q_α and of all inner vertices of Q_α that have a neighbour in $G_D^{\mathcal{Q}} - Q_\alpha$ and set $V := \bigcup_{\alpha \in \kappa} V_\alpha$. Pick $\alpha \in \kappa$ such that $|V_\alpha|$ is maximal. By Lemma 5.7 (with $p = 2k+3$) every vertex of $G_D^{\mathcal{Q}}$ lies on a path of $\mathcal{Q}_{V(D)}$ and we have $|Q_\beta| < |V_\alpha|$ for all $\beta \in V(D) \setminus \kappa$.

The paths of $\mathcal{Q}$ are induced in G as $\mathcal{Q}[W]$ is $(2k+3)$-attached in $G[W]$ for every inner bag W of $\mathcal{W}$. Hence V_α contains precisely the vertices of Q_α that are not inner vertices of degree 2 in $G_D^{\mathcal{Q}}$. By the same argument as in the proof of Claim 7.1.4 (ii) the vertices of V_α that are not ends of Q_α have average degree at least $2 + |N_D(\alpha)|(2 + \varepsilon_\alpha)$ in $G_D^{\mathcal{Q}}$.

We want to show that the average degree in $G_D^{\mathcal{Q}}$ taken over all vertices of V_α is larger than $2 + 2|N_D(\alpha)|$. Clearly the end vertices of Q_α have degree at least 1 in $G_D^{\mathcal{Q}}$ so both lack at most $1 + 2|N_D(\alpha)| \leq 3|N_D(\alpha)|$ incident edges to the desired degree. On the other hand, the degree of every vertex of V_α that is not an end of Q_α is on average at least $|N_D(\alpha)| \cdot \varepsilon_\alpha$ larger than desired. But $\varepsilon_\alpha \geq 1/k$ by Claim 7.1.2 (iii) and by Lemma 6.5 (iv) the path $Q_\alpha[W]$ contains a vertex of V_α for every inner bag of $\mathcal{W}$, in particular, $|V_\alpha| \geq n_1/2 > 6k+2$ and hence $(|V_\alpha| - 2)\varepsilon_\alpha > 6$.

This shows that there are more than $2|V_\alpha| \cdot |N_D(\alpha)|$ edges in $G_D^{\mathcal{Q}}$ that have one end on Q_α and the other on another path of $\mathcal{Q}_{V(D)}$. By Lemma 5.1 these edges can only end on paths of $\mathcal{Q}_{N_D(\alpha)}$ so by the pigeon hole principle there is $\beta \in N_D(\alpha)$ such that $G_D^{\mathcal{Q}}$ contains more than $2|V_\alpha|$ edges with one end on Q_α and the other on Q_β.

Hence the society (H, Ω) obtained from $(G_{\alpha\beta}^{\mathcal{Q}}, Q_\alpha Q_\beta^{-1})$ by suppressing all inner vertices of Q_α and Q_β that have degree 2 in $G_{\alpha\beta}^{\mathcal{Q}}$ has more than $2|V_\alpha|+2|V_\beta|-2$ edges and all its $|V_\alpha|+|V_\beta|$ vertices are in $\bar{\Omega}$. So by Lemma 6.2 (H, Ω) cannot be rural. But it is trivially 4-connected as all its vertices are in $\bar{\Omega}$ and must therefore contain a cross by Theorem 6.1. The paths of $\mathcal{Q}$ are induced so this cross consists of two edges which both have one end on Q_α and the other on Q_β. Such a cross gives rise to a linkage $\mathcal{Q}'$ from the left to the right adhesion set of some inner bag W of $\mathcal{W}$ such that the induced permutation of $\mathcal{Q}'$ maps some element of $V(D) \setminus \{\alpha\}$ (not necessarily β) to α and maps every $\gamma \notin V(D)$ to itself. Since α has a neighbour outside D

this is not an automorphism of $\Gamma(\mathcal{W},\mathcal{P})[\lambda]$ and therefore $\mathcal{Q}'$ is a twisting disturbance contradicting the stability of $(\mathcal{W},\mathcal{P})$.

It remains to show that D is rich w.r.t. $\mathcal{Q}$. Suppose that it is not. By the same argument as for Claim 7.1.4 (i) there must be $\alpha \in \kappa$ such that the inner vertices of Q_α that have a neighbour in $G_\lambda - G_D^{\mathcal{Q}}$ and in $G_D^{\mathcal{Q}} - Q_\alpha$ have average degree less than $2 + |N_D(\alpha)|(2+\varepsilon_\alpha)$ in $G_D^{\mathcal{Q}}$. Let (λ_1,λ_2) be a separation of $\Gamma(\mathcal{W},\mathcal{P})[\lambda]$ with $\lambda_1 \cap \lambda_2 = \{\alpha\}$ and $N(\alpha)\cap\lambda_2 = N(\alpha)\cap V(D)$. Let H, Z_1, Z_2, $\mathcal{P}'$, q_1, and q_2 be as in the statement of Lemma 6.5. We shall obtain contradicting upper and lower bounds for the number

$$x := \sum_{v \in V(P'_\alpha \cup Q_\alpha)} \left(d_H(v) - d_{P'_\alpha \cup Q_\alpha}(v)\right).$$

Denote by $H_1,\dots,H_m$ the blocks of H that are not a single edge and for $i = 1,\dots,m$ let V_i be the set of vertices of $C_i := H_i \cap (P'_\alpha \cup Q_\alpha)$ that are a cut-vertex of H or are incident with some edge of H_i that is not in $P'_\alpha \cup Q_\alpha$ and set $V := \bigcup_{i=1}^m V_i$. By definition we have $d_H(v) = d_{P'_\alpha \cup Q_\alpha}(v)$ for all vertices v of $(P'_\alpha \cup Q_\alpha) - V$.

Note that H is adjacent to at most $|N(\alpha)\cap\theta|$ vertices of $\mathcal{P}_\theta$ by Lemma 6.5 (ii) and Lemma 5.1. So Claim 7.1.2 (iii) and the connectivity of G imply that every vertex of H has degree at least $2k+3-|N(\alpha)\cap\theta| \geq 2|N(\alpha)\cap\lambda|+3$ in G_λ.

To obtain an upper bound for x let $i \in \{1,\dots,m\}$. By Lemma 6.5 (vi) C_i is a cycle and the society $(H_i,\Omega(C_i))$ is rural where $\Omega(C_i)$ denotes one of two cyclic permutations that C_i induces on its vertices. Since $|N(\alpha)\cap\lambda| \geq 2$ every vertex of $H_i - C_i$ has degree at least 6 in H_i by the previous paragraph. This remains true if we suppress all vertices of C_i that have degree 2 in H_i. The society obtained in this way is still rural by Lemma 6.3. Since we suppressed precisely those vertices of C_i that are not in V_i Lemma 6.2 implies $\sum_{v\in V_i} d_{H_i}(v) \leq 4|V_i| - 6$. By definition of V we have $d_H(v) = d_{P'_\alpha \cup Q_\alpha}(v)$ for all vertices v of $P'_\alpha \cup Q_\alpha$ that are not in V. Hence we have

$$x = \sum_{v\in V}\left(d_H(v) - d_{P'_\alpha\cup Q_\alpha}(v)\right) = \sum_{i=1}^m \sum_{v\in V_i}\left(d_{H_i}(v) - d_{C_i}(v)\right) \leq \sum_{i=1}^m (2|V_i| - 6).$$

Let us now obtain a lower bound for x. Clearly $G_D^{\mathcal{P}} \subseteq G_{\lambda_2}^{\mathcal{P}}$ and $G_D^{\mathcal{Q}} \subseteq G_{\lambda_2}^{\mathcal{Q}}$. To show that $d_{G_D^{\mathcal{Q}}}(v) = d_{G_{\lambda_2}^{\mathcal{Q}}}(v)$ for all $v \in V(H)$ (we follow the general convention that a vertex has degree 0 in any graph not containing it) it remains to check that an edge of G_λ that has precisely one end in H but is not in $G_D^{\mathcal{Q}}$ cannot be in $G_{\lambda_2}^{\mathcal{Q}}$. Such an edge e must be in a $\mathcal{Q}$-bridge that attaches to Q_α and some Q_β with $\beta \in \lambda \setminus V(D)$. But $N(\alpha) \cap \lambda_2 = V(D)$ and hence $\beta \in \lambda_1$. So e is an edge of $G^{\mathcal{Q}}_{\lambda_1}$ but not on Q_α and therefore not in $G^{\mathcal{Q}}_{\lambda_2}$. This already implies $d_{G^{\mathcal{P}}_D}(v) = d_{G^{\mathcal{P}}_{\lambda_2}}(v)$ for all $v \in V(H)$ since $G^{\mathcal{Q}}_D \subseteq G^{\mathcal{P}}_D$ and $G^{\mathcal{P}}_{\lambda_2} = H \cup G^{\mathcal{Q}}_{\lambda_2}$ (see the proof of Lemma 6.5 (i) for the latter identity). The next equality follows directly from the definition of H.

$$d_H(v) + d_{G^{\mathcal{Q}}_{\lambda_2}}(v) = d_{G^{\mathcal{P}}_{\lambda_2}}(v) + d_{Q_\alpha}(v) \quad \forall v \in V(H).$$

Denote by U_1 the set of inner vertices of P_α that have a neighbour in both $G_\lambda - G^{\mathcal{P}}_D$ and $G^{\mathcal{P}}_D - P_\alpha$ and by U_2 the set of inner vertices of Q_α that have a neighbour in both $G_\lambda - G^{\mathcal{Q}}_D$ and $G^{\mathcal{Q}}_D - Q_\alpha$. In other words, U_1 and U_2 are the sets of those vertices of P_α and Q_α, respectively, that are relevant for the richness of α in D. Set $V' := (V \setminus \{q_1, q_2\}) \cup (Z_1 \cap Z_2))$, $V_P := V' \cap V(P'_\alpha)$, and $V_Q := V' \cap V(Q_\alpha)$. Then $U_1 = (V \cap Z_1) \cup (Z_1 \cap Z_2) = V' \cap Z_1 \subseteq V_P$ and $U_2 = (V \cap Z_2) \cup (Z_1 \cap Z_2) \subseteq V_Q$.

By our earlier observation every vertex of H has degree at least $2|N(\alpha) \cap \lambda| + 3$ in G_λ and therefore every vertex of $V_P \setminus Z_1$ must have at least this degree in $G^{\mathcal{P}}_{\lambda_2}$. Since $U_1 \subseteq V_P$ and α is rich in D this means that

$$\sum_{v \in V_P} d_{G^{\mathcal{P}}_D}(v) \geq |V_P| \left(2 + |N_D(\alpha)| \cdot (2 + \varepsilon_\alpha)\right).$$

Similarly, we have $U_2 \subseteq V_Q \subseteq V(Q_\alpha)$ and every vertex $v \in V_Q \setminus Z_2$ satisfies $d_{G^{\mathcal{Q}}_D}(v) = 2 = d_{Q_\alpha}(v)$. So by the assumption that α is not rich in D w.r.t. $\mathcal{Q}$ we have

$$\sum_{v \in V_Q} \left(d_{G^{\mathcal{Q}}_D}(v) - d_{Q_\alpha}(v)\right) < |V_Q| \cdot |N_D(\alpha)| \cdot (2 + \varepsilon_\alpha).$$

Observe that

$$2|N(\alpha) \cap \lambda| + 3 = 2 + |N(\alpha) \cap \lambda_1| \cdot (2 + \varepsilon_\alpha) + |N(\alpha) \cap \lambda_2| \cdot (2 + \varepsilon_\alpha)$$

and recall that $N_D(\alpha) = N(\alpha) \cap \lambda_2$. Combining all of the above we get

$$
\begin{aligned}
x &\geq \sum_{v \in V'} \big(d_H(v) - d_{P'_\alpha \cup Q_\alpha}(v)\big) \\
&= \sum_{v \in V'} \Big(d_{G_D^{\mathcal{P}}}(v) - d_{G_D^{\mathcal{Q}}}(v) + d_{Q_\alpha}(v) - d_{P'_\alpha \cup Q_\alpha}(v)\Big) \\
&= \sum_{v \in V_P} d_{G_D^{\mathcal{P}}}(v) + \sum_{v \in V' \setminus V_P} d_{G_D^{\mathcal{P}}}(v) - \sum_{v \in V_Q} \Big(d_{G_D^{\mathcal{Q}}}(v) - d_{Q_\alpha}(v)\Big) - 2|V'| - 2m \\
&> |V_P| \cdot |N_D(\alpha)| \cdot (2 + \varepsilon_\alpha) + 2|V_P| + |V' \setminus V_P| \cdot (2|N(\alpha) \cap \lambda| + 3) \\
&\quad - |V_Q| \cdot |N_D(\alpha)| \cdot (2 + \varepsilon_\alpha) - 2|V'| - 2m \\
&= |V' \setminus V_Q| \cdot |N_D(\alpha)| \cdot (2 + \varepsilon_\alpha) + |V' \setminus V_P| \cdot |N(\alpha) \cap \lambda_1| \cdot (2 + \varepsilon_\alpha) - 2m \\
&> 2|V' \setminus V_Q| + 2|V' \setminus V_P| - 2m = \sum_{i=1}^{m} (2|V_i| - 6)
\end{aligned}
$$

This shows that D is rich w.r.t. $\mathcal{Q}$ as defined above. So Claim 7.1.7 holds. □

By Claim 7.1.1 the graph $\Gamma(\mathcal{W}, \mathcal{P})[\lambda]$ has a component. This component has a rich block D by Claim 7.1.5. By Claim 7.1.6 and Claim 7.1.7 we have a triangle in D and $|D| + |N(D)| \geq 2k + 3$. This contradicts Claim 7.1.2 (iv) and thus concludes the proof of Theorem 1.1. □

8 Discussion

In this section we first show that Theorem 1.1 is almost best possible (see Proposition 8.1 below) and then summarise where our proof uses the requirement that the graph G is $(2k+3)$-connected.

Proposition 8.1. *For all integers k and N with $k \geq 2$ there is a graph G which is not k-linked such that*

$$\kappa(G) \geq 2k + 1, \qquad \mathrm{tw}(G) \leq 2k + 10, \quad \textit{and} \quad |G| \geq N.$$

Proof. We reduce the assertion to the case $k = 2$, that is, to the claim that there is a graph H which is not 2-linked but satisfies

$$\kappa(H) = 5, \qquad \mathrm{tw}(H) \leq 14, \quad \text{and} \quad |H| \geq N.$$

For any $k \geq 3$ let K be the graph with $2k - 4$ vertices and no edges. We claim that $G := H * K$ (the disjoint union of H and K where every vertex of H is joined to every vertex of K by an edge) satisfies the assertion for k.

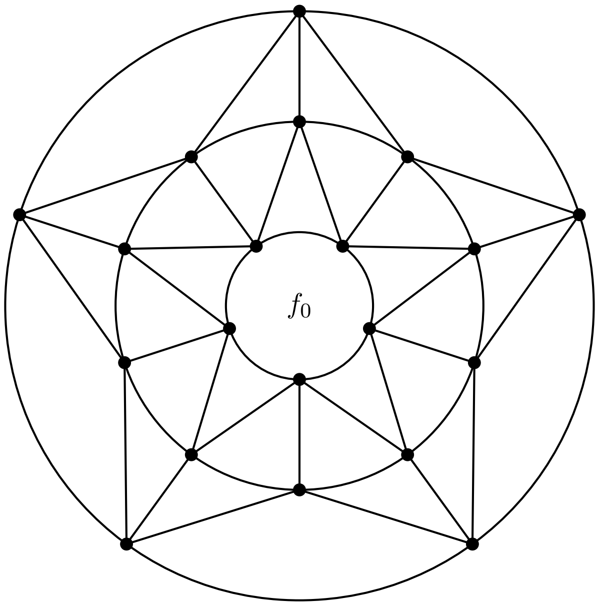

Figure 2: The 5-connected graph H_0 and its inner face f_0.

Clearly $|G| = |H| + 2k - 4 \geq N$. Taking a tree-decomposition of H of minimal width and adding $V(K)$ to every bag gives a tree-decomposition of G, so $\mathrm{tw}(G) \leq \mathrm{tw}(H) + 2k - 4 \leq 2k + 10$. To see that G is $(2k+1)$-connected, note that it contains the complete bipartite graph with partition classes $V(H)$ and $V(K)$, so any separator X of G must contain $V(H)$ or $V(K)$. In the former case we have $|X| \geq N$ and we may assume that this is larger than $2k$. In the latter case we know that $G - X \subseteq H$, in particular $X \cap V(H)$ is a separator of H and hence must have size at least 5, implying $|X| \geq |K| + 5 = 2k + 1$ as required.

Finally, G is not k-linked: By assumption there are vertices s_1, s_2, t_1, t_2 of H such that H does not contain disjoint paths P_1 and P_2 where P_i ends in s_i and t_i for $i = 1, 2$. If G was k-linked, then for any enumeration $s_3, \ldots, s_k, t_3, \ldots, t_k$ of the $2k - 4$ vertices of $V(K)$ there were disjoint paths $P_1, \ldots, P_k$ in G such that P_i has end vertices s_i and t_i for $i = 1, \ldots, k$. In particular, P_1 and P_2 do not contain a vertex of K and are hence contained in H, a contradiction.

It remains to give a counterexample for $k = 2$. The planar graph H_0 in Figure 2 is 5-connected. Denote the 5-cycle bounding the outer face of H_0 by

C_1 and the 5-cycle bounding f_0 by C_0. Then $(V(H_0 - C_0), V(H_0 - C_1))$ forms a separation of H_0 of order 10, in particular, H_0 has a tree-decomposition of width 14 where the tree is K_2. Draw a copy H_1 of H_0 into f_0 such that the cycle C_0 of H_0 gets identified with the copy of C_1 in H_1. Since $H_0 \cap H_1$ has 5 vertices, the resulting graph is still 5-connected and has a tree-decomposition of width 14. We iteratively paste copies of H_0 into the face f_0 of the previously pasted copy as above until we end up with a planar graph H such that

$$\kappa(H) = 5, \qquad \text{tw}(H) \leq 14, \quad \text{and} \quad |H| \geq N.$$

Still the outer face of H is bounded by a 5-cycle C_1, so we can pick vertices s_1, s_2, t_1, t_2 in this order on C_1 to witness that H is not 2-linked (any s_1–t_1 path must meet any s_2–t_2 path by planarity). □

Where would our proof of Theorem 1.1 fail for a $(2k+2)$-connected graph G? There are several instances where we invoke $(2k+3)$-connectivity as a substitute for a minimum degree of at least $2k+3$. The only place where minimum degree $2k+2$ does not suffice is the proof of Claim 7.1.4. We need minimum degree $2k+3$ there to get the small "bonus" ε_α in our notion of richness. Richness only allows us to make a statement about the inner vertices of a path and the purpose of this bonus is to compensate for the end vertices. Therefore the arguments involving richness in the proofs of Claim 7.1.6 and Claim 7.1.7 would break down if we only had minimum degree $2k+2$.

But even if the suppose that G has minimum degree at least $2k+3$ there are still two places where our proof of Theorem 1.1 fails: The first is the proof of Claim 7.1.3 and the second is the application of Lemma 5.7 in the proof of Claim 7.1.7.

We use Claim 7.1.3 in the proof of Claim 7.1.6, to show that no component of $\Gamma(\mathcal{W}, \mathcal{P})[\lambda]$ can be a single vertex or a single edge. In both cases we do not use the full strength of Claim 7.1.3. So although we formally rely on $(2k+3)$-connectivity for Claim 7.1.3 we do not really need it here.

However, the application of Lemma 5.7 in the proof of Claim 7.1.7 does need $(2k+3)$-connectivity. Our aim there is to obtain a contradiction to Claim 7.1.2 (iv) which inherits the bound $2k+3$ from the token game in Lemma 4.10. This bound is sharp: Let H be the union of a triangle $D = d_1d_2d_3$ and two edges d_1a_1 and d_2a_2 and set $A := \{a_1, a_2\}$. Clearly $H - A = D$ is connected and A is marginal in H. For $k = 3$ we have $2|D| + |N(D)| = 8 = 2k+2$. Let L be the pairing with edges $(a_1, 0)(a_2, 0)$ and $(d_i, 0)(d_i, \infty)$ for $i = 1, 2$. It is not hard to see that there is no L-movement on H as the two tokens from A can never meet.

So the best hope of tweaking our proof of Theorem 1.1 to work for $(2k+2)$-connected graphs is to provide a different proof for Claim 7.1.7. This would also be a chance to avoid relinkages, that is, most of Section 5, and the very technical Lemma 6.5 altogether as they only serve to establish Claim 7.1.7.

References

[1] T. Böhme, K. Kawarabayashi, J. Maharry, and B. Mohar. Linear Connectivity Forces Large Complete Bipartite Minors. *J. Combin. Theory B* **99** (2009), 557–582.

[2] T. Böhme, J. Maharry, and B. Mohar. $K_{a,k}$ Minors in Graphs of Bounded Tree-Width. *J. Combin. Theory B* **86** (2002), 133–147.

[3] B. Bollobás and A. Thomason. Highly Linked Graphs. *Combinatorica* **16** (1996), 313–320.

[4] R. Diestel. *Graph Theory.* Third edition. Graduate Texts in Mathematics, 173. Springer, Heidelberg, 2005.

[5] R. Diestel, K. Kawarabayashi, T. Müller, and P. Wollan. On the excluded minor structure theorem for graphs of large tree width. *J. Combin. Theory B* **102** (2012), 1189–1210.

[6] W. W. Johnson. Notes on the "15" Puzzle. I. *Amer. J. Math.* **2** (1879), 397–399.

[7] L. Jørgensen. Contraction to K_8. *J. Graph Theory* **18** (1994), 431–448.

[8] H. A. Jung. Eine Verallgemeinerung des n-fachen Zusammenhangs für Graphen. *Math. Ann.* **187** (1970), 95–103.

[9] K. Kawarabayashi, S. Norine, R. Thomas, and P. Wollan. K_6 minors in 6-connected graphs of bounded tree-width. *arXiv*:1203.2171 [math.CO] (2012).

[10] K. Kawarabayashi, S. Norine, R. Thomas, and P. Wollan. K_6 minors in large 6-connected graphs. *arXiv*:1203.2192 [math.CO] (2012).

[11] A. Kostochka. The minimum Hadwiger number for graphs with a given mean degree of vertices. (Russian). *Metody Diskret. Analiz.* **38** (1982), 37–58.

[12] D. G. Larman and P. Mani. On the Existence of Certain Configurations Within Graphs and the 1-Skeletons of Polytopes. *Proc. London Math. Soc.* **20** (1974), 144–160.

[13] B. Oporowski, J. Oxley, and R. Thomas. Typical Subgraphs of 3- and 4-connected graphs. *J. Combin. Theory B* **57** (1993), 239–257.

[14] N. Robertson and P. D. Seymour. Graph Minors. IX. Disjoint Crossed Paths. *J. Combin. Theory B* **49** (1990), 40–77.

[15] N. Robertson and P.D. Seymour. Graph Minors XIII. The Disjoint Paths Problem. *J. Combin. Theory B* **63** (1995), 65–110.

[16] N. Robertson and P.D. Seymour. Graph Minors. XVI. Excluding a non-planar graph. *J. Combin. Theory B* **89** (2003), 43–76.

[17] R. Thomas. Configurations in Large t-connected Graphs [presentation]. *SIAM Conference on Discrete Mathematics*, Austin, Texas, USA, June 14–17 2010. <`https://client.blueskybroadcast.com/siam10/DM/SIAM_DM_07_IP6/slides/Thomas_IP6_6_16_10_2pm.pdf`> (Accessed 2014-02-21).

[18] R. Thomas and P. Wollan. An improved linear edge bound for graph linkages. *Europ. J. Comb.* **26** (2005), 309–324.

[19] A. Thomason. An Extremal Function for Complete Subgraphs. *Math. Proc. Camb. Phil. Soc.* **95** (1984), 261–265.

[20] C. Thomassen. 2-linked graphs. *Europ. J. Comb.* **1** (1980), 371–378.

[21] C. Thomassen. *personal communication* (2014).

[22] R. M. Wilson. Graph Puzzles, Homotopy, and the Alternating Group. *J. Combin. Theory B* **16** (1974), 86–96.

Linkages in Large Graphs

Jan-Oliver Fröhlich* Ken-ichi Kawarabayashi†‡
Theodor Müller§ Julian Pott¶ Paul Wollan‖

Abstract

We prove that every large enough $(2k+3)$-connected graph G is k-linked. In an earlier paper [1] we proved this for graphs of bounded tree-width. We complete the unbounded tree-width case by finding a large linear decomposition of some subgraph neatly contained in G to which we can apply our earlier result.

1 Introduction

This article is the second part of our two part series that shows that all large enough $(2k+3)$-connected graphs are k-linked. We refer to the first part [1] as Part **I** and as we often mention particular passages of Part **I** we use **I**–X to refer to X in Part **I**. For a thorough introduction see Section **I**–1. We restate the main result of Part **I**.

Theorem 1.1 (Theorem **I**–1.1)**.** *For all integers k and w, there exists an integer N such that a graph G is k-linked if*

$$\kappa(G) \geq 2k+3, \qquad \mathrm{tw}(G) \leq w, \quad \text{and} \quad |G| \geq N.$$

In this article we prove the following theorem.

*Fachbereich Mathematik, Universität Hamburg, Hamburg, Germany. jan-oliver.froehlich@math.uni-hamburg.de
†National Institute of Informatics, Tokyo, Japan.
‡JST, ERATO, Kawarabayashi Large Graph Project.
§Fachbereich Mathematik, Universität Hamburg, Hamburg, Germany.
¶Fachbereich Mathematik, Universität Hamburg, Hamburg, Germany.
‖Department of Computer Science, University of Rome "La Sapienza", Rome, Italy. wollan@di.uniroma1.it. Research supported by the European Research Council under the European Unions Seventh Framework Programme (FP7/2007 - 2013)/ERC Grant Agreement no. 279558.

Theorem 1.2. *For all integers k there is an integer w such that all $(2k+3)$-connected graphs with tree-width at least w are k-linked.*

Together, Theorem **I**–1.1 and Theorem 1.2 yield our main result.

Theorem 1.3. *For all integers k there is an integer N such that all $(2k+3)$-connected graphs on at least N vertices are k-linked.*

Proof. Let k be given and let w be the tree-width returned by Theorem 1.2. Let N be the integer returned by Theorem 1.1 for k and w. Let G be a $(2k+3)$-connected graph on at least N vertices. If G has tree-width w it is k-linked by Theorem 1.2. If the tree-width of G is less than w it is k-linked by Theorem 1.1. □

2 Excluded substructures

In this section we collect structures such that a sufficiently connected graph containing any of these structures is k-linked. In particular, it suffices to contain a subdivided copy of $K_{2k,2k}$, a subgraph isomorphic to $K_{2k-2,2k}$, or a K_{3k} as a minor for a $(2k+3)$-conected graph to be k-linked. We also show this for a slightly more complicated structure that is introduced later.

First we have Lemma **I**–7.1:

Lemma 2.1. *Every $2k$-connected graph containing a $TK_{2k,2k}$ is k-linked.* □

Lemma 2.2. *Every $(2k+2)$-connected graph containing a subgraph isomorphic to $K_{2k-2,2k}$ is k-linked.*

Proof. In the considered graph G let A and B be disjoint vertex sets with $|A| = 2k-2$ and $|B| = 2k$ such that every vertex in A is adjacent to every vertex in B. Let $X = \{x_1, \ldots, x_k\}$ and $Y = \{y_1, \ldots, y_k\}$ be two disjoint sets, each containing k distinct vertices of G. We need to find a system of k disjoint X–Y paths linking x_i to y_i for $i = 1, \ldots, k$.

Let Z be a maximal subset of $A \setminus (X \cup Y)$ of size at most 2 (note that Z may be empty). By the connectedness of G, we find a system $\mathcal{P}$ of $2k$ disjoint $(X \cup Y)$–$(A \cup B)$ paths in $G - Z$. For each $x \in X \cup Y$ denote by P_x the unique path of $\mathcal{P}$ starting in x. Let I_A and I_B be the sets of indices $i \in \{1, \ldots, k\}$ such that both of P_{x_i} and P_{y_i} end in A and both end in B, respectively, and denote the remaining indices by J. Let A_1 and B_1 the sets of vertices of A and B, respectively, in which a path of $\mathcal{P}$ ends. Set $A_0 := A \setminus A_1$ and $B_0 := B \setminus B_1$.

For every $i \in J$ there is an edge e_i between the last vertices of P_{x_i} and P_{y_i} so $L_i := (P_{x_i} \cup P_{y_i}) + e_i$ is an x_i–y_i path. Clearly

$$|B_0| = 2k - |B_1| = |A_1| = 2|I_A| + |J| \geq |I_A|$$

So we can pick distinct vertices $(v_i)_{i \in I_A}$ in B_0. For every $i \in I_A$, the last vertices of P_{x_i} and P_{y_i} are adjacent to v_i, by edges e_i and f_i say, so $L_i := P_{x_i} \cup e_i \cup f_i \cup P_{y_i})$ is an x_i–y_i path.

If $|A_0| \geq |I_B|$ holds as well, then doing the symmetric construction for $i \in I_B$ yields an x_i–y_i path for every such i. So for each $i = 1, \dots, k$, the path L_i is comprised of two paths of $\mathcal{P}$ and at most one additional vertex unique to i hence the constructed paths are disjoint.

So we may assume that $|A_0| < |I_B|$. This implies

$$|A_0| = 2k - 2 - |A_1| = |B_1| - 2 = 2|I_B| + |J| - 2 \geq 2|A_0| + |J|$$

and therefore $|A_0| = |J| = 0$. In particular, $|I_A| = k - 1$, $|I_B| = 1$, and $Z = \emptyset$. The latter implies that $A \subseteq X \cup Y$, so we may assume that x_1 and y_1 are not in A but all other elements of $X \cup Y$ are. As $G - A$ is 2-connected[1], it contains two internally disjoint x_1–y_1 paths. One of these paths misses at least $k - 1$ vertices, $z_2, \dots, z_k$ say, of B and we denote it by L_1. For $2 \leq i \leq k$ let L_i be the x_i–y_i path $x_i z_i y_i$ of length 2. Then $\{L_i \mid 1 \leq i \leq k\}$ is an X–Y linkage as desired. □

For positive integers $r \geq 3$, define a graph H_r as follows (see Figure 1). Let $P_1, \dots, P_r$ be r disjoint ('horizontal') paths of length $r - 1$, say $P_i = v_1^i \dots v_r^i$. Let $V(H_r) = \bigcup_{i=1}^r V(P_i)$, and let

$$E(H_r) = \bigcup_{i=1}^{r} E(P_i) \cup \left\{ v_j^i v_j^{i+1} \mid i, j \text{ odd};\ 1 \leq i < r;\ 1 \leq j \leq r \right\}$$
$$\cup \left\{ v_j^i v_j^{i+1} \mid i, j \text{ even};\ 1 \leq i < r;\ 1 \leq j \leq r \right\}.$$

We call the paths P_i the *rows* of H_r; the paths induced by the vertices $\{v_j^i, v_{j+1}^i : 1 \leq i \leq r\}$ for an odd index j are its *columns*.

The 6-cycles in H_r are its *bricks*. In the natural plane embedding of H_r, these bound faces of H. The outer cycle of the unique maximal 2-connected subgraph of H_r is the *boundary cycle* of H_r. The *diagonal* of H_r is the set of vertices $\{v_i^i \mid 1 \leq i \leq r\}$. A set $X \subseteq V(H)$ is a *pseudo diagonal* if there is

[1] In fact, $G - A$ is 4-connected.

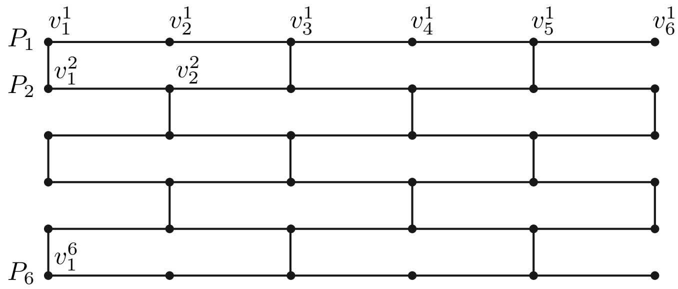

Figure 1: The graph H_6

at most one vertex in each row and each column while for any two vertices $v_i^j, v_{i'}^{j'} \in X$ it holds that $i < i'$ implies $j < j'$.

Any subdivision H of the unique maximal 2-connected subgraph of H_r will be called an *r-wall*, or a *wall of size r* and with the *branch vertices* of H we refer to its vertices of degree 3, these are the branch vertices of the defining subdivision, except the ones of degree 2 on the boundary cycle. The *bricks* and the *boundary cycle* of H are its subgraphs that form subdivisions of the bricks and the boundary cycle of H_r, respectively. The *diagonal* of H are the vertices in the diagonal of the defining H_r and a set X is a *pseudo diagonal* of H if it corresponds to a pseudo diagonal in the appropriate H_r. An embedding of H in a surface Σ is a *flat* embedding, and H is *flat* in Σ, if the boundary cycle C of H bounds an open disc $D(H)$ in Σ such that all its bricks B_i bound disjoint, open discs $D(B_i)$ in Σ with $D(B_i) \subseteq D(H)$ for all i.

A plane graph G is called *face restricted* 2*-linked* if for every two pairs $S = (s_1, s_2)$ and $T = (t_1, t_2)$ of vertices such that no face of G contains all four vertices in its boundary there is an S–T linkage. A planar graph is *face restricted* 2*-linked* if all its possible embeddings into the plane are.

Lemma 2.3. *Any* 3*-connected, planar,* 3*-regular graph is face restricted* 2*-linked.*

Proof. First note that all embeddings of a 3-connected planar graph G are equivalent. That is for every embedding the same sets of vertices appear on the boundary of a face in the same order. Let $S = (s_1, s_2)$, and $T = (t_1, t_2)$ be two pairs of vertices such that no face of some embedding of G contains all four vertices in its boundary. It suffices to show that there is an S–T linkage. As all embeddings of a 3-connected planar graph G are equivalent, there is no embedding of G with these four vertices on the boundary of the outer face (in any order). In terms of [9] we have that $(G,(s_1,t_1,s_2,t_2))$ is not rural. As G is 3-connected, we can apply [9, 2.4] to $(G,(s_1,t_1,s_2,t_2))$ and either get a tripod with paths P_1, P_2, P_3 say or a cross. The cross is the desired S–T linkage and thus me may assume that we indeed have the tripod.

Any embedding σ of G into the plane embedds the tripod and the embedded tripod has precisely three faces F_{ij} with $1 \leq i < j \leq 3$, each with the union of the paths $\sigma(P_i)$ and $\sigma(P_j)$ as its boundary. As each of the three paths lies on the boundary of a face of the embedded tripod that contains a vertex from $\sigma(\{s_1,t_1,s_2,t_2\})$, it is clear that none of these three faces contains all four vertices. We only describe how to obtain a cross from the tripod for the case that there is one face, say F_{12}, containing three of the four vertices, say $\sigma(s_1), \sigma(t_1), \sigma(s_2)$. The other cases can be solved similarly.

Let L be 3 disjoint $\{s_1,t_1,s_2\}$–$(P_1 \cup P_2)$ paths in G. As G is 3-regular L is disjoint from $P_1 \cap P_2$ as the two vertices in $P_1 \cap P_2$ both already have three neighbours in $P_1 \cup P_2 \cup P_3$—one in each path. We may assume that two of the paths in L end on P_1 and one ends on P_2.[2] By the definition of a tripod, there is a t_2–P_3 path P avoiding $P_1 \cup P_2$. It is immediate that there is an S–T linkage independently from how L and P particularly attach to the tripod. □

Lemma 2.4. *Let H be an r-wall with branch vertices B. Let $a,b,c \in B$ be distinct branch vertices, and let $D \subseteq B$ be a pseudo diagonal of H of cardinality at least* 31*. Let $F \subseteq E(H)$ be an edge set satisfying one of the following statements.*

(i) $F = \emptyset$.

(ii) F is the edge set of a maximal B-path and both end-vertices of this path are in $\{a,b,c\}$.

Then there is a vertex $d \in D$ such that there is an (a,c)–(b,d) linkage in $H - F$.

Proof. It is sufficent to prove this lemma after surpressing all vertices of degree 2 in H and thus we assume that H is 3-regular. As the proof for the case that $F = \emptyset$ is easier we give here only the proof for if F satisfies (ii). For the following conclusion note that there is no edge in H whose endvertices

[2]For $i = 1,2$ there is a path from P_i to $\{s_1,t_1,s_2\}$ as required by the definition of a tripod and we can follows this path to its first vertex on L.

have two common neighbours. $H - F$ has precisely two vertices of degree 2 and after surpressing these we obtain a graph H' that is 3-connected, planar, and 3-regular and thus by Lemma 2.3 face restricted 2-linked. Let x, y be the endvertices of the edge in F and let x', y' be a choice of the two neighbours of x and y in $H - F$ respectively such that there are 5 distinct vertices in $X := \{a, b, c, x', y'\}$. Such a choice is possible as no two vertices have two common neighbours in H. There is a vertex $d \in D$ that does not share a face boundary in any embedding of H with any of the vertices in X as every vertex of H lies in the boundary of at most 3 faces, no face contains more than 2 vertices from D in its boundary and $|D| \geq 31 > 6|X|$. By the face restricted 2-linkedness of H' there is for every choice of 3 vertices $v_1, v_2, v_3 \in X - \{x, y\}$ a (v_1, v_3)–(v_2, d) linkage in $H - F$. For the right choice of v_1, v_2, v_3 this can be extended to an (a, c)–(b, d) linkage in H. □

In [10, Section 7] Robertson and Seymour show that one can remove the vertical edges of four consecutive columns and every other edge of two horizontal rows to obtain from a wall H of size r one of size $r - 2$. As their and our definition of size differ slightly, we only need to remove the vertical edges of two consecutive columns. We denote the graph that is obtained in this way by $H - [a]$ if a is one of the vertices that is not contained any longer in the subwall Note that we can remove any vertex, branch vertex or not, in this way from the wall and that whenever D is a pseudo diagonal of H, then we need to remove at most 6 vertices from D to obtain a pseudo diagonal of $H - [a]$.[3] We write $H - [a, b, c]$ to denote the subwall where a, b, c are removed in this way one after another.

Lemma 2.5. *Let $k \geq 1$ and let G be a 3-connected graph containing a wall H with branch vertices B some pseudo diagonal $D \subseteq B$ of cardinality at least $6k + 31$. For any set of three vertices $\{a, b, c\} \subseteq V(G - H)$ there is a vertex $d \in D$ and an (a, c)–(b, d) linkage that avoids at least k vertices from D.*

Proof. For a graph G as in the statement of the lemma let L be an $\{a, b, c\}$–B linkage that has as few edges outside of H as possbile. Let L link a and a', b and b', and c and c'. By the choice of L if two paths from L intersect with a subdivided edge of H, then both end vertices of that subdivided edge, will be end vertices of paths form L. Let F be the set of maximal B-paths in H such that for each path in F the set of its inner vertices

[3] With our definition of size and only removing edges in two consecutive columns it would be sufficient to remove 4 vertices from D. We keep the 6 here to remove a potential stumbling block if you take out edges of 4 consecutive rows in your mental calculations.

has non-empty intersection with each path from L. We have $|F| \leq 1$. As $|D| \geq 6k + 18$ and the three vertices in $B \cap L$ disallow at most 18 vertices from any pseudo diagonal there is a choice of k vertices $d_1, \ldots, d_k \in D$ such that $H' = H - [d_1, \ldots, d_k]$ still contains $B \cap L$ as branch vertices. As we need to remove at most 6 vertices from a pseudo diagonal when removing one vertex d_i from H we have a pseudo digaonel $D' \subseteq D$ of H' with cardinality at least 31. Let B' be the branch vertices of H' and let F' be the maximal B'-paths in H' that contain any edge from F. We have $|F'| \leq 1$, as each B-path in F is either not contained in H' or a subpath of some B'-path in H'. A Straight forward application of Lemma 2.4 for the wall H', the branch vertices a', b', c', and the edge set of the maximal B'-path F' yields the desired linkage of G since all the maximal B'-paths that are not in F' can be used to extend L if necessary. □

Lemma 2.6. *Let G be a $(2k+3)$-connected graph that contains a wall H. If G contains a complete bipartite subgraph with bipartition (A, D) such that $|A| = 2k - 3$ and A is disjoint from H, while $|D| \geq 6k + 31$ and D is a pseudo diagonal of H, then G is k-linked.*

Proof. For a graph G with H, A, and D as in the statement, let $X = \{x_1, \ldots, x_k\}$ and $Y = \{y_1, \ldots, y_k\}$ be two disjoint sets, each containing k distinct vertices of G. We need to find a system of k disjoint X–Y paths linking x_i to y_i for $i = 1, \ldots, k$.

We will complete this proof separately for the following three cases.

Case 1: $|A \setminus (X \cup Y)| = 0$

Case 2: $|A \setminus (X \cup Y)| = 1$

Case 3: $|A \setminus (X \cup Y)| \geq 2$

Case 1: As $|A \setminus (X \cup Y)| = 0$ we have $X \cup Y \setminus A = \{a, b, c\}$ for some vertices a, b, c. If $a = x_1$ and $b = y_1$ we apply Lemma 2.5 to the 6-connected graph $G - A$, the wall H and the subset D of the diagonal of H of size at least $6k + 31$. The resulting linkage L avoids k vertices from D and thus can be extended to an X–Y linkage as for each pair (x_i, y_i) that remains to be paired there is a unique vertex in D that L avoids. If a, b, c are not linked with each other, then any set of 3 disjoint $\{a, b, c\}$–D paths in $G - A$ can in the same way be extended to an X–Y linkage.

Case 2: As $|A \setminus (X \cup Y)| = 1$ we have $X \cup Y \setminus A = \{a, b, c, d\}$ for some vertices a, b, c, d. Let v be the unique vertex in $A \setminus (X \cup Y)|$. If $a = x_1, b = y_1$ and $c = x_2, d = y_2$ we pick two disjoint paths P, Q in $G - A$ where P is a

a–b path and Q is a c–d path. These paths exist as $G - A$ is 6-connected and thus 2-linked [5]. If one of P or Q meets D in at least $k+2$ vertices, say P, let d_1 be the first vertex on P in D and let d_2 be the last vertex on P in B. Then the linkage $\{aPd_1vd_2Pb, Q\}$ can easily, as in Case 1, be extended to our desired X–Y linkage. If neither P nor Q meets B in at least $k+2$ vertices, then there are $|D| - 2k - 4 \geq 2k + 46 \geq k - 2$ vertices disjoint from $P \cup Q$ in D and we extend $\{P, Q\}$ to an X–Y linkage using the $|D| - 2k - 4 \geq k - 2$ vertices to link the remaining pairs (x_i, y_i) in A. If only one or none of the possible pairs of a, b, c, d needs to be linked, let P, Q, R, S be 4 disjoint paths from $\{a, b, c, d\}$ to D in $G - A$. Again it is easy to extend these to an appropriate X–Y linkage.

Case 3: If $|A \setminus (X \cup Y)|$ contains at least 2 vertices, then there is enough space to link up everything without the need of special structures. Let $Z \subseteq A \setminus (X \cup Y)$ have cardinality two. Let $\mathcal{P}$ be a linkage of size $2k$ from $X \cup Y$ to $(A \setminus Z) \cup D$ in $G - Z$. It is easy to join up the paths from $\mathcal{P}$ to obtain the desired X–Y linkage. □

The next theorem follows[4] from (5.4) in [10].

Theorem 2.7 (Robertson & Seymour). *Every $2k$-connected graph that has K_{3k} as a minor is k-linked.*

3 Near Embeddings and Path-Decompositions

As we could identify 'containing a K_{3k}-minor' as a sufficient condition for a $(2k+3)$-connected graph to be k-linked, we can apply the structure theorem for graphs with forbidden minors [4]. In the following we give a brief overview of this theorem.

3.1 The excluded minor structure theorem

We use the same introduction to near embeddings as in [4] but leave out some minor parts that are irrelevant in our context.

A *vortex* is a pair $V = (G, \Omega)$, where G is a graph and $\Omega =: \Omega(V)$ is a linearly ordered set $(w_1, \ldots, w_n)$ of vertices in G. These vertices are the *society vertices* of the vortex; their number n is its *length*. We do not always distinguish notationally between a vortex and its underlying graph or the vertex set of that graph; for example, a *subgraph of V* is just a subgraph

[4]To see this, set $\zeta = 2n$, $k = 3n$, $Z = \{s_1, \ldots, s_n, t_1, \ldots, t_n\}$, and $Z_i = \{s_i, t_i\}$ for $i = 1, \ldots, n$ (note that our k is their n).

of G, a *subset of* V is a subset of $V(G)$, and so on. Also, we will often use Ω to refer to the linear order of the vertices $w_1, \dots, w_n$ as well as the set of vertices $\{w_1, \dots, w_n\}$.

A path-decomposition $\mathcal{D} = (X_1, \dots, X_m)$ of G is a *decomposition* of our vortex V if $m = n$ and $w_i \in X_i$ for all i. For any subset $\Omega' = \{w_{n_1}, \dots, w_{n_l}\}$ of Ω with $w_{n_1} = w_1$ and with $w_{n_i} \leq_\Omega w_{n_{i+1}}$ for all $i < l$, the path-decomposition

$$\left(\bigcup_{i=n_1}^{n_2-1} X_i, \bigcup_{i=n_2}^{n_3-1} X_i, \dots, \bigcup_{i=n_l}^{n} X_i \right)$$

is the *decomposition of* (G, Ω') *induced by* $\mathcal{D}$. When $n > 1$, the *adhesion* of our decomposition $\mathcal{D}$ of V is the maximum value of $|X_i \cap X_{i+1}|$, taken over all $1 \leq i < n$. We define the *adhesion* of a vortex V as the minimum adhesion of any decomposition of that vortex.

We write $Z_i := (X_i \cap X_{i+1}) \setminus \Omega$ for all $1 \leq i < n$, when $\mathcal{D}$ is a decomposition of a vortex V as above. These Z_i are the *adhesion sets* of $\mathcal{D}$. We call $\mathcal{D}$ *linked* if

- all these Z_i have the same size;
- there are $|Z_i|$ disjoint Z_{i-1}–Z_i paths in $G[X_i] - \Omega$, for all $1 < i < n$;
- $X_i \cap \Omega = \{w_{i-1}, w_i\}$ for all $1 \leq i \leq n$, where $w_0 := w_1$.

Note that $X_i \cap X_{i+1} = Z_i \cup \{w_i\}$, for all $1 \leq i < n$

Given a subset D of a surface Σ, we write $\mathring{D}$, ∂D, and $\overline{D}$ for the topological interior, boundary, and closure, of D in Σ, respectively. For positive integers $\alpha_0, \alpha_1, \alpha_2$ and $\alpha := (\alpha_0, \alpha_1, \alpha_2)$, a graph G is α*-nearly embeddable* in Σ if there is a subset $A \subseteq V(G)$ with $|A| \leq \alpha_0$ such that there are integers $\alpha' \leq \alpha_1$ and $n \geq \alpha'$ for which $G - A$ can be written as the union of $n + 1$ edge-disjoint graphs $G_0, \dots, G_n$ with the following properties:

(i) For all $1 \leq i \leq j \leq n$ and $\Omega_i := V(G_i \cap G_0)$, the pairs $(G_i, \Omega_i) =: V_i$ are vortices, and $G_i \cap G_j \subseteq G_0$ when $i \neq j$.

(ii) The vortices $V_1, \dots, V_{\alpha'}$ are disjoint and have adhesion at most α_2; we denote the set of these vortices by $\mathcal{V}$. We will sometimes refer to these vortices as *large* vortices.

(iii) The vortices $V_{\alpha'+1}, \dots, V_n$ have length at most 3; we denote the set of these vortices by ν. These are the *small* vortices of the near-embedding.

(iv) There are closed discs in Σ, with disjoint interiors $D_1, \dots, D_n$, and an embedding $\sigma : G_0 \hookrightarrow \Sigma - \bigcup_{i=1}^{n} \mathring{D}_i$ such that $\sigma(G_0) \cap \partial D_i = \sigma(\Omega_i)$ for all i and the linear ordering of Ω_i is compatible with the natural cyclic ordering of its image (i.e., coincides with the linear ordering of $\sigma(\Omega_i)$ induced by $[0,1)$ when ∂D_i is viewed as a suitable homeomorphic copy of $[0,1]/\{0,1\}$). For $i = 1, \dots, n$ we think of the disc D_i as *accommodating* the (unembedded) vortex V_i, and denote D_i by $D(V_i)$.

We call $(\sigma, G_0, A, \mathcal{V}, \nu)$ an *α-near embedding* of G in Σ, or just a *near-embedding*, with *apex set* A. For any integer α' we shorten $(\alpha', \alpha', \alpha')$-near embedding to α'-near embedding. Note that an $(\alpha_0, \alpha_1, \alpha_2)$-near embedding also is an α'-near embedding for any $\alpha' \geq \max\{\alpha_0, \alpha_1, \alpha_2\}$.

Given a near-embedding $(\sigma, G_0, A, \mathcal{V}, \nu)$ of G, let G_0' be the graph resulting from G_0 by joining any two nonadjacent vertices $u, v \in V(G_0)$ that lie in a common small vortex $V \in \nu$; the new edge uv of G_0' will be called a *virtual edge*. By embedding these virtual edges disjointly in the disc $D(V)$ accommodating their vortex V, we extend the embedding $\sigma\colon G_0 \hookrightarrow \Sigma$ to an embedding $\sigma'\colon G_0' \hookrightarrow \Sigma$. We shall not normally distinguish G_0' from its image in Σ under σ'. For a subset $D \subseteq \Sigma$ the graph $G_0'[D]$ is the subgraph of G_0' that is induced by the vertices of G_0' in D.

A cycle C in Σ is *flat* if it bounds an open disc $D(C)$ in Σ. Disjoint cycles $C_1, \dots, C_n$ in Σ are *concentric* if they bound open discs $D(C_1) \supseteq \dots \supseteq D(C_n)$. A set $\mathcal{P}$ of path intersects $C_1, \dots, C_n$ *orthogonally*, and is *orthogonal to* $C_1, \dots, C_n$, if every path P in $\mathcal{P}$ intersects each of the cycles in a (possibly trivial but non-empty) subpath of P. Let G be a graph embedded in a surface Σ, and Ω a subset of its vertices. Let $C_1, \dots, C_n$ be concentric cycles of G. The cycles $C_1, \dots, C_n$ *enclose* Ω if $\Omega \subseteq \overline{D(C_n)}$. They *tightly enclose* Ω if, in addition, the following holds:

> *For all $1 \leq k \leq n$ every family $C'_{n-k+1}, \dots, C'_n$ of concentric cycles that lie in the subgraph of G embedded in $D(C_{n-k+1})$ and enclose Ω satisfies $C_{n-k+1} = C'_{n-k+1}$.*

For a near-embedding $(\sigma, G_0, A, \mathcal{V}, \nu)$ of a graph G in a surface Σ and concentric cycles $C_1, \dots, C_n$ in G_0', a vortex $V \in \mathcal{V}$ is *(tightly) enclosed* by these cycles if they (tightly) enclose $\Omega(V)$. When we speak of the *genus* of a surface Σ we always mean its Euler genus, the number $2 - \chi(\Sigma)$. An $(\alpha_0, \alpha_1, \alpha_2)$-near embedding $(\sigma, G_0, A, \mathcal{V}, \nu)$ of a graph G in some surface Σ is called *(β, r)-rich* for integers $3 \leq \beta \leq r$ if the following statements hold:

(i) G_0' contains r-wall H and H is flat w.r.t the embedding σ.

(ii) If $\Sigma \not\cong S^2$, the representativity of G_0' in Σ is at least β, that is, every genus-reducing curve in Σ has an intersection with G_0' of cardinality at least β.

(iii) For every vortex $V \in \mathcal{V}$ there are β concentric cycles $C_1(V), \dots, C_\beta(V)$ in G_0' tightly enclosing V. For $D_i(V) := D(C_i(V))$ with $i = 1, \dots, \beta$ we have $\Omega(V) \subseteq D_\beta(V)$ and $\overline{D(H)} \cap \overline{D_1(V)} = \emptyset$. Moreover, for distinct large vortices $V, W \in \mathcal{V}$, the discs $\overline{D_1(V)}$ and $\overline{D_1(W)}$ are disjoint.

(iv) Let $V \in \mathcal{V}$, say $\Omega(V) = (w_1, \dots, w_n)$. Then there is a linked decomposition of V with linkage $\mathcal{L}$ of adhesion at most α_2 and a path P in $V \cup (\bigcup \nu)$ with $V(P \cap G_0) = \Omega(V)$ that avoids all the paths of $\mathcal{L}$, and traverses $w_1, \dots, w_n$ in this order.

(v) For every vortex $V \in \mathcal{V}$, its set of society vertices $\Omega(V)$ is linked in G_0' to branch vertices of H by a set $\mathcal{P}(V)$ of β disjoint paths having no inner vertices in H.

(vi) For every vortex $V \in \mathcal{V}$, the paths in $\mathcal{P}(V)$ intersect the cycles $C_1(V)$, ..., $C_\beta(V)$ orthogonally.

Now we have the terminology to state the structure theorem from [4].

Theorem 3.1. *For every non-planar graph R and integers $3 \le \beta \le r$ there exist integers $\alpha_0 = \alpha_0(R, \beta)$, $\alpha_1 = \alpha_1(R)$ and $w = w(\alpha_0, R, \beta, r)$ such that the following holds with $\alpha = (\alpha_0, \alpha_1, \alpha_1)$. Every graph G of tree-width $\mathrm{tw}(G) \ge w$ that does not contain R as a minor has an α-near, (β, r)-rich embedding in some surface Σ in which R cannot be embedded.*

Whenever we talk about $\alpha_0(R, \beta)$, $\alpha_1(R)$, or $w(\alpha_0, R, \beta, r)$ we refer to the values provided in this theorem. Let us finish this section with some terminology useful when talking about near embeddings. For a vertex set $X \subseteq V(G)$ the sum

$$\mathcal{E}_G(X) := \sum_{x \in X} (d_G(x) - 6)$$

is the *excess* of X in G. The *excess* of a subgraph of G is the excess of its vertex set. We may omit the index of $\mathcal{E}_G$ if it is clear from the context or all options in the current context have the same excess.

Let $G = (V, E)$ be a graph that is embedded in some surface Σ. Any graph on n vertices that is maximal with the property of being embeddable in Σ has precisely $3n - 3\chi(\Sigma)$ edges. Hence it holds that $\mathcal{E}(V) \le -6\chi(\Sigma)$. Clearly, the sum of the excesses of two disjoint sets is the excess of their union. Thus, if for some $X \subseteq V$ and $c \in \mathbb{N}$ we have $\mathcal{E}(X) \geq c - 6\chi(\Sigma)$ we also have $\mathcal{E}(V \setminus X) \leq -c$. The *excess* of a face F with boundary cycle C is defined to be the number $\mathcal{E}(V(C)) + 2|C| - 6$. Note that one can think of this as the maximal potential excess as this is just the excess of $V(C)$ in the graph obtained from G by embedding disjointly as many edges as possible in the face F. The *excess* $\mathcal{E}(V)$ of a vortex $V \in \mathcal{V}$ in a near embedding $(\sigma, G_0', A, \mathcal{V}, \nu)$ is the excess of the face of G_0' that contains the accommodating disc of that vortex.

3.2 Forcing edges into a surface

The terminology regarding linear decompositions used in this section is thoroughly defined in Section **I**–3. We now use the results from Section 2 to show that for an α-near (β, r)-rich embedding $(\sigma, G_0, A, \mathcal{V}, \nu)$ of a $(2k+3)$-connected graph G that is not k-linked one of the large vortices has arbitrarily large negative excess if r is large enough.

This vortex together with its β tightly enclosing cycles will have a rich enough structure to find a long regular stable decomposition whose auxiliary graph is 'large'[5] enough to solve our linkage problem.

A subgraph $H \subseteq G$ together with a linear decomposition $(\mathcal{W}, \mathcal{P})$ of H is a *decomposed subgraph* of G. We also call H a *decomposed subgraph* and $(\mathcal{W}, \mathcal{P})$ *its decomposition.* A decomposed subgraph $H \subseteq G$ is *a stable regular subgraph* if its decomposition $(\mathcal{W}, \mathcal{P})$ is stable and regular. Its *length*, *attachedness*, and *adhesion* are the length, attachedness and adhesion of $(\mathcal{W}, \mathcal{P})$, respectively. The stable regular subgraph H with decomposition $(\mathcal{W}, \mathcal{P})$ has a *collar* $H^{\mathcal{P}}_{\Pi}$ for some path $\Pi \subseteq \Gamma(\mathcal{W}, \mathcal{P})$ with *leaf* j_1 and *collarsize* $n \in \mathbb{N}$ if

1. The path $\Pi = j_1 j_2 \dots j_n$ in $\Gamma(\mathcal{W}, \mathcal{P})$ has length $n-1$, and j_1 has degree 1 in $\Gamma(\mathcal{W}, \mathcal{P})$, and j_i has degree 2 in $\Gamma(\mathcal{W}, \mathcal{P})$ for $1 < i \leq n$. The j_i are the indices of paths of $\mathcal{P} \subseteq G$. To reduce the clutter in our notation we shall always assume that $j_i = i$ for $1 \leq i \leq n$.

2. Every path from $G - H$ to $H^{\mathcal{P}}_{\Gamma(\mathcal{W},\mathcal{P})-\Pi}$ meets all paths P_i with $1 \leq i \leq n$ or the first or last adhesion set of $(\mathcal{W}, \mathcal{P})$.

3. By our assumption in Property 1 the path $\Pi_i := 1\Pi(i-1)$ is the initial subpath of length $i-2$ of Π. For every foundational linkage of $\mathcal{W}$ if the path Q_i starting in the endvertex of P_i is contained in $H^{\mathcal{P}}_{\Gamma(\mathcal{W},\mathcal{P})-\Pi_i}$, then $Q_i = P_i$.

[5] or 'rich'

4. $H_{\Pi}^{\mathcal{P}}$ restricted to the inner bags of $(\mathcal{W}, \mathcal{P})$ can be embedded into a disc with the first and last adhesion set on the boundary and such that the image of P_i separates the disc into two disc, one of which contains the images of the paths with smaller indices and one of which contains the paths with larger indices.

5. Every vertex in the collar $H_{\Pi}^{\mathcal{P}}$ is adjacent to at most $2k-3$ trivial paths of $\mathcal{P}$.

6. For every inner bag W of $(\mathcal{W}, \mathcal{P})$, the path $P_1[W]$ is incident with $2|P_1[W]| + 6n$ edges in $G[W] - E(P_1)$.

Let us give some intuition for these properties. The collar of size n corresponds to the graph enclosed by n concentric cycles around a vortex. Property 2 captures the separating property of a cycle that encloses a vortex. Namely, the vortex is separated from the wall by the union of the cycle and the apex set (and the small vortices but we will get rid of those by some other argument). Note that every vertex in an inner bag of $\mathcal{W}$ that also has some neighbour that does not lie in any inner bag of $\mathcal{W}$ lies either in the first or last adhesion set of $(\mathcal{W}, \mathcal{P})$ or in P_1. This defines a demarcation line and on the 'inside' we will have enough strutcure to apply the results from Part **I** while the 'outside' is packed away neatly enough to not cause any complications for us. Property 3 captures the 'tightly' from the β tightly enclosing cycles around a vortex. Property 4 is an obvious consequence of the fact that the tightly enclosing cycles are embedded in a disc. Property 5 is used in combination with the high connectedness and thus high minimum degree of G to ensure that all vertices in the collar not in the first or last adhesion set of $\mathcal{W}$ nor on P_1 have degree at least 6 in the collar itself. Property 6 ensures that there are enough edges in H as to ensure a structure in $\Gamma(\mathcal{W}, \mathcal{P})$ that lets us solve any linkage problem.

With these definitions we are able to state and prove the main link between the bounded tree-width case and the unbounded tree-width case:

Lemma 3.2. *For all integers N and k there exists an integer w such that the following holds. Every $(2k+3)$-connected graph G of tree-width $\mathrm{tw}(G) \geq w$ is either k-linked or contains a stable regular subgraph of length N, attachedness $2k+3$, with collarsize $6k$, and adhesion at most $\alpha_0(K_{3k+6}, 6k) + 2\alpha_1(K_{3k+6}) + 12k$.*

Proof. Let us first introduce the values of some integers $N_1, \ldots, N_6$ that are dependent only on the integers N and k. In the following we need numbers

$\alpha_0 := \alpha_0(K_{3k+6}, 6k)$ and $\alpha_1 := \alpha_1(K_{3k+6})$ as given by Theorem 3.1. We set $\alpha^+ := \alpha_0 + 2\alpha_1 + 12k$ as this term is used several times later on, it is the upper bound for the adhesion we are aiming for.

Let

$$\chi(k) := \min\{\chi(\Sigma) \mid K_{3k+6} \text{ can not be embedded into } \Sigma\}$$

and let $\rho := 6k\binom{\alpha_0}{2k} + 2k\binom{\alpha_0}{2k-2}$. Note that these values only depend on the choice of k. Let

N_6 be the integer returned by **I**–3.5 for
$a := 2k,$
$l := N,$
$p := 2k + 3,$
$w := \alpha^+,$
$N_5 - 1$ be the integer returned by [6, Lemma 3.5] for
$l := N_6,$
$p := \alpha^+,$
$q := 2k + 3,$

$$N_4 := N_5\left(2\rho + \binom{\alpha^+}{2k-2}\right),$$

$$N_3 := N_4^{\alpha^+ + 1}(\alpha^+)!$$

$$N_2 := N_3(72k + 6\rho + 6) + 6,$$

$$N_1 := N_2\alpha_1 + 2 + (6k + 31)\binom{\alpha_0}{2k-3} + 25\rho - 6\chi(k)$$

We will show that the conclusion of the lemma follows if we set the value of w to $w(\alpha_0, K_{3k+6}, 6k, N_1)$ as provided by Theorem 3.1. Let G be a $(2k+3)$-connected graph with tree-width at least w. By Theorem 2.7 we may assume that K_{3k+6} is not a minor of G. Set $\alpha := (\alpha_0, \alpha_1, \alpha_1)$. By the structure theorem (3.1) there is an α-near $(6k, N_1)$-rich embedding $(\sigma, G_0, A, \mathcal{V}, \nu)$ of G into some surface Σ in which K_{3k+6} can not be embedded.

Claim 3.2.1. *If $|\nu| \geq 2k\binom{\alpha_0}{2k}$ then G is k-linked.*

Proof. Each small vortex $X \in \nu$ contains a vertex v_X not in its society. Since G is $(2k+3)$-connected Menger's Theorem yields a fan with centre v_X

of size $2k+3$ whose leaves are contained in the flat wall. As $|\Omega(X)| \leq 3$ there are at least $2k$ internally disjoint v_X–A paths in that fan, and their union is a fan S_X that has all its non leaves in X and all its at least $2k$ leaves in A. For distinct small vortices X, Y we have that $V(S_X \cap S_Y) \subseteq A$. If there are $2k\binom{\alpha_0}{2k}$ small vortices, then we have at least that many such fans and, by the pigeon hole principle, there is one $2k$-subset of A that is the set of leaves for $2k$ such fans. These two vertex sets of size $2k$ are the branch vertices of a $TK_{2k,2k}$ and with Lemma 2.1 this implies that G is k-linked. □

Claim 3.2.2. *If the number of vertices in G_0' with $2k-2$ neighbours in A is more than $2k\binom{\alpha_0}{2k-2}$ then G is k-linked.*

Proof. Similar to the proof of Claim 3.2.1 that many vertices with $2k-2$ neighbours in A would force one of the $2k-2$ subsets of A to be completely adjacent to a set of $2k$ vertices. In this case we obtain a $K_{2k,2k-2}$ as a subgraph of G and this is sufficient for the k-linkedness of G by Lemma 2.2. □

Claim 3.2.3. *If the number of vertices in some pseudo diagonal of a flat r-wall in G_0 that have $2k-3$ neighbours in A is at least $(6k+31)\binom{\alpha_0}{2k-3}$, then G is k-linked.*

Proof. If there are more such vertices, then there is a pseudo diagonal with cardinality $6k+31$ that is completely adjacent to a $(2k-3)$-subset of A. With Lemma 2.6 the claim follows. □

Claim 3.2.4. *There is a large vortex $V \in \mathcal{V}$ tightly enclosed by $C_1(V), \ldots, C_{6k}(V)$ such that the face with boundary $C_1(V)$ in $G_0'[\Sigma \setminus \mathring{D}_1(V)]$ has excess at most $-N_2$.*

Proof. Let H be a flat N_1-wall in G_0' as given by the $(6k, N_1)$-richness of the near embedding. By replacing H with a subwall of size at least $N_1 - 18\rho$ if necessary, we may ensure that H does not contain a vertex that lies in the society of a small vortex and thus $H \subseteq G_0$. For every vortex $V \in \mathcal{V}$ there are $6k$ concentric cycles $C_1(V), \ldots, C_{6k}(V)$ in G_0' tightly enclosing V and bounding open discs $D_1(V) \supseteq \ldots \supseteq D_{6k}(V)$, such that $D_{6k}(V)$ contains $\Omega(V)$ and $\overline{D(H)}$ does not meet $\overline{D_1(V)}$. For distinct large vortices $V, W \in \mathcal{V}$, the discs $\overline{D_1(V)}$ and $\overline{D_1(W)}$ are disjoint.

Let $G_i(V) := G_0'[\overline{D_i(V)}]$ be the restriction of G_0' to the vertices that are mapped into $\overline{D_i(V)}$ by σ and let

$$X := V\left(G_0' \setminus \bigcup_{V \in \mathcal{V}} G_1(V)\right)$$

be the set of vertices of G_0' not in $G_1(V)$ for any $V \in \mathcal{V}$. Note that $V(H) \subseteq X$, that is the vertice of our flat wall are contained in X. Let $G'' := G_0'[\Sigma \setminus \bigcup_{V \in \mathcal{V}} \mathring{D_1}(V)]$ be the subgraph of G_0' that misses the vertices in the interior of the discs $D_1(V)$ for $V \in \mathcal{V}$. We now have $G'' \cap G_1(V) = C_1(V)$ for any $V \in \mathcal{V}$. Let F_V be the face of G'' that contains the accommodating disc of $V \in \mathcal{V}$. We have that $F_V = D_1(V)$.

Let G''' be the graph obtained from G'' by additionally embedding disjointly as many edges as possible, that is $|C_1(V)| - 3$, in the face F_V for each $V \in \mathcal{V}$. As G''' is still embeddable into Σ we have

$$\begin{aligned}
-6\chi(k) \geq -6\chi(\Sigma) &\geq \mathcal{E}_{G'''}(V(G''')) \\
&= \mathcal{E}_{G''}(V(G'')) + \sum_{V \in \mathcal{V}} (2|C_1(V)| - 6) \\
&\geq \mathcal{E}_{G''}(X) + \sum_{V \in \mathcal{V}} \mathcal{E}_{G''}(C_1(V)) + \sum_{V \in \mathcal{V}} (2|C_1(V)| - 6) \\
&\geq \mathcal{E}_{G''}(X) + \sum_{V \in \mathcal{V}} \mathcal{E}_{G''}(F_V).
\end{aligned}$$

By Claim 3.2.1, Claim 3.2.2 all but at most ρ vertices in X have degree at least 6 in G_0' while by Claim 3.2.3 from the at least $N_1 - 18\rho - 2$ vertices on the diagonal of H, which is also a pseudo diagonal of H at least $N_1 - 3\rho - 2 - (6k+31)\binom{\alpha_0}{2k-3} - \rho$ are neither incident with at least $2k-3$ vertices in A nor with a virtual edge and thus have degree 7 in G_0' and also in G_0. With $M := \min_{V \in \mathcal{V}} (\mathcal{E}(F_V))$ we have

$$-6\chi(k) \geq \mathcal{E}_{G''}(X) + \alpha_1 M \geq N_1 - 18\rho - 2 - (6k+31)\binom{\alpha_0}{2k-3} - 7\rho + \alpha_1 M$$

which implies

$$N_1 \leq 2 + (6k+31)\binom{\alpha_0}{2k-3} + 25\rho - \alpha_1 M - 6\chi(k).$$

By the choice of N_1 we have one vortex $V \in \mathcal{V}$ such that F_V has excess at most $-N_2$ in G''. □

Claim 3.2.5. *There is a linked vortex (V', Ω') with $V' = V \cup G_1(V) \cup A$*[6] *and decomposition $\mathcal{W}$ with foundational linkage $\mathcal{P}$ of adhesion at most $\alpha^+ = \alpha_0 + 2\alpha_1 + 12k$.*

Furthermore, the vertices of Ω' are precisely $V(C_1(V) - v)$ for some vertex $v \in C_1(V)$ and the excess of (V', Ω') is at most $-N_2 + 6$. $\mathcal{P}$ necessarily contains precisely $6k$ non-trivial paths that are disjoint from V. Let $P_1, \dots, P_{6k}$ be those paths and let them be ordered naturally. $P_1 = C_1(V) - v$ in this case.

Proof. Using the ideas from [4, Lemma 15] we can find for any vertex v of $C_1(V)$ a set S_v of $6k$ vertices that contains one vertex from each cycle C_i such that there is a curve γ_v in Σ from v to $\Omega(V)$ meeting G_0' only in S_v and in a vertex $\omega_v \in \Omega(V)$. We may assume that for any two vertices v, w their curves γ_v and γ_w either are disjoint or $\gamma_v \cap \gamma_w = \gamma_v[x, 1] = \gamma_w[y, 1]$ for some values $x, y \in [0, 1]$.[7] There is an edge $v_0 v_\infty$ in $C_1(V)$ such that ω_{v_0} is Ω-minimal and ω_{v_∞} is Ω-maximal in $\{\omega_v \mid v \in C_1(V)\}$. For any $v \in C_1(V)$ let A_v be the adhesion set of V that contains $\omega_v \in \Omega(V)$. Then $S_v \cup A_v \cup S_{v_0} \cup A_{v_0} \cup A$ for $v \in C_1(V) - v_0$ are the adhesion sets of a slim decomposition $\mathcal{W}$ of V' and as $C_1(V), \dots, C_{6k}(V)$ are tightly enclosing every foundational linkage of $\mathcal{W}$ contains $P_1 := C_1(V) - v_0$. The natural order of the $6k$ paths disjoint from V is provided by the order they intersect with the curves γ_v for any $v \in C_1(V) - v_0$. The adhesion of $\mathcal{W}$ is at most α^+. □

Claim 3.2.6. *There is a contraction $\mathcal{W}'$ of $\mathcal{W}$ of length N_3 with trivial paths θ such that every vertex v in an inner bag of $\mathcal{W}'$ that is incident with a virtual edge or with more than $2k - 3$ vertices in the set θ is necessarily contained in θ and $A \subseteq \theta$.*

Proof. We give an algorithm to construct edge disjoint paths $Q_1, \dots, Q_{N_3}$ with the following properties (i) to (iii).

(i) $\bigcup_{1 \le i \le N_3} Q_i$ covers precisely the edges of $C_1(V) - v_0$.

(ii) For all $q_i \in Q_i$ and $q_j \in Q_j$ with $i < j$ we have $q_i \le_{\Omega'} q_j$.

Let $v \le_{\Omega'} w$ be the end vertices of Q_i for some $1 \le i \le N_3$. Let W_{j_v} be the bag of $\mathcal{W}$ whose left adhesion set contains S_v and let W_{j_w} be the

[6] Note that this is neither a subgraph of G nor of G_0'. We will transform it into a subgraph of G later on by restricting it to a section that contains no virtual edges.

[7] To see this just pick one curve after another and follow any previous curve from the first point they intersect.

bag of $\mathcal{W}$ whose right adhesion set contains S_w and let $W'_i := W_{[j_v,j_w]}$ and $X_i := G_1(V)[W'_i]$.

(iii) $\sum_{v\in Q_i} d_{X_i}(v) \geq 4|Q_i| + 72k + 6\rho$.

As $N_2 - 6 = N_3(72k + 12\rho + 6)$ the following algorithm provides such paths. Start at the Ω'-smallest vertex in $C_1(V) - v_0$ and add it to Q_1. Extend Q_1 along $C_1(V)$ with Ω'-increasing vertices until the excess of Q_1 in G'' is at most $-72k - 12\rho - 12$. Start the process with the last vertex added to the previous path for the next path. It is clear that we end up with enough paths with this construction as every vertex contributes at most -6 to the excess and thus we 'lose' at most $6N_3$ excess to work with at the vertices that lie in two paths. And (iii) holds as Q_i has at most ρ vertices with negative excess in G'_0 and the constant -12 allows for enough slack to ignore the end vertices of Q_i in the computation.

Thus $\mathcal{E}_{X_i\cap G}(W'_i) \geq -2|Q_i| - 6|W'_i \cap \Omega(V)|$ as we 'lose' at most 6 from the sum for any vertex in $Q_i, S_v, S_w, W'_i \cap \Omega(V)$, and also at most 6ρ for the virtual edges or vertices incident with at least $2k-2$ vertices in the apex set.

The restriction of σ to $X_i \cap G$ embeds $X_i \cap G$ into a closed disc in such a way that the boundary of the disc has precisely the vertices $Q_i \cup S_v \cup S_w \cup (\Omega(V) \cap W'_i)$ mapped to it.

Suppose for a contradiction that $S_v \cap S_w \neq \emptyset$, then the boundary contains precisely the images of the vertices in $Q_i \cup S_v \cup S_w$ and $\mathcal{E}_{X_i\cap G}(W'_i) \geq -2|Q_i|$ as we don't have extra vertices from $W'_i \cap \Omega(V)$ that contribute to a lower excess. As the excess of a graph embedded into the plane is at most -12 and we could add $|Q_i \cup S_v \cup S_w| - 3$ edges to the drawing of X_i provided by σ it holds that

$$\begin{aligned} -12 &\geq \mathcal{E}_{X_i\cap G}(W'_i) + 2|Q_i \cup S_v \cup S_w| - 6 \\ &\geq -2|Q_i| + 2|Q_i \cup S_v \cup S_w| - 6 \\ &\geq -6. \end{aligned}$$

With this contradiction we showed that S_v and S_w are disjoint. The slim decomposition $(V(W'_1), \ldots, V(W'_{N_3}))$ has length $N_3 = N_4^{\alpha^+ + 1}\alpha^+!$. By [6, Lemma 3.4, Lemma 3.5] we have a contraction $\mathcal{W}^{\text{temp}}$ of it of length N_4 that additionally satisfies **I**–(L7), that is any path of a foundational linkage is either trivial or contains an edge in every inner bag of $\mathcal{W}^{\text{temp}}$.

As $N_4 = N_5\left(2\rho + \binom{\alpha^+}{2k-2}\right)$ there are indices i, j with $j - i = N_5$ such that $W_{[i,j]}$ contains no virtual edges or vertices incident with more than

$2k-3$ trivial paths (Note that the apex set is contained in the trivial path by construction). Let $\mathcal{W}'$ be the contraction of $\mathcal{W}^{\text{temp}}$ along $1, i, i+1, \ldots, j$. $\mathcal{W}'$ is a slim decomposition of length N_5 with adhesion at most α^+, satisfying **I**–(L7). Every foundational linkage of $\mathcal{W}'$ contains precisely $6k$ non-trivial paths that are disjoint from V. □

As in Part **I**, for a linkage $\mathcal{Q}$ with trivial paths θ in a graph H, the union of θ with a proper rerouting of $\mathcal{Q} \setminus \theta$ obtained from applying Lemma **I**–3.3 to $\mathcal{Q} \setminus \theta$ in $H - \theta$ is a *bridge stabilisation of* $\mathcal{Q}$ *in* H.

Claim 3.2.7. *Every bridge stabilisation* Q' *of a foundational linkage of* $\mathcal{W}'$ *is* $(2k+3)$*-attached.*

Proof. We follow the proof of Lemma **I**–3.7 (iii). For any inner bag W of $\mathcal{W}'$ and the set Z consisting of the left and right adhesion set of W and the vertices of $P_1[W]$ it holds that there is an x–Z fan of size $2k+3$ in $G[W]$ for any $x \in W \setminus Z$ as G is $(2k+3)$-connected and Z separates x from the flat wall.

We stay in the notation from **I**–3.7 (iii). Suppose B is a non-trivial hosted $\mathcal{Q}'$-bridge and let Q' be the non-trivial path to which it attaches. If Q' is the leaf P_1 of the collar, then B has at most one attachment on $\mathcal{Q}'$, as P_1 is a subgraph of the tightly enclosing cycle $C_1(V)$. Consequently, B is attached to at least $2k+2$ paths of θ as G is $(2k+3)$-connected. We can follow the remainder of the proof verbatim as $P_1 \subseteq Y \setminus X$. □

Let us recap this situation in the light of [3, Section 3]. We have a decomposed subgraph $H \subseteq V \cup G_1(V)$ for some large vortex V of G with slim decomposition $(\mathcal{W}', \mathcal{Q}')$ say that has length N_5, attachedness $2k+3$, adhesion at most α^+, and $6k$ non-trivial paths of $\mathcal{Q}'$ are disjoint from V.

The slim decomposition $\mathcal{W}'$ with attachedness $p = 2k+3$ and adhesion $q \leq \alpha^+$ has length $\lambda{+}1 = N_5$. According to [6, Lemma 3.5] we have a regular decomposition $(\mathcal{W}'', \mathcal{P}')$ of length N_6 which originates from a contraction of $\mathcal{W}'$. The proof of Theorem **I**–3.5 invokes p-connectedness two times. The first time to obtain a regular decomposition with the properties of $(\mathcal{W}'', \mathcal{P})$ and the second time to prove that every bridge stabilisation of $\mathcal{P}$ is $(2k+3)$-attached. As we already have the regular decomposition $(\mathcal{W}'', \mathcal{P})$ we can disregard the first instance. And with Claim 3.2.7 we can disregard the second instance, too.

Thus we have a contraction $(\mathcal{X}, \mathcal{Q})$ of $(\mathcal{W}'', \mathcal{P})$ that is a stable regular decomposition of adhesion at most α^+, attachedness $2k+3$, and length N. It remains to show that it has collarsize $6k$. Every foundational linkage of $\mathcal{X}$

contains precisely $6k$ non-trivial paths that are disjoint from V. We may assume that $\mathcal{Q}$ contains the path $P_i \subseteq C_i(V)$ for $1 \leq i \leq 6k$. Then Property 1 and Property 2 hold as $C_1(V), \ldots, C_{6k}(V)$ enclose V and Property 3 holds as they do so tightly. As $H^{\mathcal{Q}}_{\Pi}[W_{1,N-1}]$ is embedded into the disc $D_1(V)$ by σ it is planar and Property 4 holds. Property 5 holds by Claim 3.2.6. Finally, Property 6 holds as the decomposition constructed in the proof of Claim 3.2.6 with bags $V(W_i')$ already had enough edges and moving to a contraction only increases the count. □

Foregoing some details, we say that S is 'properly attached' to a decomposed subgraph H if there is a linkage from S to H such that each path of the linkage intersects with H only in the first and last bag of its decomposition. In the following we construct such a properly attaching linkage without destroying to much of the length of H nor of the size of its collar.

In detail the definition looks as follows. Let H be a stable regular subgraph of some graph G with decomposition $(\mathcal{W}, \mathcal{P})$ and let $\lambda \subseteq \mathcal{P}$ be the non-trivial paths in $\mathcal{P}$. Let C be a component of $\Gamma(\mathcal{W}, \mathcal{P})[\lambda]$, then $N(C)$ are the (trivial) neighbours of C in $\Gamma(\mathcal{W}, \mathcal{P})$. A set S that intersects with $V(H)$ only in the first and last bag of $\mathcal{W}$ can be *attached properly to* C if there is a linkage L from S to $\Gamma^{\mathcal{P}}_{C \cup N(C)}$ such that all paths of L end in the first or last adhesion set of $\mathcal{W}$. In this case L *attaches* S *properly to* C.

The main idea to construct such a linkage that properly attaches a set S of size $2k$ to our decomposed subgraph H with collarsize $6k$ is the following. We take $2k$ disjoint S–H paths and reroute them along the $2k$ outer most paths of the collar to the first and last adhesion set of the decomposition of H. In the construction we will lose some length and collarsize $2k$. First we provide a single purpose lemma to reroute linkages along other linkages.

Lemma 3.3. *Let G be a graph, let $S, T \subseteq V(G)$ be of cardinality $k \in \mathbb{N}$, and let $\mathcal{H}$ be a set of pairwise disjoint connected subgraphs of G such that each $H \in \mathcal{H}$ contains a vertex from T. Let $\mathcal{P}$ be a set of k disjoint paths such that each path of $\mathcal{P}$ has one end in S and either has its other end in T or meets at least k distinct graphs from $\mathcal{H}$. Then there are k disjoint S–T paths in $\bigcup(\mathcal{P} \cup \mathcal{H})$.*

Proof. The proof will be one straight forward application of Menger's Theorem. Suppose there are not k disjoint S–T paths in $\bigcup(\mathcal{P} \cup \mathcal{H})$. By Menger's Theorem there is a set X of at most $k-1$ vertices that separates S and T in $\bigcup(\mathcal{P} \cup \mathcal{H})$. By cardinality one path $P \in \mathcal{P}$ contains no vertex from X. If P ends in T, then this contradicts the fact that X separates S and T. If P meets k graphs from $\mathcal{H}$, then there is one, $H \in \mathcal{H}$ say, that is disjoint from

X. Thus $P \cup H$ contain an S–T path avoiding X, again contradicting that X is a separator. $\square$

Lemma 3.4. *For integers k, $n > 2k$, N, and α let G be a $(2k+3)$-connected graph and let $S \subseteq V(G)$ be a vertex set of cardinality $2k$. Let H be a stable regular subgraph of G with decomposition $(\mathcal{W}, \mathcal{P})$ with collarsize n of length $(2k+1)N + 8k$, attachedness $2k+3$, and adhesion at most $\alpha + n$. Let S intersect with $V(H)$ only in the first and last bag of $\mathcal{W}$. Then there is a stable regular subgraph H' of G with decomposition $(\mathcal{W}', \mathcal{P}')$ of length N, with attachedness $2k+3$, collarsize $n-2k$, and adhesion at most $\alpha+n-2k$ such that S can be properly attached to the component of $\Gamma(\mathcal{W}', \mathcal{P}')[\lambda']$ that contains the leaf of the collar where λ' denotes the non-trivial paths in $\mathcal{P}'$. In particular, each path of $\mathcal{P}'$ is a subpath of a path of $\mathcal{P}$.*

Proof. Set $\Gamma := \Gamma(\mathcal{W}, \mathcal{P})$ and let $\Pi = 1, \ldots, n$ be the path in Γ such that $H^{\mathcal{P}}_{\Pi}$ is the collar. It is clear that every contraction of length m of $(\mathcal{W}, \mathcal{P})$ is a decomposition witnessing that H is a stable regular subgraph with collarsize n of length m and attachedness $2k+3$.

Claim 3.4.1. *For all $0 \le i < n$ the graph $H_i := H^{\mathcal{P}}_{\Gamma-\{1,\ldots,i\}}$[8] with decomposition $(\mathcal{W}', \mathcal{P} \setminus \{P_1, \ldots, P_i\})$ where the jth bag of $\mathcal{W}'$ is $W \cap V(H_i)$ and W is the jth bag of $\mathcal{W}$ is a stable regular subgraph of G with the same length as H and at least its attachedness and collarsize $n-i$.*

Let $P'_i := P_i[W']$ and $P'_{i+1} := P_{i+1}[W']$ and let X be the set consisting of the four end vertices of P'_i and P'_{i+1}. For $i = 0$ we have $H_0 = H$ and thus the statement holds by assumption. As Property 1 to Property 5 are trivially true for any H_i with decomposition $(\mathcal{W}', \mathcal{P} \setminus \{P_1, \ldots, P_i\})$ it remains to show that Property 6 holds.

Suppose the claim does not hold and let $1 \le i < n$ be the smallest i such that Property 6 is violated as witnessed by a bag W', say. That is, for some inner bag W' of the decomposition $(\mathcal{W}', \mathcal{P} \setminus \{P_1, \ldots, P_i\})$, the path P'_{i+1} is incident with less than $2|P'_{i+1}| + 6(n-i)$ edges in $G[W'] - E(P_{i+1})$. By Property 4 the graph $L := H^{\mathcal{P}}_{\{i,i+1\}}[W']$ is planar and the society $(L, P'_i(P'_{i+1})^{-1})$ is rural.

All vertices of $L - (P'_i \cup P'_{i+1})$ have degree at least 6 in L by Property 5 and the $(2k+3)$-connectedness of G. By Lemma **I**–6.2 we thus have

$$\sum_{v \in P'_i \cup P'_{i+1}} d_L(v) \le 4|P'_i| + 4|P'_{i+1}| - 6.$$

[8] Assuming that $\{1, \ldots, 0\}$ is the empty set we have $H_0 = H$.

By Property 6 for H_{i-1} we also get

$$4|P'_i| - 2 + 6(n - (i-1)) \leq \sum_{v \in P'_i} d_L(v).$$

As the paths P'_i and P'_{i+1} are disjoint the union commutes with the summation of the vertex degrees and thus we can reduce to

$$\sum_{v \in P'_{i+1}} d_L(v) \leq 4|P'_{i+1}| - 6n + 6i - 10.$$

As every vertex in $L - (X \cup P_1)$ has degree at least 6 in the restriction $H^{\mathcal{P}}_{\Pi}[W']$ of the collar to the bag W' and the vertices in X have at least degree one in L, we have for $d'(v) := d_{H^{\mathcal{P}}_{\Pi}[W']}(v)$

$$\begin{aligned} 6|P'_{i+1}| &\leq \sum_{v \in P'_{i+1}} d'(v) + 10 \\ &= \sum_{v \in P'_{i+1}} d_L(v) + 10 \qquad + \sum_{v \in P'_{i+1}} \big(d'(v) - d_L(v)\big) \\ &< \sum_{v \in P'_{i+1}} d_L(v) + 10 \qquad +2\left|P'_{i+1}\right| + 6(n-i) \end{aligned}$$

which reduces to

$$4\left|P'_{i+1}\right| - 6n + 6i - 10 < \sum_{v \in P'_{i+1}} d_L(v),$$

contradicting the upper bound. This means that Property 6 holds for H_i with collar $H^{\mathcal{P}}_{\Pi - \{1,\ldots,i\}}$ of size $n - i$ and thus concludes the proof of Claim 3.4.1.

By the pigeonhole principle there is an interval $W_i, \ldots, W_{i+N}$ of bags of $(\mathcal{W}, \mathcal{P})$ with $4k \leq i \leq 2kN + 4k$ such that the contraction $(\mathcal{W}', \mathcal{P}')$ of $(\mathcal{W}, \mathcal{P})$ along $i+1, \ldots, i+N$ contains a vertex from S in an inner bag only if it is a trivial path. Let C be the component of $\Gamma(\mathcal{W}', \mathcal{P}')$ that contains the leaf of the collar. Let X be the intersection of the first adhesion set of $(\mathcal{W}', \mathcal{P}')$ with $\Gamma^{\mathcal{P}'}_{C \cup N(C)}$ and let Y be the intersection of the last adhesion set of $(\mathcal{W}', \mathcal{P}')$ with $\Gamma^{\mathcal{P}'}_{C \cup N(C)}$. Let $L \subseteq H$ be the union of the paths $P_1, \ldots, P_{2k}$ and the first and last bag of $(\mathcal{W}', \mathcal{P}')$.

Claim 3.4.2. *There is a set $\mathcal{H}$ of $2k$ disjoint X–Y paths in L, and each path in $\mathcal{H}$ contains a path P'_i with $1 \leq i \leq 2k$ as a subpaths.*

Proof. We prove this claim with a straight forward application of Menger's Theorem. Any set $U \subseteq V(L)$ of at most $2k-1$ vertices misses one path P of the $2k$ paths $P_1, \ldots, P_{2k}$ of $(\mathcal{W}, \mathcal{P})$ contained in L. Also such a set U misses a path P' of the $2k$ paths $P_{2k+1}, \ldots, P_{4k}$ of $(\mathcal{W}, \mathcal{P})$ which are not contained in L (but whose end vertices are). Furthermore, U misses bags W_f and W_l of $\mathcal{W}$ that are contained in the first and last bag of $\mathcal{W}'$, respectively, as only consecutive bags of $\mathcal{W}$ have non-empty intersection with L and both W'_0 and W'_N each contain at least $4k$ bags of $\mathcal{W}$.

In the union $P \cup P'[V(L)] \cup G[W_f] \cup G[W_l] \subseteq L$ there is an X–Y path that avoids U. By Menger's Theorem we have a set $\mathcal{H}$ of $2k$ disjoint X–Y paths in L. Clearly, each path in $\mathcal{H}$ contains a path $\mathcal{P}'_i$ with $1 \leq i \leq 2k$ as a subgraph. □

For every path $P \in \mathcal{H}$ the set $V(P) \cup X \cup Y$ separates $\Gamma^{\mathcal{P}'}_{C \cup N(C)}$ from $G-H$ by Property 2. Thus every path from $G-H$ to $\Gamma^{\mathcal{P}'}_{C \cup N(C)}$ meets either all paths in $\mathcal{H}$ or ends in $X \cup Y$. Let $\mathcal{L}$ be a $(S \cup T)$–$\Gamma^{\mathcal{P}'}_{C \cup N(C)}$ linkage of size $2k$ which exists as G is $(2k+3)$-connected. Every path in $\mathcal{L}$ meets either all $2k$ paths in $\mathcal{H}$ or ends in $X \cup Y$. By Lemma 3.3 there is an S–$(X \cup Y)$ linkage $\mathcal{L}'$ all whose paths lie in $\bigcup(\mathcal{L} \cup \mathcal{H})$. This completes the proof as $H[\bigcup \mathcal{W}']$ with decomposition $(\mathcal{W}', \mathcal{P}')$ is a stable regular subgraph of G with collarsize $n-2k$, of length N, attachedness $2k+3$ and $\mathcal{L}'$ attaches S properly to C. □

For the following proof of Theorem 1.2 we closely follow the proof of Theorem **I**–1.1. Before the arguments from that proof work in our setting we need to adapt some lemmas. The complete Section **I**–4 'Token Movements' will be required here.

In the following let G be a p-connected graph that contains a stable regular subgraph H with decomposition $(\mathcal{W}, \mathcal{P})$ of length N with attachedness p and collarsize n. Set $\theta := \{\alpha \mid P_\alpha \text{ is trivial}\}$ and $\lambda := \{\alpha \mid P_\alpha \text{ is non-trivial}\}$.

Lemma 3.5 (from **I**–6.4)**.** **I**–*6.4 holds also if G is a stable regular subgraph of a p-connected graph with decomposition $(\mathcal{W}, \mathcal{P})$, with attachedness p, collarsize at least* 4*, and length at least* 3*.*

Proof. If both α and β are in the collar of G, the result follows by Property 4. If one of α or β is not in the collar, then we follow the proof from **I**–6.4. □

Lemma 3.6 (**I**–6.5)**.** **I**–*6.5 holds also if G is a stable regular subgraph of a p-connected graph with decomposition $(\mathcal{W}, \mathcal{P})$ with attachedness p and collarsize at least* 4*.*

Proof. As the leaf 1 of the collar is not a cut-vertex in $\Gamma(\mathcal{W},\mathcal{P})$ and every relinkage contains $\mathcal{P}_1$ as one of its paths we may follow the proof from **I**–6.5 verbatim. $\square$

We extend the definition of rich from Section **I**–7 by one special case: the leaf of a collar is always rich. One reason to have Property 6 in the definition of the collar as is, was to make the following arguments about 'richness' work.

We now restate our main theorem before proceeding with its proof.

Theorem 1.2. For every $k \in \mathbb{N}$ there is $w \in \mathbb{N}$ such that every $(2k+3)$-connected graph with tree-width at least w is k-linked.

Proof. Let k be given, let $\alpha_0 := \alpha_0(K_{3k+6}, 6k)$, let $\alpha_1 := \alpha_1(K_{3k+6})$, and let f be the function from the statement of Lemma **I**–4.10 with $n := \alpha_0 + 2\alpha_1 + 10k$. Let

$$N := \max\{2k(k+3)+1, 12k+4, 2f(k)+1\}$$

Let w be the integer returned from Lemma 3.2 when asked for G to be k-linked or containing a stable regular subgraph of length $2kN+8k$, attachedness $2k+3$, collarsize $6k$, and adhesion at most $n+2k$.

Let G be a $(2k+3)$-connected graph with tree-width at least w. We want to show that G is k-linked. Let $S = (s_1, \ldots, s_k)$ and $T = (t_1, \ldots, t_k)$ be disjoint k-tuples of distinct vertices of G. Suppose there is no S–T linkage in G. By Lemma 3.2 we may assume that G contains a stable regular subgraph of length $2kN+8k$, attachedness $2k+3$, adhesion at most $n+2k$, and with collarsize $6k$. By Lemma 3.4 the graph G contains a stable regular subgraph H with decomposition $(\mathcal{W},\mathcal{P})$ of length N, attachedness $2k+3$, with collarsize $4k$ and adhesion at most n such that there are paths $Q_1, \ldots, Q_{2k}$ attaching $S \cup T$ properly to Γ_0 where Γ_0 is the component of $\Gamma(\mathcal{W},\mathcal{P})[\lambda]$ that contains the leaf 1 of the collar.

It remains to show that with this given decomposition we can solve any linkage problem. This was researched extensively in Part **I** and we will follow the arguments given there and outline the difference where necessary. In particular we follow the proof of Theorem **I**–1.1 in Section **I**–7. We show the following claims only for the component Γ_0 of $\Gamma(\mathcal{W},\mathcal{P})[\lambda]$ containing the leaf 1 of the collar which is sufficient as we linked S and T properly to Γ_0.

Claim 7.1.1. holds as the leaf of the collar is not trivial and thus Γ_0 is not empty.

Claim 7.1.2. holds: We can ignore the first paragraph of its proof that uses the connectedness of G since we properly attached $S \cup T$ with Lemma 3.4. The remainder of the proof given there can be applied verbatim to our situation.

Claim 7.1.3. is true as all the $4k$ paths from the collar are in Γ_0.

Claim 7.1.4. holds for $\Gamma \subseteq \Gamma_0$ if $1 \notin V(\Gamma)$ with the same proofs.

Claim 7.1.5. reads Γ_0 contains a rich block in our case. We need to tend to the leaf of the collar individually here but that is easy. Its neighbour $2 \in \Gamma_0$ is rich in $\Gamma(\mathcal{W}, \mathcal{P})[\{1, 2, 3\}]$ by Claim 7.1.4. It is not rich in $\Gamma[\{1, 2\}]$ and thus it is rich in $\Gamma[\{2, 3\}]$. From here we follow the proof of Claim **I**–7.1.4 verbatim picking a sink of the directed tree R that is pointed to by $(2, \Gamma(\mathcal{W}, \mathcal{P})[\{2, 3\}])$.

Claim 7.1.6. holds as $\Gamma(\mathcal{W}, \mathcal{P})[\{1, 2\}]$ is not rich and we can thus follow the proof from Part **I** verbatim.

Claim 7.1.7. holds as $\Gamma(\mathcal{W}, \mathcal{P})[\{1, 2\}]$ is not rich.

Γ_0 is a component of $\Gamma(\mathcal{W}, \mathcal{P})[\lambda]$. It contains a rich block D by Claim 7.1.5. By Claim 7.1.6 and Claim 7.1.7 we have a triangle in D and $|D| + |N(D)| \geq 2k + 3$. This contradicts Claim 7.1.2 (iv) and thus concludes the proof of Theorem 1.2. □

References

[1] J. Fröhlich, K. Kawarabayashi, Th. Müller, J. Pott and P. Wollan: Linkages in Large Graphs of Bounded Tree-Width, *preprint* (2014).

[2] T. Böhme, K. Kawarabayashi, J. Maharry and B. Mohar: Linear Connectivity Forces Large Complete Bipartite Minors, *J. Combin. Theory B* **99** (2009), 557–582.

[3] T. Böhme, J. Maharry and B. Mohar: $K_{a,k}$ Minors in Graphs of Bounded Tree-Width, *J. Combin. Theory B* **86** (2002), 133–147.

[4] R. Diestel, K. Kawarabayashi, T. Müller and P. Wollan: On the excluded minor structure theorem for graphs of large tree-width, *J. Combin. Theory B* **102** (2012), 1189–1210.

[5] H.A. Jung: Eine Verallgemeinerung des n-fachen Zusammenhangs für Graphen, *Mathematische Annalen* **182** (1970), 95–103.

[6] K. Kawarabayashi, S. Norine, R. Thomas and P. Wollan: K_6 minors in 6-connected graphs of bounded tree-width, *arXiv*:**1203.2171** (2012).

[7] K. Kawarabayashi, S. Norine, R. Thomas and P. Wollan: K_6 minors in large 6-connected graphs, *arXiv*:**1203.2192** (2012).

[8] B. Oporowski, J. Oxley and R. Thomas: Typical Subgraphs of 3- and 4-connected graphs, *J. Combin. Theory B* **57** (1993), 239–257.

[9] N. Robertson and P. D. Seymour: Graph Minors. IX. Disjoint Crossed Paths, *J. Combin. Theory B* **49** (1990), 40–77.

[10] N. Robertson and P. D. Seymour: Graph Minors. XIII. The Disjoint Paths Problem, *J. Combin. Theory B* **63** (1995), 65–110.

Summary

We look at two common phenomena in infinite graphs: limit structures and ubiquitous patterns. The ubiquitous patterns are either assumed or developed by Ramsey-type arguments. Surprisingly, both phenomena have analogues in finite but arbitrarily large graphs.

Let G be a locally finite graph. In its Freudenthal compactification $|G|$ arcs and circles arise naturally as limit structures. They describe structures which in finite graphs are similarly described by paths and cycles. For example, a *Hamilton circle* in $|G|$ is a topological circle that contains all the (infinitely many) vertices of G. For a thorough introduction see [2]. These topological limit structures are studied in the three articles 'Dual trees must share their ends' (p. 1), 'Orthogonality and minimality in the homology of locally finite graphs' (p. 19), and 'Extending cycles locally to Hamilton cycles' (p. 27).

In the article 'Transitivity conditions in infinite graphs' (p. 40) we assume a high degree of symmetry, and show that this leads to ubiquity: the graphs characterized are shown to be made up of few types of finite subgraphs that arise everywhere. As these are organized into a tree-structure we are able to give a concise structural characterization of the graph properties that were originally defined merely in terms of symmetry.

The article 'Edge-disjoint double rays in infinite graphs: a Halin type result' (p. 73) confirms a conjecture of Andreae [1] that a graph that contains k edge-disjoint double rays for every $k \in \mathbb{N}$ also contains infinitely many edge-disjoint double rays. We obtain this infinite set of edge-disjoint double rays as a limit structure, while Lemma 29 (p. 86) explicitly describes a Ramsey-type phenomenon required in the steps of its construction.

The articles 'Linkages in large graphs of bounded tree-width' and 'Linkages in large graphs' (p. 89 and p. 155) are concerned with finite but arbitrarily large graphs. Interestingly, limits and ubiquity have a finite analogue here. In both articles we examine areas of graphs that converge to trees as the graphs get larger. Ramsey-type arguments allow us to obtain a tree-structure all whose parts express the same ubiquitous patterns.

All articles come with their own detailed and more specific introduction.

Literatur

[1] Th. Andreae. Über maximale Systeme von kantendisjunkten unendlichen Wegen in Graphen. *Math. Nachr.*, 101:219–228, 1981.

[2] R. Diestel. Locally finite graphs with ends: a topological approach. *Discrete Math.*, 310–312: 2750–2765 (310); 1423–1447 (311); 21–29 (312), 2010–11. arXiv:0912.4213.

Zusammenfassung

Wir betrachten zwei Eigenschaften, die man in unendlichen Graphen beobachten kann: Limiten und Muster die allgegenwärtig sind. Die allgegenwärtigen Muster werden in unseren Betrachtungen entweder vorrausgesetzt oder durch Ramsey-artige Argumente entwickelt. Überraschenderweise finden wir beide Phänomene in endlichen aber beliebig großen Graphen wieder.

In der Freudenthal-Kompaktifizierung $|G|$, eines lokal-endlichen Graphen G, treten Bögen und (topologische) Kreise auf natürlicher Weise als Limiten auf. Sie beschreiben dort Strukturen, die in endlichen Graphen von Wegen und Kreisen beschrieben werden. Zum Beispiel ist ein *Hamiltonkreis* in $|G|$ ein topologischer Kreis, der alle Ecken von G enthält. Eine umfangreiche Einleitung dazu gibt es in [2].

Diese topologischen Limiten werden in den drei Artikeln „Dual trees must share their ends" (S. 1), „Orthogonality and minimality in the homology of locally finite graphs" (S. 19) und „Extending cycles locally to Hamilton cycles" (S. 27) untersucht.

In dem Artikel „Transitivity conditions in infinite graphs" (S. 40) setzen wir ein hohes Maß an Symmetrie vorraus und zeigen, dass diese Vorrausetzung ausreicht um allgegenwärtige Muster zu erzwingen. In der Tat sind die charakterisierten Graphen nur aus wenigen Arten sehr einfacher Teilgraphen aufgebaut die überall auftreten. Da wir zeigen können, dass diese Teilgraphen in einer Baumstruktur zusammengefügt sind, können wir eine präzise strukturelle Charakterisierung der Grapheigenschaften geben, die vorher ausschließlich durch die Forderung nach hoher Symmetrie definiert wurden.

Der Artikel „Edge-disjoint double rays in infinite graphs: a Halin type result" (S. 73) bestätigt die Vermutung von Andreae [1], dass ein Graph der k kantendisjunkte Doppelstrahlen für jedes $k \in \mathbb{N}$ enthält, auch unendlich viele kantendisjunkte Doppelstrahlen enthält. Wir erhalten die unendliche Menge kantendisjunkter Doppeltstrahlen als einen Limes, wobei Lemma 29 (S. 86) ein Ramsey-artiges Phänomen beschreibt, dass in den Schritten der Limeskonstruktion nötig ist.

Die Artikel „Linkages in large graphs of bounded tree-width" und „Linkages in large graphs" (S. 89 und S. 155) beschäftigen sich mit endlichen aber beliebig großen Graphen. Interessanterweise können wir Limiten und Allgegenwärtigkeit hier wiederfinden. In beiden Artikeln untersuchen wir Gebiete von Graphen, die gegen Bäume konvergieren, wenn die Graphen größer werden. Ramsey-artige Argumente lassen uns in diesen beinahe Bäumen eine Struktur finden, deren einzelne Teile alle das selbe allgegenwärtige Muster aufweisen.

Jeder Artikel beginnt mit einer eigenen ausführlicheren Einleitung.

Literatur

[1] Th. Andreae. Über maximale Systeme von kantendisjunkten unendlichen Wegen in Graphen. *Math. Nachr.*, 101:219–228, 1981.

[2] R. Diestel. Locally finite graphs with ends: a topological approach. *Discrete Math.*, 310–312: 2750–2765 (310); 1423–1447 (311); 21–29 (312), 2010–11. arXiv:0912.4213.

Lebenslauf

Entfällt aus datenschutzrechtlichen Gründen.