Linkages in Large Graphs and Matroid Union

Dissertation

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

vorgelegt
im Fachbereich Mathematik
von

Jan-Oliver Fröhlich
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. Christian Reiher

Hamburg, den 19.03.2014

Prof. Dr. Armin Iske
Leiter des Fachbereichs Mathematik

Contents

1 Linkages in Large Graphs of Bounded Tree-Width **3**

1 Introduction . . . 3
2 Outline . . . 5
3 Stable decompositions . . . 7
4 Token movements . . . 19
5 Relinkages . . . 33
6 Rural societies . . . 38
7 Constructing a linkage . . . 52
8 Discussion . . . 63
9 References . . . 65

2 Infinite matroid union **68**

1 Introduction . . . 68
2 Preliminaries . . . 71
3 The union of arbitrary infinite matroids . . . 72
4 Matroid union . . . 74
5 Base packing in co-finitary matroids . . . 87
6 References . . . 88

3 On the intersection of infinite matroids **90**

1 Introduction . . . 90
2 Preliminaries . . . 94
3 From infinite matroid intersection to the infinite Menger theorem 95
4 From infinite matroid union to infinite matroid intersection . . 98
5 The graphic nearly finitary matroids . . . 101
6 References . . . 105

Appendices **106**

Summary . . . 106
Zusammenfassung . . . 107
Die Entwicklung der Linkage-Arbeit . . . 108
Academic CV . . . 110

Linkages in Large Graphs of Bounded Tree-Width

Jan-Oliver Fröhlich* Ken-ichi Kawarabayashi†‡
Theodor Müller§ Julian Pott¶ Paul Wollan‖

February 2014

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

*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.

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 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*

$$\kappa(G) \geq 2k+3, \qquad \mathrm{tw}(G) < w, \quad \textit{and} \quad |G| \geq N.$$

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 there 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 $B_i, B_{i+1}, \dots, B_{i+a}$ such that no vertex of $s_1, \dots, s_k, t_1, \dots, t_k$ 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 $s_1, \dots, s_k, t_1, \dots, t_k$ 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 $B_i \cup \dots \cup B_{i+a}$. 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 $(A_L, B_L) \leq (A', B')$ and $(A'', B'') \leq (A_R, B_R)$. To verify that $(A_L, B_L) \leq (A'', B'')$ we need to show $A_L \subseteq A' \cap A_R$ and $B_L \supseteq B' \cup B_R$. All required inclusions follow from $(A_L, B_L) \leq (A', B')$ and $(A_L, B_L) \leq (A_R, B_R)$. 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. $\square$

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 $S' - vP'w$.*

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.

[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.

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 $\mathcal{W} = (W_0, \dots, W_l)$ be an ordered tuple of subsets of $V(G)$. Then l is the *length* of $\mathcal{W}$, the sets W_i with $0 \le i \le l$ are its *bags*, and the sets $W_{i-1} \cap W_i$ with $1 \le i \le l$ are its *adhesion sets*. We refer to the bags W_i with $1 \le i \le l-1$ 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 $W_0, \dots, W_l$ where l is the length of $\mathcal{W}$. For indices $0 \le j \le k \le l$ 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 $0 \le i \le j \le k \le l$, 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 $P_1, \dots, P_q$ 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 $\mathcal{Q} = \{Q_1, \dots, Q_q\}$ 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 $\mathcal{Q}$ 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

[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.

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 $\lambda := \{\alpha \mid P_\alpha \text{ is non-tivial}\}$ and $\theta := \{\alpha \mid P_\alpha \text{ is trivial}\}$. 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 $\lambda := \{\alpha \mid P_\alpha \text{ is non-trivial}\}$.

(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 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 $i_1, \dots, i_n$ of $1, \dots, l$, the *contraction of* $(\mathcal{W},\mathcal{P})$ *along* $i_1, \dots, i_n$ is the pair $(\mathcal{W}',\mathcal{P}')$ defined as follows. We let $\mathcal{W}' := (W'_0, W'_1, \dots, W'_n)$ with $W'_0 := W_{[0,\ i_1-1]}$,

$$W'_j := W_{[i_j,\ i_{j+1}-1]} \quad \text{for} \quad j = 1, \dots, n-1,$$

$W_n := W_{[i_n,l]}$, 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 $i_j \leq i < i_{j+1}$ 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 $i_j \leq i < i_{j+1}$ 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 $i_j \leq i < i_{j+1}$. □

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 $W_{k+1}, \dots, W_{k+l-1}$ of $\mathcal{W}$ contains a disturbance. If not, then by Lemma 3.6, the contraction of $(\mathcal{W}, \mathcal{P})$ along the sequence $k + 1, k + 2, \dots, k + l$ 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 $k_0 := k, k_1, \dots, k_{w!} := k' + 1$, such that we have $k_j - k_{j-1} \geq l$ for $j \in \{1, \dots, w!\}$. For each $j \in \{0, \dots, w! - 1\}$ one of the at least $l-1$ consecutive inner bags $W_{k_j+1}, W_{k_j+2}, \dots, W_{k_{j+1}-1}$ 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 $j \in \{0, \dots, w! - 1\}$ 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 $\pi_{j_1} \circ \pi_{j_1-1} \ldots \circ \pi_{j_0} = \mathrm{id}$.

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 $j \in \{j_0, \dots, j_1\}$. Of all the restrictions of $\mathcal{Q}$ to the bags $W_k, \dots, W_{k'}$ 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,

[3] Let $(G, \cdot)$ be a group of order n and $g_1, \dots, g_n \in G$. 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.

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 $\square$

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 $(\mathcal{W}, \mathcal{P})$. $\square$

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

$$\big((X_0, \dots, X_n = Y_0, \dots, Y_m), (M_1, \dots M_n, N_1, \dots, N_m)\big)$$

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

$$\big((X_0 \cup Z, \dots, X_n \cup Z), (M_1, \dots M_n)\big)$$

is a movement and we denote it by $(\mathcal{X} \cup Z, \mathcal{M})$.

Let $(\mathcal{X}, \mathcal{M})$ be a movement. For $i = 1, \dots, n$ 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. $\square$

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 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.*

[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.

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. $\square$

Claim 4.10.2. *Suppose that* $|V(D)\backslash X| \geq 2$ *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 $z \in V(D)\backslash X$ and let $X' := (X\backslash\{y'\}) \cup \{z\}$. Let $\varphi : X \to X'$ be any bijection with $\varphi|_{X\backslash V(D)} = \mathrm{id}_{X\backslash V(D)}$ 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 $X'' := (X'\backslash\{y\}) \cup \{y'\}$ 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. □

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 $y_ly'y$ 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 $i = 1, \dots, k-1$ 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 $\Gamma(\mathcal{W}, \mathcal{P})$. $\square$

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. $\square$

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 $G^{\mathcal{Q}'}_D \subseteq G^{\mathcal{Q}}_D$, in particular, $G^{\mathcal{Q}}_D \subseteq G^{\mathcal{P}}_D$.*

Proof. Clearly $G^{\mathcal{Q}}_D$ and $G^{\mathcal{Q}'}_D$ are induced subgraphs of G so it suffices to show $V(G^{\mathcal{Q}'}_D) \subseteq V(G^{\mathcal{Q}}_D)$. Suppose for a contradiction that there is a vertex $w \in V(G^{\mathcal{Q}'}_D) \setminus V(G^{\mathcal{Q}}_D)$. We have $G^{\mathcal{Q}'}_D \cap \mathcal{Q}' = \mathcal{Q}'_{V(D)} \subseteq G^{\mathcal{Q}}_D$ so w must be an inner vertex of a $\mathcal{Q}'$-bridge B'. But w is in $G_\lambda - G^{\mathcal{Q}}_D$ 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 $G^{\mathcal{Q}}_D \supseteq \mathcal{Q}'_{V(D)}$. 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 $G^{\mathcal{Q}}_D \subseteq G^{\mathcal{P}}_D$ 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 $|N(\alpha) \cap \theta| \leq p - 3$ 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 $i = 1, \ldots, l - 1$. 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 $\mathcal{Q}'_{V(D)} \subseteq G^{\mathcal{Q}}_D$ 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 $G^{\mathcal{Q}}_{\lambda_0}$ such that $G^{\mathcal{Q}}_{\lambda_0} - v$ contains $|\lambda_0|$ disjoint paths from the first to the last adhesion set of $\mathcal{W}$ and v has a neighbour in $G_\lambda - G^{\mathcal{Q}}_{\lambda_0}$.

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 $G^{\mathcal{Q}'}_D$ 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 $G^{\mathcal{Q}'}_D$ such that v has a neighbour in $G_\lambda - G^{\mathcal{Q}'}_D$ and $G^{\mathcal{Q}'}_D - v$ contains an X–Y linkage $\mathcal{Q}''$ where X and Y denote the intersection of $V(G^{\mathcal{Q}'}_D)$ with the first and last adhesion set of $\mathcal{W}$, respectively.

By Lemma 5.4 we have $G^{\mathcal{Q}''}_D \subseteq G^{\mathcal{Q}'}_D \subseteq G^{\mathcal{Q}}_D$ and thus $\mathcal{Q}''$ is a $(\mathcal{Q}, V(D))$-relinkage as well. This implies $G^{\mathcal{Q}''}_D = G^{\mathcal{Q}'}_D$ by the minimality of $G^{\mathcal{Q}'}_D$. 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 $G_\lambda - G^{\mathcal{Q}'}_D$ and there is a w–$\mathcal{Q}'_{\lambda\setminus V(D)}$ path R that avoids $G^{\mathcal{Q}'}_D$. 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 $G^{\mathcal{Q}}_D$. Then the following statements are true.*

- *(i) Either $2|D| + |N(D) \cap \theta| \ge p$ or $V(G^{\mathcal{Q}}_D) = V(\mathcal{Q}_{V(D)})$ and $\kappa \ne \emptyset$.*
- *(ii) Either $2|D| + |N(D)| \ge p$ or there is $\alpha \in \kappa$ such that $|Q_\beta| \le |V \cap V(Q_\alpha)| + 1$ for all $\beta \in V(D) \setminus \kappa$.*

Note that $V(G^{\mathcal{Q}}_D) = V(\mathcal{Q}_{V(D)})$ 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 $G^{\mathcal{Q}}_D$ 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 $G^{\mathcal{Q}}_D$ that have a neighbour in $G_\lambda - G^{\mathcal{Q}}_D$. Clearly $Z \subseteq V(\mathcal{Q}_\kappa) \cup X \cup Y$. Moreover, $G^{\mathcal{Q}}_D - z$ 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 $G^{\mathcal{Q}}_D$ 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 $G^{\mathcal{Q}}_D$ 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 $G^{\mathcal{Q}}_D$, 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 $(X, V(G^{\mathcal{Q}}_D)) \in \mathcal{S}$ and $(V(G^{\mathcal{Q}}_D), Y) \in \mathcal{S}$ so for any vertex v of $G^{\mathcal{Q}}_D - (X \cup Y)$ 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 $G^{\mathcal{Q}}_D$. So $S_L \cup S_R \cup V(\mathcal{Q}_{N(D)\cap\theta})$ separates v from $G - G^{\mathcal{Q}}_D$ 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 $V(G^{\mathcal{Q}}_D) = Z_0$ Since every separator of a separation of $\mathcal{S}$ consists of one vertex from each path of $\mathcal{Q}_{V(D)}$ this means $V(\mathcal{Q}_{V(D)}) \subseteq V(G^{\mathcal{Q}}_D) = Z_0 \subseteq V(\mathcal{Q}_{V(D)})$. If $\kappa = \emptyset$, then $X \cup Y \cup V(\mathcal{Q}_{N(D)\cap\theta})$ separates $G^{\mathcal{Q}}_D - (X \cup Y)$ from $G - G^{\mathcal{Q}}_D$ 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 $\theta := \{\alpha \mid P_\alpha$ is trivial$\}$ and $\lambda := \{\alpha \mid P_\alpha$ 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. $\square$

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 $G^{\mathcal{P}}_{\alpha\beta} \cup \mathcal{P}_{N_{\alpha\beta}}$. $\square$

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. □

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. □

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. $\square$

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 $(G^{\mathcal{P}}_{\alpha\beta}, P_\alpha P_\beta^{-1})$. $\square$

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.

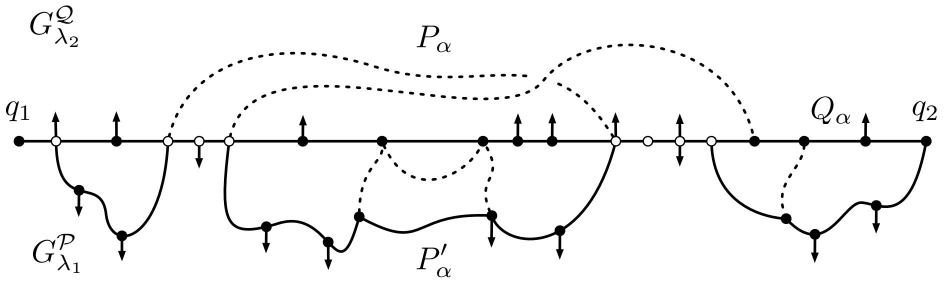

Figure 1: The graph $H = G^{\mathcal{P}}_{\lambda_2} \cap G^{\mathcal{Q}}_{\lambda_1}$.

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 dotted 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.

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 $\lambda_0 := \lambda$). □

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 $Q_\alpha \subseteq G^{\mathcal{Q}}_D \cap H \subseteq G^{\mathcal{Q}}_{\lambda_2} \cap G^{\mathcal{Q}}_{\lambda_1} = Q_\alpha$ 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 $G^{\mathcal{Q}}_D$-path and z_1 has a neighbour in $G_\lambda - G^{\mathcal{P}}_{\lambda_2} \subseteq G_\lambda - G^{\mathcal{P}}_D$ so P is an α-outlet.

Before we prove the converse implication let us show that $H \subseteq G^{\mathcal{P}}_D$. If some vertex v of $H \subseteq G^{\mathcal{P}}_{\lambda_2}$ is not in $G^{\mathcal{P}}_D$, 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 $G_\lambda - G^{\mathcal{P}}_D \subseteq G_\lambda - H$. 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. $\square$

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 $q'_1q'_2$, 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 $Q_\alpha r_1 R r_2 Q_\alpha$ 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 $P' := q'_1Pt_1P_\alpha t_2Pq'_2$ 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 $z_1 R z'_1$ contains all attachments of B' and no inner vertex of $z_1 R z'_1$ 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. □

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. $\square$

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'. $\square$

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 $b_1, \dots, b_k \in B_0$ 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 $i = 1, \dots, k$ 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. $\square$

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 $i = 1, \ldots, k$

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 $G_\Gamma^\mathcal{P}$. In other words, $G_\Gamma^\mathcal{P}$ contains at least $|Z||N_\Gamma(\alpha)|(2 + \varepsilon_\alpha)$ edges with one end on P_α and the other in $G_\Gamma^\mathcal{P} - P_\alpha$.

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 $G_{\Gamma_j}^\mathcal{P}$ 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 $G_{\Gamma_j}^\mathcal{P} - P_\alpha$.

By assumption and Claim 6.5.1 the path P_α separates $G_{\Gamma_i}^\mathcal{P}$ from $G_{\Gamma_j}^\mathcal{P}$ 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 $G_{\Gamma_i}^\mathcal{P} - P_\alpha \subseteq G_\lambda - G_{\Gamma_j}^\mathcal{P}$. Then the only reason that z is not also in Z_j is that it has no neighbour in $G_{\Gamma_j}^\mathcal{P} - P_\alpha$, 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 $G_{\Gamma_j}^\mathcal{P}$ 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 $G_D^\mathcal{P}$ that may have a neighbour in $G_\lambda - G_D^\mathcal{P}$ 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 $G_\lambda - G_D^\mathcal{P}$. 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

$$\begin{aligned}
x &= \sum_{v\in V} \left(d_{G^{\mathcal{P}}_D}(v) - d_{\mathcal{P}_{V(D)}}(v) \right) \\
&\leq \sum_{\alpha\in V(D)} \sum_{\beta\in N_D(\alpha)} \sum_{v\in V_\alpha} (d_{G^{\mathcal{P}}_{\alpha\beta}}(v) - d_{P_\alpha}(v)) \\
&= \sum_{\alpha\beta\in E(D)} \sum_{v\in V_\alpha\cup V_\beta} d_{G^{\mathcal{P}}_{\alpha\beta}}(v) - \sum_{\alpha\in V(D)} |N_D(\alpha)| \cdot (2|V_\alpha| - 2) \\
&\leq \sum_{\alpha\beta\in E(D)} (4|V_\alpha| + 4|V_\beta| - 6) - \sum_{\alpha\in V(D)} |N_D(\alpha)| \cdot (2|V_\alpha| - 2) \\
&= \sum_{\alpha\in V(D)} |N_D(\alpha)| \, (4|V_\alpha| - 3) - \sum_{\alpha\in V(D)} |N_D(\alpha)| \cdot (2|V_\alpha| - 2) \\
&< \sum_{\alpha\in V(D)} 2|N_D(\alpha)| \cdot |V_\alpha|.
\end{aligned}$$

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

$$\begin{aligned}
x &= \sum_{\alpha\in V(D)} \sum_{v\in V_\alpha} \left(d_{G^{\mathcal{P}}_D}(v) - d_{P_\alpha}(v) \right) \\
&\geq \sum_{\alpha\in V(D)} (|V_\alpha| - 2) \cdot |N_D(\alpha)| \cdot (2 + \varepsilon_\alpha) \\
&\geq \sum_{\alpha\in V(D)} |N_D(\alpha)| \cdot (2|V_\alpha| - 4 + 4k\varepsilon_\alpha) \\
&\geq \sum_{\alpha\in V(D)} 2|N_D(\alpha)| \cdot |V_\alpha|. \qquad \square
\end{aligned}$$

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 $G_\lambda - G^{\mathcal{Q}}_D$ and in $G^{\mathcal{Q}}_D - Q_\alpha$ have average degree at least $2 + |N_D(\alpha)|(2 + \varepsilon_\alpha)$ in $G^{\mathcal{Q}}_D$.

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 $G^{\mathcal{Q}}_D - Q_\alpha$ 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 $G^{\mathcal{Q}}_D$ 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 $G^{\mathcal{Q}}_D$. 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 $G^{\mathcal{Q}}_D$.

We want to show that the average degree in $G^{\mathcal{Q}}_D$ 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 $G^{\mathcal{Q}}_D$ 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 $G^{\mathcal{Q}}_D$ 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 $G^{\mathcal{Q}}_D$ 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 $(G^{\mathcal{Q}}_{\alpha\beta}, Q_\alpha Q_\beta^{-1})$ by suppressing all inner vertices of Q_α and Q_β that have degree 2 in $G^{\mathcal{Q}}_{\alpha\beta}$ 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 $G_\lambda - G^{\mathcal{Q}}_D$ and in $G^{\mathcal{Q}}_D - Q_\alpha$ have average degree less than $2 + |N_D(\alpha)|(2+\varepsilon_\alpha)$ in G^{Q}_D. 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 $H_1, \ldots, H_m$ the blocks of H that are not a single edge and for $i = 1, \ldots, m$ 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 $i \in \{1, \ldots, m\}$. 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

$$x = \sum_{v \in V} \left(d_H(v) - d_{P'_\alpha \cup Q_\alpha}(v) \right) = \sum_{i=1}^m \sum_{v \in V_i} \left(d_{H_i}(v) - d_{C_i}(v) \right) \leq \sum_{i=1}^m \left(2|V_i| - 6 \right).$$

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 $G_{\lambda_1}^{\mathcal{Q}}$ but not on Q_α and therefore not in $G_{\lambda_2}^{\mathcal{Q}}$. This already implies $d_{G_D^{\mathcal{P}}}(v) = d_{G_{\lambda_2}^{\mathcal{P}}}(v)$ for all $v \in V(H)$ since $G_D^{\mathcal{Q}} \subseteq G_D^{\mathcal{P}}$ and $G_{\lambda_2}^{\mathcal{P}} = H \cup G_{\lambda_2}^{\mathcal{Q}}$ (see the proof of Lemma 6.5 (i) for the latter identity). The next equality follows directly from the definition of H.

$$d_H(v) + d_{G_{\lambda_2}^{\mathcal{Q}}}(v) = d_{G_{\lambda_2}^{\mathcal{P}}}(v) + d_{Q_\alpha}(v) \quad \forall v \in V(H).$$

Denote by U_1 the set of inner vertices of P_α that have a neighbour in both $G_\lambda - G_D^{\mathcal{P}}$ and $G_D^{\mathcal{P}} - P_\alpha$ and by U_2 the set of inner vertices of Q_α that have a neighbour in both $G_\lambda - G_D^{\mathcal{Q}}$ and $G_D^{\mathcal{Q}} - Q_\alpha$. 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 $G_{\lambda_2}^{\mathcal{P}}$. Since $U_1 \subseteq V_P$ and α is rich in D this means that

$$\sum_{v \in V_P} d_{G_D^{\mathcal{P}}}(v) \geq |V_P| \left(2 + |N_D(\alpha)| \cdot (2 + \varepsilon_\alpha)\right).$$

Similarly, we have $U_2 \subseteq V_Q \subseteq V(Q_\alpha)$ and every vertex $v \in V_Q \setminus Z_2$ satisfies $d_{G_D^{\mathcal{Q}}}(v) = 2 = d_{Q_\alpha}(v)$. So by the assumption that α is not rich in D w.r.t. $\mathcal{Q}$ we have

$$\sum_{v \in V_Q} \Big(d_{G_D^{\mathcal{Q}}}(v) - d_{Q_\alpha}(v)\Big) < |V_Q| \cdot |N_D(\alpha)| \cdot (2 + \varepsilon_\alpha).$$

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

$$\begin{aligned}
x &\geq \sum_{v \in V'} \left(d_H(v) - d_{P'_\alpha \cup Q_\alpha}(v)\right) \\
&= \sum_{v \in V'} \Big(d_{G_D^{\mathcal{P}}}(v) - d_{G_D^{\mathcal{Q}}}(v) + d_{Q_\alpha}(v) - d_{P'_\alpha \cup Q_\alpha}(v)\Big) \\
&= \sum_{v \in V_P} d_{G_D^{\mathcal{P}}}(v) + \sum_{v \in V' \setminus V_P} d_{G_D^{\mathcal{P}}}(v) - \sum_{v \in V_Q} \Big(d_{G_D^{\mathcal{Q}}}(v) - d_{Q_\alpha}(v)\Big) - 2|V'| - 2m \\
&> |V_P| \cdot |N_D(\alpha)| \cdot (2 + \varepsilon_\alpha) + 2|V_P| + |V' \setminus V_P| \cdot (2|N(\alpha) \cap \lambda| + 3) \\
&\quad - |V_Q| \cdot |N_D(\alpha)| \cdot (2 + \varepsilon_\alpha) - 2|V'| - 2m \\
&= |V' \setminus V_Q| \cdot |N_D(\alpha)| \cdot (2 + \varepsilon_\alpha) + |V' \setminus V_P| \cdot |N(\alpha) \cap \lambda_1| \cdot (2 + \varepsilon_\alpha) - 2m \\
&> 2|V' \setminus V_Q| + 2|V' \setminus V_P| - 2m = \sum_{i=1}^{m} (2|V_i| - 6)
\end{aligned}$$

This shows that D is rich w.r.t. $\mathcal{Q}$ as defined above. So Claim 7.1.7 holds. $\square$

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. $\square$

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*

$$\kappa(G) \geq 2k + 1, \qquad \mathrm{tw}(G) \leq 2k + 10, \quad \textit{and} \quad |G| \geq N.$$

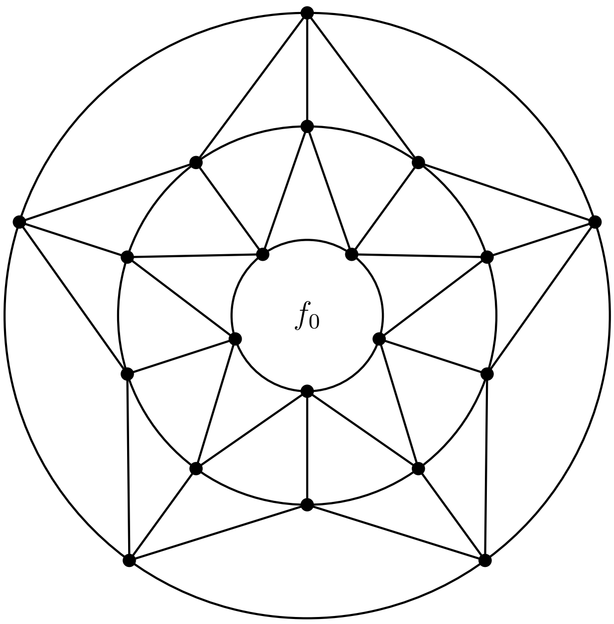

Figure 2: The 5-connected graph H_0 and its inner face f_0.

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.

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 whereP_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

$$\kappa(H) = 5, \qquad \mathrm{tw}(H) \leq 14, \quad \text{and} \quad |H| \geq N.$$

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). $\square$

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.

Infinite matroid union

Elad Aigner-Horev* Johannes Carmesin
Jan-Oliver Fröhlich

University of Hamburg
9 July 2012

Abstract

We consider the problem of determining whether the union of two infinite matroids is a matroid. We introduce a superclass of the finitary matroids, the *nearly finitary matroids*, and prove that the union of two nearly finitary matroids is a nearly finitary matroid.

On the other hand, we prove that the union of two arbitrary infinite matroids is not necessarily a matroid. Indeed, we show (under a weak additional assumption) that the nearly finitary matroids are essentially the largest class of matroids for which one can have a union theorem.

We then extend the base packing theorem for finite matroids to finite families of co-finitary matroids. This, in turn, yields a matroidal proof for the tree-packing results for infinite graphs due to Diestel and Tutte.

1 Introduction

Recently, Bruhn, Diestel, Kriesell, Pendavingh and Wollan [4] found axioms for infinite matroids in terms of independent sets, bases, circuits, closure and (relative) rank. These axioms allow for duality of infinite matroids as known from finite matroid theory, which settled an old problem of Rado. With these new axioms it is possible now to look which theorems of finite matroid theory have infinite analogues.

Here, we shall look at the *matroid union theorem* which is a classical result in finite matroid theory [6, 7]. It says that, given finite matroids $M_1 = (E_1, \mathcal{I}_1)$ and $M_2 = (E_2, \mathcal{I}_2)$, the set system

$$\mathcal{I}(M_1 \vee M_2) = \{I_1 \cup I_2 \mid I_1 \in \mathcal{I}_1,\ I_2 \in \mathcal{I}_2\} \tag{1}$$

*Research supported by the Minerva foundation.

is the set of independent sets of a matroid, the *union matroid* $M_1 \vee M_2$, and specifies a rank function for this matroid.

The matroid union theorem has important applications in finite matroid theory. For example, it can be used to provide short proofs for the *base covering and packing theorem* (discussed below more broadly), or to the *matroid intersection* theorem [6].

While the union of two finite matroids is always a matroid, it is not true that the union of two infinite matroids is always a matroid (see Proposition 1.1 below). The purpose of this paper is to study for which matroids their union is a matroid.

1.1 Our results

In this section, we outline our results with minimal background, deferring details until later sections. First we prove the following.

Proposition 1.1. *If M and N are infinite matroids, then $\mathcal{I}(M_1 \vee M_2)$ is not necessarily a matroid.*

One of the matroids involved in the proof of this proposition is finitary. Nevertheless, in Section 4.2, we establish a union theorem (see Theorem 1.2 below) for a superclass of the finitary matroids which we call *nearly finitary* matroids, defined next.

For any matroid M, taking as circuits only the finite circuits of M defines a (finitary) matroid with the same ground set as M. This matroid is called the *finitarization* of M and denoted by M^{fin}.

It is not hard to show that every basis B of M extends to a basis B^{fin} of M^{fin}, and conversely every basis B^{fin} of M^{fin} contains a basis B of M. Whether or not $B^{\text{fin}} \setminus B$ is finite will in general depend on the choices for B and B^{fin}, but given a choice for one of the two, it will no longer depend on the choice for the second one.

We call a matroid M *nearly finitary* if every base of its finitarization contains a base of M such that their difference is finite.

The class of nearly finitary matroids contains all finitary matroids, but not only. For example, the set system $\mathcal{C}(M) \cup \mathcal{B}(M)$ consisting of the circuits of an infinite-rank finitary matroid M together with its bases forms the set of circuits of a nearly finitary matroid that is not finitary (see Proposition 4.13). In [2] we characterize the graphic nearly finitary matroids; this also gives rise to numerous examples of nearly finitary matroids that are not finitary.

We show that the class of finitary matroids is closed under union (Section 4.2). In Section 4.3 we prove the same for the larger class of nearly finitary matroids, which is our main result:

Theorem 1.2. [Nearly finitary union theorem]
If M_1 and M_2 are nearly finitary matroids, then $M_1 \vee M_2$ is a matroid and in fact nearly finitary.

Theorem 1.2 is essentially best possible as follows.

First, the non-finitary matroid involved in the proof of Proposition 1.1 is a countable direct sum of infinite circuits and loops. This is essentially the simplest example of a matroid that is not nearly finitary.

Second, we show in Section 4.3.1 that for every matroid N that is not nearly finitary and that satisfies a (weak) additional assumption there exists a finitary matroid M such that $\mathcal{I}(M \vee N)$ is not a matroid. Thus in essence, not only is the class of nearly finitary matroids maximal with the property of having a union theorem; it is not even possible to add a matroid that is not nearly finitary to the class of finitary matroids without invalidating matroid union.

More precisely, we prove the following counterpart to Theorem 1.2.

Proposition 1.3. *Let N be a matroid that is not nearly finitary. Suppose that the finitarization of N has an independent set I containing only countably many N-circuits such that I has no finite subset meeting all of these circuits. Then there exists a finitary matroid M such that $\mathcal{I}(M \vee N)$ is not a matroid.*

A simple consequence of Theorem 1.2 is that $M_1 \vee \cdots \vee M_k$ is a nearly finitary matroid whenever $M_1, \ldots, M_k$ are nearly finitary. On the other hand, (by Observation 4.10) a countable union of nearly finitary matroids need not be a matroid.

In finite matroid theory, the *base covering* and *base packing* theorems are two well-known applications of the finite matroid union theorem. The former extends to finitary matroids in a straightforward manner (see Corollary 5.1).

In Section 5, we extend the finite base packing theorem to finite families of *co-finitary* matroids; i.e., matroids whose dual is finitary. The finite base packing theorem asserts that *a finite matroid M admits k disjoint bases if and only if $k \cdot \mathrm{rk}(X) + |E(M) \setminus X| \geq k \cdot \mathrm{rk}(M)$ for every $X \subseteq E(M)$* [7], where rk denotes the rank function of M. For infinite matroids, this rank condition is too crude. We reformulate this condition using the notion of *relative rank* introduced in [4] as follows: given two subsets $B \subseteq A \subseteq E(M)$, the *relative rank of A with respect to B* is denoted by $\mathrm{rk}(A|B)$, satisfies $\mathrm{rk}(A|B) \in \mathbb{N} \cup \{\infty\}$, and is given by

$$\mathrm{rk}(A|B) = \max\{|I \setminus J| : \ J \subseteq I, \ I \in \mathcal{I}(M) \cap 2^A, \ J \text{ maximal in } \mathcal{I}(M) \cap 2^B\}.$$

Theorem 1.4. *A co-finitary matroid M with ground set E admits k disjoint bases if and only if $|Y| \geq k \cdot \mathrm{rk}(E|E - Y)$ for all finite sets $Y \subseteq E$.*

Theorem 1.4 does not extend to arbitrary infinite matroids. Indeed, for every integer k there exists a finitary matroid with no three disjoint bases and satisfying $|Y| \geq k \cdot \mathrm{rk}(E|E-Y)$ for every $Y \subseteq E$ [1, 5].

This theorem gives a short matroidal proof of a result of Diestel and Tutte [5, Theorem 8.5.7] who showed that the well-known tree-packing theorem for finite graphs due to Nash-Williams and Tutte [5] extends to infinite graphs with so-called topological spanning trees.

2 Preliminaries

Notation and terminology for graphs are that of [5], for matroids that of [6, 4], and for topology that of [3].

Throughout, G always denotes a graph where $V(G)$ and $E(G)$ denote its vertex and edge sets, respectively. We write M to denote a matroid and write $E(M)$, $\mathcal{I}(M)$, $\mathcal{B}(M)$, and $\mathcal{C}(M)$ to denote its ground set, independent sets, bases, and circuits, respectively.

It will be convenient to have a similar notation for set systems. That is, for a set system $\mathcal{I}$ over some ground set E, an element of $\mathcal{I}$ is called *independent*, a maximal element of $\mathcal{I}$ is called a *base* of $\mathcal{I}$, and a minimal element of $\mathcal{P}(E) \setminus \mathcal{I}$ is called *circuit* of $\mathcal{I}$. A set system is *finitary* if an infinite set belongs to the system provided each of its finite subsets does; with this terminology, M is finitary provided that $\mathcal{I}(M)$ is finitary.

We review the definition of a matroid as given in [4]. A set system $\mathcal{I}$ is the set of independent sets of a matroid if it satisfies the following *independence axioms* [4]:

(I1) $\emptyset \in \mathcal{I}$.

(I2) $\lceil \mathcal{I} \rceil = \mathcal{I}$, that is, $\mathcal{I}$ is closed under taking subsets.

(I3) Whenever $I, I' \in \mathcal{I}$ with I' maximal and I not maximal, there exists an $x \in I' \setminus I$ such that $I + x \in \mathcal{I}$.

(IM) Whenever $I \subseteq X \subseteq E$ and $I \in \mathcal{I}$, the set $\{I' \in \mathcal{I} \mid I \subseteq I' \subseteq X\}$ has a maximal element.

In [4], an equivalent axiom system to the independence axioms is provided and is called the *circuit axioms system*; this axiom system characterises a matroid in terms of its circuits. Of these circuit axioms, we shall make frequent use of the so called *(infinite) circuit elimination axiom* phrased here for a matroid M:

(C) Whenever $X \subseteq C \in \mathcal{C}(M)$ and $\{C_x \mid x \in X\} \subseteq \mathcal{C}(M)$ satisfies $x \in C_y \Leftrightarrow x = y$ for all $x, y \in X$, then for every $z \in C \setminus \left(\bigcup_{x\in X} C_x\right)$ there exists a $C' \in \mathcal{C}(M)$ such that $z \in C' \subseteq \left(C \cup \bigcup_{x\in X} C_x\right) \setminus X$.

3 The union of arbitrary infinite matroids

In this section, we prove Proposition 1.1. That is, we show that there exists infinite matroids M and N whose union is not a matroid.

As the nature of M and N is crucial for establishing the tightness of Theorem 1.2, we prove Proposition 1.1 in two steps as follows.

In Claim 3.1, we treat the relatively simpler case in which M is finitary and N is co-finitary and both have uncountable ground sets. Second, then, in Claim 3.2, we refine the argument as to have M both finitary and co-finitary and N co-finitary and both on countable ground sets.

Claim 3.1. *There exists a finitary matroid M and a co-finitary matroid N such that $\mathcal{I}(M \vee N)$ is not a matroid.*

Proof. Set $E = E(M) = E(N) = \mathbb{N} \times \mathbb{R}$. Next, put $M := \bigoplus_{n\in\mathbb{N}} M_n$, where $M_n := U_{1,\{n\}\times\mathbb{R}}$. The matroid M is finitary as it is a direct sum of 1-uniform matroids. For $r \in \mathbb{R}$, let N_r be the circuit matroid on $\mathbb{N} \times \{r\}$; set $N := \bigoplus_{r\in\mathbb{R}} N_r$. As N is a direct sum of circuits, it is co-finitary. (see Figure 1).

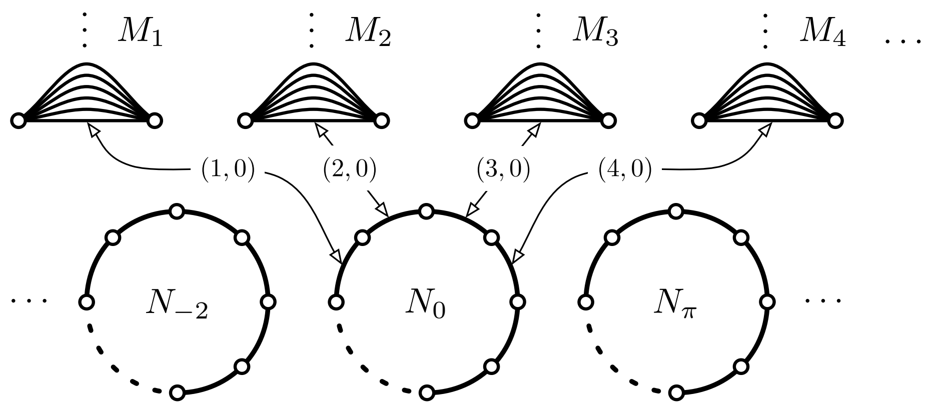

Figure 1: $M = \bigoplus_{n\in\mathbb{N}} M_n$ and $N = \bigoplus_{r\in\mathbb{R}} N_r$.

We show that $\mathcal{I}(M \vee N)$ violates the axiom (IM) for $I = \emptyset$ and $X = E$; so that $\mathcal{I}(M \vee N)$ has no maximal elements. It is sufficient to show that a set $J \subseteq E$ belongs to $\mathcal{I}(M \vee N)$ if and only if it contains at most countably many circuits of N. For if so, then for any $J \in \mathcal{I}(M \vee N)$ and any circuit $C = \mathbb{N} \times \{r\}$ of N with $C \nsubseteq J$ (such a circuit exists) we have $J \cup C \in \mathcal{I}(M \vee N)$.

The point to observe here is that every independent set of M is countable, (since every such set meets at most one element of M_n for each $n \in \mathbb{N}$), and that every independent set of N misses uncountably many elements of E (as any such set must miss at least one element of N_r for each $r \in \mathbb{R}$).

Suppose $J \subseteq E$ contains uncountably many circuits of N. Since each independent set of N misses uncountably many elements of E, every set $D = J \setminus J_N$ is uncountable whenever $J_N \in \mathcal{I}(J)$. On the other hand, since each independent set of M is countable, we have that $D \notin \mathcal{I}(M)$. Consequently, $J \notin \mathcal{I}(M \vee N)$, as required.

We may assume then that $J \subseteq E$ contains only countably many circuits of N, namely, $\{C_{r_1}, C_{r_2}, \ldots\}$. Now the set $J_M = \{(i, r_i) : i \in \mathbb{N}\}$ is independent in M; consequently, $J \setminus J_M$ is independent in N; completing the proof. □

We proceed with matroids on countable ground sets.

Claim 3.2. *There exist a matroid M that is both finitary and co-finitray, and a co-finitary matroid N whose common ground is countable such that $\mathcal{I}(M \vee N)$ is not a matroid.*

Proof. For the common ground set we take $E = (\mathbb{N} \times \mathbb{N}) \cup L$ where $L = \{\ell_1, \ell_2, \ldots\}$ is countable and disjoint to $\mathbb{N} \times \mathbb{N}$. The matroids N and M are defined as follows. For $r \in \mathbb{N}$, let N_r be the circuit matroid on $\mathbb{N} \times \{r\}$. Set N to be the matroid on E obtained by adding the elements of L to the matroid $\bigoplus_{r \in \mathbb{N}} N_r$ as loops. Next, for $n \in \mathbb{N}$, let M_n be the 1-uniform matroid on $(\{n\} \times \{1, 2, \ldots, n\}) \cup \{\ell_n\}$. Let M be the matroid obtained by adding to the matroid $\bigoplus_{n \in \mathbb{N}} M_n$ all the members of $E \setminus E(\bigoplus_{n \in \mathbb{N}} M_n)$ as loops

We show that $\mathcal{I}(M \vee N)$ violates the axiom (IM) for $I = \mathbb{N} \times \mathbb{N}$ and $X = E$. It is sufficient to show that

(a) $I \in \mathcal{I}(M \vee N)$; and that

(b) every set J satisfying $I \subset J \subseteq E$ is in $\mathcal{I}(M \vee N)$ if and only if it misses infinitely many elements of L.

To see that $I \in \mathcal{I}(M \vee N)$, note that the set $I_M = \{(n, n) \mid n \in \mathbb{N}\}$ is independent in M and meets each circuit $\mathbb{N} \times \{r\}$ of N. In particular, the set $I_N := (\mathbb{N} \times \mathbb{N}) \setminus I_M$ is independent in N, and therefore $I = I_M \cup I_N \in \mathcal{I}(M \vee N)$.

Let then J be a set satisfying $I \subseteq J \subseteq E$, and suppose, first, that $J \in \mathcal{I}(M \vee N)$. We show that J misses infinitely many elements of L.

There are sets $J_M \in \mathcal{I}(M)$ and $J_N \in \mathcal{I}(N)$ such that $J = J_M \cup J_N$. As J_N misses at least one element from each of the disjoint circuits of N in I, the set $D := I \setminus J_N$ is infinite. Moreover, we have that $D \subseteq J_M$, since $I \subseteq J$. In particular, there is an infinite subset $L' \subseteq L$ such that $D + l$ contains a circuit of M for every $\ell \in L'$. Indeed, for every $e \in D$ is contained in some M_{n_e}; let then $L' = \{\ell_{n_e} : e \in D\}$ and note that $L' \cap J = \emptyset$. This shows that J_M and L' are disjoint and thus J and L' are disjoint as well, and the assertion follows.

Suppose, second, that there exists a sequence $i_1 < i_2 < \dots$ such that J is disjoint from $L' = \{\ell_{i_r} : r \in \mathbb{N}\}$. We show that the superset $E \setminus L'$ of J is in $\mathcal{I}(M \vee N)$. To this end, set $D := \{(i_r, r) \mid r \in \mathbb{N}\}$. Then, D meets every circuit $\mathbb{N} \times \{r\}$ of N in I, so that the set $J_N := \mathbb{N} \times \mathbb{N} \setminus D$ is independent in N. On the other hand, D contains a single element from each M_n with $n \in L'$. Consequently, $J_M := (L \setminus L') \cup D \in \mathcal{I}(M)$ and therefore $E \setminus L' = J_M \cup J_N \in \mathcal{I}(M \vee N)$. □

While the union of two finitary matroids is a matroid, by Proposition 4.1, the same is not true for two co-finitary matroids.

Corollary 3.3. *The union of two co-finitary matroids is not necessarily a matroid.*

4 Matroid union

In this section, we prove Theorem 1.2. The main difficulty in proving this theorem is the need to verify that given two nearly finitary matroids M_1 and M_2, that the set system $\mathcal{I}(M_1 \vee M_2)$ satisfies the axioms (IM) and (I3).

To verify the (IM) axiom for the union of two nearly finitary matroids we shall require the following theorem, proved below in Section 4.2.

Proposition 4.1. *If M_1 and M_2 are finitary matroids, then $M_1 \vee M_2$ is a finitary matroid.*

To verify (IM) for the union of finitary matroids we use a compactness argument (see Section 4.2). More specifically, we will show that $\mathcal{I}(M_1 \vee M_2)$ is a finitary set system whenever M_1 and M_2 are finitary matroids. It is then an easy consequence of Zorn's lemma that all finitary set systems satisfy (IM).

The verification of axiom (I3) is dealt in a joint manner for both matroid families. In the next section we prove the following.

Proposition 4.2. *The set system $\mathcal{I}(M_1 \vee M_2)$ satisfies* (I3) *for any two matroids M_1 and M_2.*

Indeed, for finitary matroids, Proposition 4.2 is fairly simple to prove. We, however, require this proposition to hold for nearly finitary matroids as well. Consequently, we prove this proposition in its full generality, i.e., for any pair of matroids. In fact, it is interesting to note that the union of infinitely many matroids satisfies (I3); though the axiom (IM) might be violated as seen in Observation 4.10).

At this point it is insightful to note a certain difference between the union of finite matroids to that of finitary matroids in a more precise manner. By the finite matroid union theorem if M admits two disjoint bases, then the union of these bases forms a base of $M \vee M$. For finitary matroids the same assertion is false.

Claim 4.3. *There exists an infinite finitary matroid M with two disjoint bases whose union is not a base of the matroid $M \vee M$ as it is properly contained in the union of some other two bases.*

Proof. Consider the infinite one-sided ladder with every edge doubled, say H, and recall that the bases of $M_F(H)$ are the ordinary spanning trees of H. Figure 2 shows two pairs of disjoint bases of $M_F(H)$. However, the union of the lower pair properly contains that of the upper pair as it additionally contains the leftmost edge of H. $\square$

Clearly, a direct sum of infinitely many copies of H gives rise to an infinite sequence of unions of disjoint bases, each properly containing the previous one. In fact, one can construct a (single) matroid formed as the union of two nearly finitary matroids that admits an infinite properly nested sequence of unions of disjoint bases.

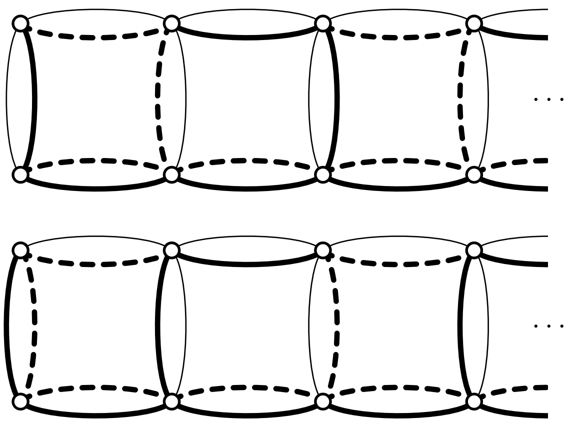

Figure 2: Two pairs of disjoint spanning trees of H (fat solid and fat dashed).

4.1 Exchange chains and the verification of axiom (I3)

In this section, we prove Proposition 4.2. Throughout this section M_1 and M_2 are matroids. It will be useful to show that the following variant of (I3) is satisfied.

Proposition 4.4. *The set $\mathcal{I} = \mathcal{I}(M_1 \vee M_2)$ satisfies the following.*

(I3') For all $I, B \in \mathcal{I}$ where B is maximal and all $x \in I \setminus B$ there exists $y \in B \setminus I$ such that $(I + y) - x \in \mathcal{I}$.

Observe that unlike in (I3), the set I in (I3') may be maximal.

We begin by showing that Proposition 4.4 implies Proposition 4.2.

Proof of Proposition 4.2 from Proposition 4.4. Let $I \in \mathcal{I}$ be non-maximal and $B \in \mathcal{I}$ be maximal. As I is non-maximal there is an $x \in E \setminus I$ such that $I + x \in \mathcal{I}$. We may assume $x \notin B$ or the assertion follows by (I2). By (I3'), applied to $I + x$, B, and $x \in (I + x) \setminus B$ there is $y \in B \setminus (I + x)$ such that $I + y \in \mathcal{I}$. $\square$

We proceed to prove Proposition 4.4. The following notation and terminology will be convenient. A circuit of M which contains a given set $X \subseteq E(M)$ is called an *X-circuit*.

By a *representation of* a set $I \in \mathcal{I}(M_1 \vee M_2)$, we mean a pair (I_1, I_2) where $I_1 \in \mathcal{I}(M_1)$ and $I_2 \in \mathcal{I}(M_2)$ such that $I = I_1 \cup I_2$.

For sets $I_1 \in \mathcal{I}(M_1)$ and $I_2 \in \mathcal{I}(M_2)$, and elements $x \in I_1 \cup I_2$ and $y \in E(M_1) \cup E(M_2)$ (possibly in $I_1 \cup I_2$), a tuple $Y = (y_0 = y, \ldots, y_n = x)$ with $y_i \neq y_{i+1}$ for all i is called an *even (I_1, I_2, y, x)-exchange chain* (or *even (I_1, I_2, y, x)-chain*) of *length n* if the following terms are satisfied.

(X1) For an even i, there exists a $\{y_i, y_{i+1}\}$-circuit $C_i \subseteq I_1 + y_i$ of M_1.

(X2) For an odd i, there exists a $\{y_i, y_{i+1}\}$-circuit $C_i \subseteq I_2 + y_i$ of M_2.

If $n \geq 1$, then (X1) and (X2) imply that $y_0 \notin I_1$ and that, starting with $y_1 \in I_1 \setminus I_2$, the elements y_i alternate between $I_1 \setminus I_2$ and $I_2 \setminus I_1$; the single exception being y_n which can lie in $I_1 \cap I_2$.

By an *odd exchange chain* (or *odd chain*) we mean an even chain with the words 'even' and 'odd' interchanged in the definition. Consequently, we say *exchange chain* (or *chain*) to refer to either of these notions. Furthermore, a subchain of a chain is also a chain; that is, given an (I_1, I_2, y_0, y_n)-chain $(y_0, \ldots, y_n)$, the tuple $(y_k, \ldots, y_l)$ is an (I_1, I_2, y_k, y_l)-chain for $0 \leq k \leq l \leq n$.

Lemma 4.5. *If there exists an (I_1, I_2, y, x)-chain, then $(I + y) - x \in \mathcal{I}(M_1 \vee M_2)$ where $I := I_1 \cup I_2$. Moreover, if $x \in I_1 \cap I_2$, then $I + y \in \mathcal{I}(M_1 \vee M_2)$.*

Remark. In the proof of Lemma 4.5 chains are used in order to alter the sets I_1 and I_2; the change is in a single element. Nevertheless, to accomplish this change, exchange chain of arbitrary length may be required; for instance, a chain of length four is needed to handle the configuration depicted in Figure 3.

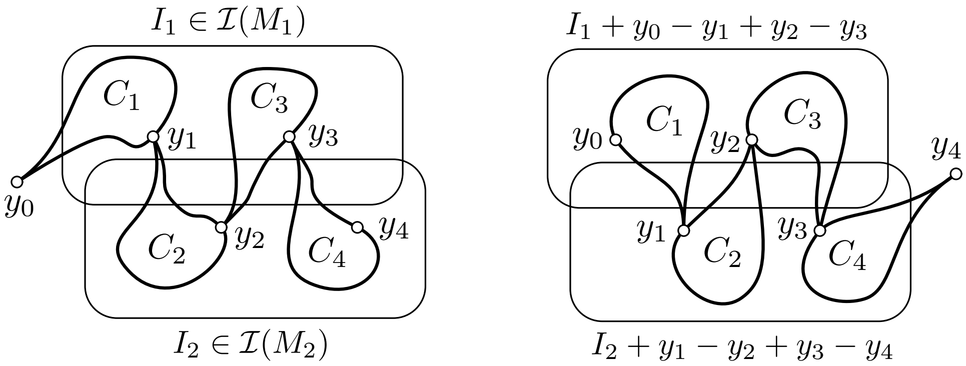

(a) the initial representation (b) the obtained representation

Figure 3: An even exchange chain of length 4.

Next, we prove Lemma 4.5.

Proof of Lemma 4.5. The proof is by induction on the length of the chain. The statement is trivial for chains of length 0. Assume $n \geq 1$ and that $Y = (y_0, \ldots, y_n)$ is a shortest (I_1, I_2, y, x)-chain. Without loss of generality, let Y be an even chain. If $Y' := (y_1, \ldots, y_n)$ is an (odd) (I_1', I_2, y_1, x)-chain where $I_1' := (I_1 + y_0) - y_1$, then $((I_1' \cup I_2) + y_1) - x \in \mathcal{I}(M_1 \vee M_2)$ by the induction hypothesis and the assertion follows, since $(I_1' \cup I_2) + y_1 = (I_1 \cup I_2) + y_0$. If also $x \in I_1 \cap I_2$, then either $x \in I_1' \cap I_2$ or $y_1 = x$ and hence $n = 1$. In the former case $I + y \in \mathcal{I}(M_1 \vee M_2)$ follows from the induction hypothesis and in the latter case $I + y = I_1' \cup I_2 \in \mathcal{I}(M_1 \vee M_2)$ as $x \in I_2$.

Since I_2 has not changed, (X2) still holds for Y', so to verify that Y' is an (I_1', I_2, y_1, x)-chain, it remains to show $I_1' \in \mathcal{I}(M_1)$ and to check (X1). To this end, let C_i be a $\{y_i, y_{i+1}\}$-circuit of M_1 in $I_1 + y_i$ for all even i. Such exist by (X1) for Y. Notice that any circuit of M_1 in $I_1 + y_0$ has to contain y_0 since $I_1 \in \mathcal{I}(M_1)$. On the other hand, two distinct circuits in $I_1 + y_0$ would give rise to a circuit contained in I_1 by the circuit elimination axiom applied to these two circuits, eliminating y_0. Hence C_0 is the unique circuit of M_1 in $I_1 + y_0$ and $y_1 \in C_0$ ensures $I_1' = (I_1 + y_0) - y_1 \in \mathcal{I}(M_1)$.

To see (X1), we show that there is a $\{y_i, y_{i+1}\}$-circuit C_i' of M_1 in $I_1' + y_i$ for every even $i \geq 2$. Indeed, if $C_i \subseteq I_1' + y_i$, then set $C_i' := C_i$; else, C_i contains an element of $I_1 \setminus I_1' = \{y_1\}$. Furthermore, $y_{i+1} \in C_i \setminus C_0$; otherwise

$(y_0, y_{i+1}, \ldots, y_n)$ is a shorter (I_1, I_2, y, x)-chain for, contradicting the choice of Y. Applying the circuit elimination axiom to C_0 and C_i, eliminating y_1 and fixing y_{i+1}, yields a circuit $C'_i \subseteq (C_0 \cup C_i) - y_1$ of M_1 containing y_{i+1}. Finally, as I'_1 is independent and $C'_i \setminus I'_1 \subseteq \{y_i\}$ it follows that $y_i \in C'_i$. □

We shall require the following. For $I_1 \in \mathcal{I}(M_1)$, $I_2 \in \mathcal{I}(M_2)$, and $x \in I_1 \cup I_2$, let

$$A(I_1, I_2, x) := \{a \mid \text{there exists an } (I_1, I_2, a, x)\text{-chain}\}.$$

This has the property that

$$\text{for every } y \notin A, \text{ either } I_1 + y \in \mathcal{I}(M_1) \text{ or the unique circuit } C_y \text{ of } M_1 \text{ in } I_1 + y \text{ is disjoint from } A. \tag{2}$$

To see this, suppose $I_1 + y \notin \mathcal{I}(M_1)$. Then there is a unique circuit C_y of M_1 in $I_1 + y$. If $C_y \cap A = \emptyset$, then the assertion holds so we may assume that $C_y \cap A$ contains an element, a say. Hence there is an (I_1, I_2, a, x)-chain $(y_0 = a, y_1, \ldots, y_{n-1}, y_n = x)$. As $a \in I_1$ this chain must be odd or have length 0, that is, $a = x$. Clearly, $(y, a, y_1, \ldots, y_{n-1}, x)$ is an even (I_1, I_2, y, x)-chain, contradicting the assumption that $y \notin A$.

Next, we prove Proposition 4.4.

Proof of Proposition 4.4. Let $B \in \mathcal{I}(M_1 \vee M_2)$ maximal, $I \in \mathcal{I}(M_1 \vee M_2)$, and $x \in I \setminus B$. Recall that we seek a $y \in B \setminus I$ such that $(I + y) - x \in \mathcal{I}(M_1 \vee M_2)$. Let (I_1, I_2) and (B_1, B_2) be representations of I and B, respectively. We may assume $I_1 \in \mathcal{B}(M_1|I)$ and $I_2 \in \mathcal{B}(M_2|I)$. We may further assume that for all $y \in B \setminus I$ the sets $I_1 + y$ and $I_2 + y$ are dependent in M_1 and M_2, respectively, for otherwise it holds that $I + y \in \mathcal{I}(M_1 \vee M_2)$ so that the assertion follows. Hence, for every $y \in (B \cup I) \setminus I_1$ there is a circuit $C_y \subseteq I_1 + y$ of M_1; such contains y and is unique since otherwise the circuit elimination axiom applied to these two circuits eliminating y yields a circuit contained in I_1, a contradiction.

If $A := A(I_1, I_2, x)$ intersects $B \setminus I$, then the assertion follows from Lemma 4.5. Else, $A \cap (B \setminus I) = \emptyset$, in which case we derive a contradiction to the maximality of B. To this end, set (Figure 4)

$$B'_1 := (B_1 \setminus b_1) \cup i_1 \qquad \text{where} \quad b_1 := B_1 \cap A \quad \text{and} \quad i_1 := I_1 \cap A$$

$$B'_2 := (B_2 \setminus b_2) \cup i_2 \qquad \text{where} \quad b_2 := B_2 \cap A \quad \text{and} \quad i_2 := I_2 \cap A$$

Since A contains x but is disjoint from $B \setminus I$, it holds that $(b_1 \cup b_2) + x \subseteq i_1 \cup i_2$ and thus $B + x \subseteq B'_1 \cup B'_2$. It remains to verify the independence of B'_1 and B'_2 in M_1 and M_2, respectively.

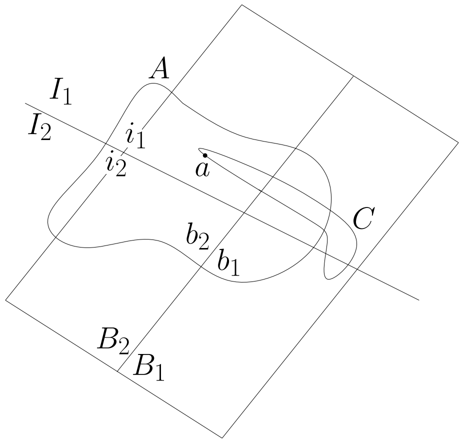

Figure 4: The independent sets I_1, at the top, and I_2, at the bottom, the bases B_1, on the right, and B_2, on the left, and their intersection with A.

Without loss of generality it is sufficient to show $B_1' \in \mathcal{I}(M_1)$. For the remainder of the proof 'independent' and 'circuit' refer to the matroid M_1. Suppose for a contradiction that the set B_1' is dependent, that is, it contains a circuit C. Since i_1 and $B_1 \setminus b_1$ are independent, neither of these contain C. Hence there is an element $a \in C \cap i_1 \subseteq A$. But $C \setminus I_1 \subseteq B_1 \setminus A$ and therefore no C_y with $y \in C \setminus I_1$ contains a by (2). Thus, applying the circuit elimination axiom on C eliminating all $y \in C \setminus I_1$ via C_y fixing a, yields a circuit in I_1, a contradiction. $\square$

Since in the proof of Proposition 4.4 the maximality of B is only used in order to avoid the case that $B+x \in \mathcal{I}(M_1 \vee M_2)$, one may prove the following slightly stronger statement.

Corollary 4.6. *For all $I, J \in \mathcal{I}(M_1 \vee M_2)$ and $x \in I \setminus J$, if $J+x \notin \mathcal{I}(M_1 \vee M_2)$, then there exists $y \in J \setminus I$ such that $(I+y)-x \in \mathcal{I}(M_1 \vee M_2)$.*

Next, the proof of Proposition 4.4, shows that for any maximal representation (I_1, I_2) of I there is $y \in B \setminus I$ such that exchanging finitely many elements of I_1 and I_2 gives a representation of $(I+y)-x$.

For subsequent arguments, it will be useful to note the following corollary. Above we used chains whose last element is fixed. One may clearly use chains whose first element is fixed. If so, then one arrives at the following.

Corollary 4.7. *For all $I, J \in \mathcal{I}(M_1 \vee M_2)$ and $y \in J \setminus I$, if $I+y \notin \mathcal{I}(M_1 \vee M_2)$, then there exists $x \in I \setminus J$ such that $(I+y)-x \in \mathcal{I}(M_1 \vee M_2)$.*

4.2 Finitary matroid union

In this section, we prove Proposition 4.1. In view of Proposition 4.2, it remains to show that $\mathcal{I}(M_1 \vee M_2)$ satisfies (IM) whenever M_1 and M_2 are finitary matroids.

The verification of (IM) for countable finitary matroids can be done using König's infinity lemma. Here, in order to capture matroids on any infinite ground set, we employ a topological approach. See [3] for the required topological background needed here.

We recall the definition of the product topology on $\mathcal{P}(E)$. The usual base of this topology is formed by the system of all sets

$$C(A, B) := \{X \subseteq E \mid A \subseteq X, B \cap X = \emptyset\},$$

where $A, B \subseteq E$ are finite and disjoint. Note that these sets are closed as well. Throughout this section, $\mathcal{P}(E)$ is endowed with the product topology and *closed* is used in the topological sense only.

We show that Proposition 4.1 can easily be deduced from Proposition 4.8 and Lemma 4.9, presented next.

Proposition 4.8. *Let $\mathcal{I} = \lceil \mathcal{I} \rceil \subseteq \mathcal{P}(E)$. The following are equivalent.*

4.8.1. $\mathcal{I}$ is finitary;

4.8.2. $\mathcal{I}$ is compact, in the subspace topology of $\mathcal{P}(E)$.

A standard compactness argument can be used in order to prove 4.8.1. Here, we employ a slightly less standard argument to prove 4.8.2 as well. Note that as $\mathcal{P}(E)$ is a compact Hausdorff space, assertion 4.8.2 is equivalent to the assumption that $\mathcal{I}$ is closed in $\mathcal{P}(E)$, which we use quite often in the following proofs.

Proof of Proposition 4.8. To deduce 4.8.2 from 4.8.1, we show that $\mathcal{I}$ is closed. Let $X \notin \mathcal{I}$. Since $\mathcal{I}$ is finitary, X has a finite subset $Y \notin \mathcal{I}$ and no superset of Y is in $\mathcal{I}$ as $\mathcal{I} = \lceil \mathcal{I} \rceil$. Therefore, $C(Y, \emptyset)$ is an open set containing X and avoiding $\mathcal{I}$ and hence $\mathcal{I}$ is closed.

For the converse direction, assume that $\mathcal{I}$ is compact and let X be a set such that all finite subsets of X are in $\mathcal{I}$. We show $X \in \mathcal{I}$ using the finite intersection property[1] of $\mathcal{P}(E)$. Consider the family $\mathcal{K}$ of pairs (A, B) where

[1] *The finite intersection property* ensures that an intersection over a family $\mathcal{C}$ of closed sets is non-empty if every intersection of finitely many members of $\mathcal{C}$ is.

$A \subseteq X$ and $B \subseteq E \setminus X$ are both finite. The set $C(A,B) \cap \mathcal{I}$ is closed for every $(A,B) \in \mathcal{K}$, as $C(A,B)$ and $\mathcal{I}$ are closed. If $\mathcal{L}$ is a finite subfamily of $\mathcal{K}$, then

$$\bigcup_{(A,B)\in\mathcal{L}} A \in \bigcap_{(A,B)\in\mathcal{L}} \left(C(A,B) \cap \mathcal{I}\right).$$

As $\mathcal{P}(E)$ is compact, the finite intersection property yields

$$\left(\bigcap_{(A,B)\in\mathcal{K}} C(A,B)\right) \cap \mathcal{I} = \bigcap_{(A,B)\in\mathcal{K}} \left(C(A,B) \cap \mathcal{I}\right) \neq \emptyset.$$

However, $\bigcap_{(A,B)\in\mathcal{K}} C(A,B) = \{X\}$. Consequently, $X \in \mathcal{I}$, as desired. $\square$

Lemma 4.9. *If $\mathcal{I}$ and $\mathcal{J}$ are closed in $\mathcal{P}(E)$, then so is $\mathcal{I} \vee \mathcal{J}$.*

Proof. Equipping $\mathcal{P}(E) \times \mathcal{P}(E)$ with the product topology, yields that Cartesian products of closed sets in $\mathcal{P}(E)$ are closed in $\mathcal{P}(E) \times \mathcal{P}(E)$. In particular, $\mathcal{I} \times \mathcal{J}$ is closed in $\mathcal{P}(E) \times \mathcal{P}(E)$. In order to prove that $\mathcal{I} \vee \mathcal{J}$ is closed, we note that $\mathcal{I} \vee \mathcal{J}$ is exactly the image of $\mathcal{I} \times \mathcal{J}$ under the union map

$$f : \mathcal{P}(E) \times \mathcal{P}(E) \to \mathcal{P}(E), \quad f(A,B) = A \cup B.$$

It remains to check that f maps closed sets to closed sets; which is equivalent to showing that f maps compact sets to compact sets as $\mathcal{P}(E)$ is a compact Hausdorff space. As continuous images of compact spaces are compact, it suffices to prove that f is continuous, that is, to check that the pre-images of subbase sets $C(\{a\}, \emptyset)$ and $C(\emptyset, \{b\})$ are open as can be seen here:

$$\begin{aligned} f^{-1}(C(\{a\},\emptyset)) &= (C(\{a\},\emptyset) \times \mathcal{P}(E)) \cup (\mathcal{P}(E) \times C(\{a\},\emptyset)) \\ f^{-1}(C(\emptyset,\{b\})) &= C(\emptyset,\{b\}) \times C(\emptyset,\{b\}). \end{aligned}$$

$\square$

Next, we prove Proposition 4.1.

Proof of Proposition 4.1. By Proposition 4.2 it remains to show that the union $\mathcal{I}(M_1) \vee \mathcal{I}(M_2)$ satisfies (IM). As all finitary set systems satisfy (IM), by Zorn's lemma, it is sufficient to show that $\mathcal{I}(M_1 \vee M_2)$ is finitary. By Proposition 4.8, $\mathcal{I}(M_1)$ and $\mathcal{I}(M_2)$ are both compact and thus closed in $\mathcal{P}(E)$, yielding, by Lemma 4.9, that $\mathcal{I}(M_1) \vee \mathcal{I}(M_2)$ is closed in $\mathcal{P}(E)$, and thus compact. As $\mathcal{I}(M_1) \vee \mathcal{I}(M_2) = \lceil \mathcal{I}(M_1) \vee \mathcal{I}(M_2) \rceil$, Proposition 4.8 asserts that $\mathcal{I}(M_1) \vee \mathcal{I}(M_2)$ is finitary, as desired. $\square$

We conclude this section with the following observation.

Observation 4.10. *A countable union of finitary matroids need not be a matroid.*

Proof. We show that for any integer $k \geq 1$, the set system

$$\mathcal{I} := \bigvee_{n \in \mathbb{N}} U_{k,\mathbb{R}}$$

is not a matroid, where here $U_{k,\mathbb{R}}$ denotes the k-uniform matroid with ground set $\mathbb{R}$.

Since a countable union of finite sets is countable, we have that the members of $\mathcal{I}$ are the countable subsets of $\mathbb{R}$. Consequently, the system $\mathcal{I}$ violates the (IM) axiom for $I = \emptyset$ and $X = \mathbb{R}$. $\square$

Above, we used the fact that the members of $\mathcal{I}$ are countable and that the ground set is uncountable. One can have the following more subtle example, showing that a countable union of finite matroids need not be a matroid.

Let $A = \{a_1, a_2, \ldots\}$ and $B = \{b_1, b_2, \ldots\}$ be disjoint countable sets, and for $n \in \mathbb{N}$, set $E_n := \{a_1, \ldots, a_n\} \cup \{b_n\}$. Then $\bigvee_{n \in \mathbb{N}} U_{1,E_n}$ is an infinite union of finite matroids and fails to satisfy (IM) for $I = A$ and $X = A \cup B = E(M)$.

4.3 Nearly finitary matroid union

In this section, we prove Theorem 1.2.

For a matroid M, let $\mathcal{I}^{\text{fin}}(M)$ denote the set of subsets of $E(M)$ containing no finite circuit of M, or equivalently, the set of subsets of $E(M)$ which have all their finite subsets in $\mathcal{I}(M)$. We call $M^{\text{fin}} = (E(M), \mathcal{I}^{\text{fin}}(M))$ the *finitarization* of M. With this notation, a matroid M is *nearly finitary* if it has the property that

$$\text{for each } J \in \mathcal{I}(M^{\text{fin}}) \text{ there exists an } I \in \mathcal{I}(M) \text{ such that } |J \setminus I| < \infty. \quad (3)$$

For a set system $\mathcal{I}$ (not necessarily the independent sets of a matroid) we call a maximal member of $\mathcal{I}$ a *base* and a minimal member subject to not being in $\mathcal{I}$ a *circuit*. With these conventions, the notions of *finitarization* and *nearly finitary* carry over to set systems.

Let $\mathcal{I} = \lceil \mathcal{I} \rceil$. The finitarization $\mathcal{I}^{\text{fin}}$ of $\mathcal{I}$ has the following properties.

1. $\mathcal{I} \subseteq \mathcal{I}^{\text{fin}}$ with equality if and only if $\mathcal{I}$ is finitary.

2. $\mathcal{I}^{\text{fin}}$ is finitary and its circuits are exactly the finite circuits of $\mathcal{I}$.

3. $(\mathcal{I}|X)^{\text{fin}} = \mathcal{I}^{\text{fin}}|X$, in particular $\mathcal{I}|X$ is nearly finitary if $\mathcal{I}$ is.

The first two statements are obvious. To see the third, assume that $\mathcal{I}$ is nearly finitary and that $J \in \mathcal{I}^{\text{fin}}|X \subseteq \mathcal{I}^{\text{fin}}$. By definition there is $I \in \mathcal{I}$ such that $J \setminus I$ is finite. As $J \subseteq X$ we also have that $J \setminus (I \cap X)$ is finite and clearly $I \cap X \in \mathcal{I}|X$.

Proposition 4.11. *The pair $M^{\text{fin}} = (E, \mathcal{I}^{\text{fin}}(M))$ is a finitary matroid, whenever M is a matroid.*

Proof. By construction, the set system $\mathcal{I}^{\text{fin}} = \mathcal{I}(M^{\text{fin}})$ satisfies the axioms (I1) and (I2) and is finitary, implying that it also satisfies (IM).

It remains to show that $\mathcal{I}^{\text{fin}}$ satisfies $(I3)$. By definition, a set $X \subseteq E(M)$ is not in $\mathcal{I}^{\text{fin}}$ if and only if it contains a finite circuit of M.

Let $B, I \in \mathcal{I}^{\text{fin}}$ where B is maximal and I is not, and let $y \in E(M) \setminus I$ such that $I + y \in \mathcal{I}^{\text{fin}}$. If $I + x \in \mathcal{I}^{\text{fin}}$ for any $x \in B \setminus I$, then we are done.

Assuming the contrary, then $y \notin B$ and for any $x \in B \setminus I$ there exists a finite circuit C_x of M in $I + x$ containing x. By maximality of B, there exists a finite circuit C of M in $B + y$ containing y. By the circuit elimination axiom (in M) applied to the circuits C and $\{C_x\}_{x \in X}$ where $X := C \cap (B \setminus I)$, there exists a circuit

$$D \subseteq \left(C \cup \bigcup_{x \in X} C_x \right) \setminus X \subseteq I + y$$

of M containing $y \in C \setminus \bigcup_{x \in X} C_x$. The circuit D is finite, since the circuits C and $\{C_x\}$ are; this contradicts $I + y \in \mathcal{I}^{\text{fin}}$. □

Proposition 4.12. *For arbitrary matroids M_1 and M_2 it holds that*

$$\mathcal{I}(M_1^{\text{fin}} \vee M_2^{\text{fin}}) = \mathcal{I}(M_1^{\text{fin}} \vee M_2^{\text{fin}})^{\text{fin}} = \mathcal{I}(M_1 \vee M_2)^{\text{fin}}.$$

Proof. By Proposition 4.11, the matroids M_1^{fin} and M_2^{fin} are finitary and therefore $M_1^{\text{fin}} \vee M_2^{\text{fin}}$ is a finitary as well, by Proposition 4.1. This establishes the first equality.

The second equality follows from the definition of finitarization provided we show that the finite members of $\mathcal{I}(M_1^{\text{fin}} \vee M_2^{\text{fin}})$ and $\mathcal{I}(M_1 \vee M_2)$ are the same.

Since $\mathcal{I}(M_1) \subseteq \mathcal{I}(M_1^{\text{fin}})$ and $\mathcal{I}(M_2) \subseteq \mathcal{I}(M_2^{\text{fin}})$ it holds that $\mathcal{I}(M_1^{\text{fin}} \vee M_2^{\text{fin}}) \supseteq \mathcal{I}(M_1 \vee M_2)$. On the other hand, a finite set $I \in \mathcal{I}(M_1^{\text{fin}} \vee M_2^{\text{fin}})$ can be written as $I = I_1 \cup I_2$ with $I_1 \in \mathcal{I}(M_1^{\text{fin}})$ and $I_2 \in \mathcal{I}(M_2^{\text{fin}})$ finite. As I_1 and I_2 are finite, $I_1 \in \mathcal{I}(M_1)$ and $I_2 \in \mathcal{I}(M_2)$, implying that $I \in \mathcal{I}(M_1 \vee M_2)$. □

With the above notation a matroid M is nearly finitary if each base of M^{fin} contains a base of M such that their difference is finite. The following is probably the most natural manner to construct nearly finitary matroids (that are not finitary) from finitary matroids.

For a matroid M and an integer $k \geq 0$, set $M[k] := (E(M), \mathcal{I}[k])$, where

$$\mathcal{I}[k] := \{I \in \mathcal{I}(M) \mid \exists J \in \mathcal{I}(M) \text{ such that } I \subseteq J \text{ and } |J \setminus I| = k\}.$$

Proposition 4.13. *If* $\operatorname{rk}(M) \geq k$*, then* $M[k]$ *is a matroid.*

Proof. The axiom (I1) holds as $\operatorname{rk}(M) \geq k$; the axiom (I2) holds as it does in M. For (I3) let $I', I \in \mathcal{I}(M[k])$ such that I' is maximal and I is not. There is a set $F' \subseteq E(M) \setminus I'$ of size k such that, in M, the set $I' \cup F'$ is not only independent but, by maximality of I', also a base. Similarly, there is a set $F \subseteq E(M) \setminus I$ of size k such that $I \cup F \in \mathcal{I}(M)$.

We claim that $I \cup F$ is non-maximal in $\mathcal{I}(M)$ for any such F. Suppose not and $I \cup F$ is maximal for some F as above. By assumption, I is contained in some larger set of $\mathcal{I}(M[k])$. Hence there is a set $F^+ \subseteq E(M) \setminus I$ of size $k+1$ such that $I \cup F^+$ is independent in M. Clearly $(I \cup F) \setminus (I \cup F^+) = F \setminus F^+$ is finite, so Lemma 4.14 implies that

$$\left|F^+ \setminus F\right| = \left|(I \cup F^+) \setminus (I \cup F)\right| \leq \left|(I \cup F) \setminus (I \cup F^+)\right| = \left|F \setminus F^+\right|.$$

In particular, $k+1 = |F^+| \leq |F| = k$, a contradiction.

Hence we can pick F such that $F \cap F'$ is maximal and, as $I \cup F$ is non-maximal in $\mathcal{I}(M)$, apply (I3) in M to obtain a $x \in (I' \cup F') \setminus (I \cup F)$ such that $(I \cup F) + x \in \mathcal{I}(M)$. This means $I + x \in \mathcal{I}(M[k])$. And $x \in I' \setminus I$ follows, as $x \notin F'$ by our choice of F.

To show (IM), let $I \subseteq X \subseteq E(M)$ with $I \in \mathcal{I}(M[k])$ be given. By (IM) for M, there is a $B \in \mathcal{I}(M)$ which is maximal subject to $I \subseteq B \subseteq X$. We may assume that $F := B \setminus I$ has at most k elements; for otherwise there is a superset $I' \subseteq B$ of I such that $|B \setminus I'| = k$ and it suffices to find a maximal set containing $I' \in \mathcal{I}(M[k])$ instead of I.

We claim that for any $F^+ \subseteq X \setminus I$ of size $k+1$ the set $I \cup F^+$ is not in $\mathcal{I}(M[k])$. For a contradiction, suppose it is. Then in $M|X$, the set $B = I \cup F$ is a base and $I \cup F^+$ is independent and as $(I \cup F) \setminus (I \cup F^+) \subseteq F \setminus F^+$ is finite, Lemma 4.14 implies

$$\left|F^+ \setminus F\right| = \left|(I \cup F^+) \setminus (I \cup F)\right| \leq \left|(I \cup F) \setminus (I \cup F^+)\right| = \left|F \setminus F^+\right|.$$

This means $k+1 = |F^+| \leq |F| = k$, a contradiction. So by successively adding single elements of $X \setminus I$ to I as long as the obtained set is still in $\mathcal{I}(M[k])$ we arrive at the wanted maximal element after at most k steps. $\square$

We conclude this section with a proof of Theorem 1.2. To this end, we shall require following two lemmas.

Lemma 4.14. *Let M be a matroid and $I, B \in \mathcal{I}(M)$ with B maximal and $B \setminus I$ finite. Then, $|I \setminus B| \leq |B \setminus I|$.*

Proof. The proof is by induction on $|B \setminus I|$. For $|B \setminus I| = 0$ we have $B \subseteq I$ and hence $B = I$ by maximality of B. Now suppose there is $y \in B \setminus I$. If $I + y \in \mathcal{I}$ then by induction

$$|I \setminus B| = |(I + y) \setminus B| \leq |B \setminus (I + y)| = |B \setminus I| - 1$$

and hence $|I \setminus B| < |B \setminus I|$. Otherwise there exists a unique circuit C of M in $I + y$. Clearly C cannot be contained in B and therefore has an element $x \in I \setminus B$. Then $(I + y) - x$ is independent, so by induction

$$|I \setminus B| - 1 = |((I + y) - x) \setminus B| \leq |B \setminus ((I + y) - x)| = |B \setminus I| - 1,$$

and hence $|I \setminus B| \leq |B \setminus I|$. □

Lemma 4.15. *Let $\mathcal{I} \subseteq \mathcal{P}(E)$ be a nearly finitary set system satisfying (I1), (I2), and the following variant of (I3):*

- *(*) For all $I, J \in \mathcal{I}$ and all $y \in I \setminus J$ with $J + y \notin \mathcal{I}$ there exists $x \in J \setminus I$ such that $(J + y) - x \in \mathcal{I}$.*

Then $\mathcal{I}$ satisfies (IM).

Proof. Let $I \subseteq X \subseteq E$ with $I \in \mathcal{I}$. As $\mathcal{I}^{\text{fin}}$ satisfies (IM) there is a set $B^{\text{fin}} \in \mathcal{I}^{\text{fin}}$ which is maximal subject to $I \subseteq B^{\text{fin}} \subseteq X$ and being in $\mathcal{I}^{\text{fin}}$. As $\mathcal{I}$ is nearly finitary, there is $J \in \mathcal{I}$ such that $B^{\text{fin}} \setminus J$ is finite and we may assume that $J \subseteq X$. Then, $I \setminus J \subseteq B^{\text{fin}} \setminus J$ is finite so that we may choose a J minimizing $|I \setminus J|$. If there is a $y \in I \setminus J$, then by (*) we have $J + y \in \mathcal{I}$ or there is an $x \in J \setminus I$ such that $(J + y) - x \in \mathcal{I}$. Both outcomes give a set containing more elements of I and hence contradicting the choice of J.

It remains to show that J can be extended to a maximal set B of $\mathcal{I}$ in X. For any superset $J' \in \mathcal{I}$ of J, we have $J' \in \mathcal{I}^{\text{fin}}$ and $B^{\text{fin}} \setminus J'$ is finite as it is a subset of $B^{\text{fin}} \setminus J$. As $\mathcal{I}^{\text{fin}}$ is a matroid, Lemma 4.14 implies

$$|J' \setminus B^{\text{fin}}| \leq |B^{\text{fin}} \setminus J'| \leq |B^{\text{fin}} \setminus J|.$$

Hence, $|J' \setminus J| \leq 2|B^{\text{fin}} \setminus J| < \infty$. Thus, we can greedily add elements of X to J to obtain the wanted set B after finitely many steps. □

Next, we prove Theorem 1.2.

Proof of Theorem 1.2. By Proposition 4.4, in order to prove that $M_1 \vee M_2$ is a matroid, it is sufficient to prove that $\mathcal{I}(M_1 \vee M_2)$ satisfies (IM). By Corollary 4.7 and Lemma 4.15 it remains to show that $\mathcal{I}(M_1 \vee M_2)$ is nearly finitary.

So let $J \in \mathcal{I}(M_1 \vee M_2)^{\text{fin}}$. By Proposition 4.12 we may assume that $J = J_1 \cup J_2$ with $J_1 \in \mathcal{I}(M_1^{\text{fin}})$ and $J_2 \in \mathcal{I}(M_2^{\text{fin}})$. By assumption there are $I_1 \in \mathcal{I}(M_1)$ and $I_2 \in \mathcal{I}(M_2)$ such that $J_1 \setminus I_1$ and $J_2 \setminus I_2$ are finite. Then $I = I_1 \cup I_2 \in \mathcal{I}(M_1 \vee M_2)$ and the assertion follows as $J \setminus (I_1 \cup I_2) \subseteq (J_1 \setminus I_1) \cup (J_2 \setminus I_2)$ is finite. □

4.3.1 Unions of non-nearly finitary matroids

In this section, we prove Proposition 1.3 asserting that a certain family of non-nearly finitary matroids does not admit a union theorem.

A matroid N is non-nearly finitary provided it has a set $I \in \mathcal{I}(N^{\text{fin}})$ with the property that no finite subset of I meets all the necessarily infinite circuits of N in I. If we additionally assume that there is one such I which contains only countably many circuits, then there exists a finitary matroid M such that $\mathcal{I}(M \vee N)$ is not a matroid.

Proof of Proposition 1.3. For N and I as in Proposition 1.3 choose an enumeration $C_1, C_2, \ldots$ of the circuits of N in I. We may assume that $I = \bigcup_{n \in \mathbb{N}} C_n$. There exist countably many disjoint subsets $Y_1, Y_2, \ldots$ of I satisfying

1. $|Y_n| \leq n$ for all $n \in \mathbb{N}$; and
2. $Y_n \cap C_i \neq \emptyset$ for all $n \in \mathbb{N}$ and all $1 \leq i \leq n$.

We construct the above sets as follows. Suppose $Y_1, \ldots, Y_n$ have already been defined. Let Y_{n+1} be a set of size at most $n+1$ disjoint to each of $Y_1, \ldots, Y_n$ and meeting the circuits $C_1, \ldots, C_{n+1}$; such exists as $\bigcup_{i=1}^{n} Y_i$ is finite and all circuits in I are infinite.

Let $L = \{l_1, l_2, \ldots\}$ be a countable set disjoint from $E(N)$. For each $n \in \mathbb{N}$ let M_n be the 1-uniform matroid on $Y_n \cup \{l_n\}$, i.e. $M_n := U_{1, Y_n \cup \{l_n\}}$. Then, $M := \bigoplus_{n \in \mathbb{N}} M_n$ is a direct sum of finite matroids and hence finitary.

We contend that $I \in \mathcal{I}(M \vee N)$ and that $\mathcal{I}(M \vee N)$ violates (IM) for I and $X := I \cup L$. By construction, Y_n contains some element d_n of C_n, for every $n \in \mathbb{N}$. So that $J_M = \{d_1, d_2, \ldots\}$ meets every circuit of N in I and is independent in M. This means that $J_N := I \setminus J_M \in \mathcal{I}(N)$ and thus $I = J_M \cup J_N \in \mathcal{I}(M \vee N)$.

It is now sufficient to show that a set J satisfying $I \subseteq J \subseteq X$ is in $\mathcal{I}(M \vee N)$ if and only if it misses infinitely many elements $L' \subseteq L$. Suppose that $J \in \mathcal{I}(M \vee N)$. There are sets $J_M \in \mathcal{I}(M)$ and $J_N \in \mathcal{I}(N)$ such that $J = J_M \cup J_N$. As $D := I \setminus J_N$ meets every circuit of N in I by independence of J_N, the set D is infinite. But $I \subseteq J$ and hence $D \subseteq J_M$. Let A be the set of all integers n such that $Y_n \cap D \neq \emptyset$. As Y_n is finite for every $n \in \mathbb{N}$, the set A must be infinite and so is $L' := \{l_n \mid n \in A\}$. Since J_M is independent in M and any element of L' forms a circuit of M with some element of J_M, we have $J_M \cap L' = \emptyset$ and thus $J \cap L' = \emptyset$ as no independent set of N meets L.

Suppose that there is a sequence $i_1 < i_2 < \dots$ such that J is disjoint from $L' = \{l_{i_n} \mid n \in \mathbb{N}\}$. We show that the superset $X \setminus L'$ of J is in $\mathcal{I}(M \vee N)$. By construction, for every $n \in \mathbb{N}$, the set Y_{i_n} contains an elements d_n of C_n. Set $D := \{d_n \mid n \in \mathbb{N}\}$. Then D meets every circuit of N in I, so $J_N := I \setminus D$ is independent in N. On the other hand, D contains exactly one element of each M_n with $n \in L'$. So $J_M := (L \setminus L') \cup D \in \mathcal{I}(M)$ and therefore $X \setminus L' = J_M \cup J_N \in \mathcal{I}(M \vee N)$. □

It is not known wether or not the proposition remains true if we drop the requirement that there are only countable many circuits in I.

5 Base packing in co-finitary matroids

In this section, we prove Theorem 1.4, which is a base packing theorem for co-finitary matroids.

Proof of Theorem 1.4. As the 'only if' direction is trivial, it remains to show the 'if' direction. For a matroid N and natural numbers k, c put

$$\mathcal{I}[N, k, c] := \{X \subseteq E(N) \mid \exists I_1, ..., I_k \in \mathcal{I}(N) \text{ with } g_c(I_1, ..., I_k) = X\},$$

where $g_c(I_1, ..., I_k) := \{e : |\{j : e \in I_j\}| \geq c\}$. The matroid M has k disjoint spanning sets if and only if M^* has k independent sets such that every element of E is in at least $k-1$ of those independent sets. Put another way, M has k disjoint bases if and only if

$$\mathcal{I}[M^*, k, k-1] = \mathcal{P}(E). \tag{4}$$

As M^* is finitary, $\mathcal{I}[M^*, k, k-1]$ is finitary by an argument similar to that in the proof of Lemma 4.9; here one may define

$$f : \mathcal{P}(E)^k \to \mathcal{P}(E); f(A_1, ..., A_k) = g_{k-1}(A_1, ..., A_k),$$

and repeat the above argument.

Thus, it suffices to show that every finite set Y is in $\mathcal{I}[M^*, k, k-1]$. To this end, it is sufficient to find k independent sets of M^* such that every element of Y is in at least $k-1$ of those; complements of which are M-spanning sets $S_1, ..., S_k$ such that these are disjoint if restricted to Y. To this end, we show that there are disjoint spanning sets $S'_1, ..., S'_k$ of $M.Y$ and set $S_i := S'_i \cup (E-Y)$. Since Theorem 1.4 is true for finite matroids [6], the sets $S'_1, ..., S'_k$ exist if and only if $|Z| \geq k \cdot \mathrm{rk}_{M.Y}(Y|Y-Z)$ for all $Z \subseteq Y$. As $|Z| \geq k \cdot \mathrm{rk}(E|E-Z)$, by assumption, and as $\mathrm{rk}(E|E-Z) = \mathrm{rk}_{M.Y}(Y|Y-Z)$ [4, Lemma 3.13], the assertion follows. □

It might be worth noting that this proof easily extends to arbitrary finite families of co-finitary matroids.

Finally, we use Proposition 4.1 (actually only the fact that (IM) is satisfied for unions of finitary matroids), to derive a base covering result for finitary matroids. The finite base covering theorem asserts that *a finite matroid M can be covered by k bases if and only if* $\mathrm{rk}(X) \geq |X|/k$ *for every* $X \subseteq E(M)$ [7].

Corollary 5.1. *A finitary matroid M can be covered by k independent sets if and only if* $\mathrm{rk}_M(X) \geq |X|/k$ *for every finite* $X \subseteq E(M)$.

This claim is false if M is an infinite circuit, implying that this result is best possible in the sense that M being finitary is necessary.

Proof. The 'only if' implication is trivial. Suppose then that each finite set $X \subseteq E(M)$ satisfies $\mathrm{rk}_M(X) \geq |X|/k$ and put $N = \bigvee_{i=1}^{k} M$; such is a finitary matroid by Proposition 4.1. If N is the free matroid, the assertion holds trivially. Suppose then that N is not the free matroid and consequently contains a circuit C; such is finite as N is finitary. Hence, $M|C$ cannot be covered by k independent sets of $M|C$ so that by the finite matroid covering theorem [6, Theorem 12.3.12] there exists a finite set $X \subseteq C$ such that $\mathrm{rk}_{M|C}(X) < |X|/k$ which clearly implies $\mathrm{rk}_M(X) < |X|/k$; a contradiction. □

References

[1] R. Aharoni and C. Thomassen. Infinite, highly connected digraphs with no two arc-disjoint spanning trees. *J. Graph Theory* **13** (1989),71–74.

[2] E. Aigner-Horev, J. Carmesin, and J. Fröhlich. On the intersection of infinite matroids. *arXiv*:1111.0606 [math.CO] (2012).

[3] M. A. Armstrong. *Basic Topology.* Corrected reprint. Undergraduate Texts in Mathematics. Springer-Verlag, New York, 1983.

[4] H. Bruhn, R. Diestel, M. Kriesell, R. Pendavingh, and P. Wollan. Axioms for infinite matroids. *arXiv*:1003.3919 [math.CO] (2010).

[5] R. Diestel. *Graph Theory.* Fourth edition. Graduate Texts in Mathematics, 173. Springer, Heidelberg, 2010.

[6] J. Oxley. *Matroid Theory.* Oxford Science Publications. Oxford University Press, New York, 1992.

[7] A. Schrijver. *Combinatorial optimization. Polyhedra and efficiency. Vol. B. Matroids, trees, stable sets.* Chapters 39–69. Algorithms and Combinatorics, 24, B. Springer-Verlag, Berlin, 2003.

On the intersection of infinite matroids

Elad Aigner-Horev* Johannes Carmesin
Jan-Oliver Fröhlich

University of Hamburg
9 July 2012

Abstract

We show that the *infinite matroid intersection conjecture* of Nash-Williams implies the infinite Menger theorem proved recently by Aharoni and Berger.

We prove that this conjecture is true whenever one matroid is nearly finitary and the second is the dual of a nearly finitary matroid, where the nearly finitary matroids form a superclass of the finitary matroids.

In particular, this proves the infinite matroid intersection conjecture for finite-cycle matroids of 2-connected, locally finite graphs with only a finite number of vertex-disjoint rays.

1 Introduction

The infinite Menger theorem[1] was conjectured by Erdős in the 1960s and proved recently by Aharoni and Berger [1]. It states that for any two sets of vertices S and T in a connected graph, there is a set of vertex-disjoint S-T-paths whose maximality is witnessed by an S-T-separator picking exactly one vertex form each of these paths.

The complexity of the only known proof of this theorem and the fact that the finite Menger theorem has a short matroidal proof, make it natural to ask whether a matroidal proof of the infinite Menger theorem exists. In this paper, we propose a way to approach this issue by proving that a certain

*Research supported by the Minerva foundation.

[1]see Theorem 3.1 below.

conjecture of Nash-Williams regarding infinite matroids implies the infinite Menger theorem.

Recently, Bruhn, Diestel, Kriesell, Pendavingh and Wollan [5] found axioms for infinite matroids in terms of independent sets, bases, circuits, closure and (relative) rank. These axioms allow for duality of infinite matroids as known from finite matroid theory, which settled an old problem of Rado. With these new axioms it is possible now to look which theorems of finite matroid theory have infinite analogues.

Here, we shall look at the *matroid intersection theorem*, which is a classical result in finite matroid theory [9]. It asserts that *the maximum size of a common independent set of two matroids M_1 and M_2 on a common ground set E is given by*

$$\min_{X \subseteq E} \mathrm{rk}_{M_1}(X) + \mathrm{rk}_{M_2}(E \setminus X), \tag{1}$$

where rk_{M_i} denotes the rank function of the matroid M_i.

In this paper, we consider the following conjecture of Nash-Williams, which first appeared in [2] and serves as an infinite analogue to the finite matroid intersection theorem[2].

Conjecture 1.1. [The infinite matroid intersection conjecture]
Any two matroids M_1 and M_2 on a common ground set E have a common independent set I admitting a partition $I = J_1 \cup J_2$ such that $\mathrm{cl}_{M_1}(J_1) \cup \mathrm{cl}_{M_2}(J_2) = E$.

Here, $\mathrm{cl}_M(X)$ denotes the *closure* of a set X in a matroid M; it consists of X and the elements spanned by X in M (see [9]). Originally, Nash-Williams's Conjecture just concerned *finitary matroids*, those all of whose circuits are finite.

1.1 Our results

Aharoni and Ziv [2] proved that Conjecture 1.1 implies the infinite analogues of König's and Hall's theorems. We strengthen this by showing that this conjecture implies the celebrated *infinite Menger theorem* (Theorem 3.1 below), which is known to imply the infinite analogues of König's and Hall's theorems [7].

Theorem 1.2. *The infinite matroid intersection conjecture for finitary matroids implies the infinite Menger theorem.*

[2] An alternative notion of infinite matroid intersection was recently proposed by Christian [6].

In finite matroid theory, an exceptionally short proof of the matroid intersection theorem employing the well-known *finite matroid union theorem* [9, 10] is known. The latter theorem asserts that for two finite matroids $M_1 = (E_1, \mathcal{I}_1)$ and $M_2 = (E_2, \mathcal{I}_2)$ the set system

$$\mathcal{I}(M_1 \vee M_2) = \{I_1 \cup I_2 \mid I_1 \in \mathcal{I}_1,\ I_2 \in \mathcal{I}_2\} \tag{2}$$

forms the set of independent sets of their *union matroid* $M_1 \vee M_2$.

In a previous paper [3, Proposition 1.1], we showed that for infinite matroids M_1 and M_2, the set system $\mathcal{I}(M_1 \vee M_2)$ is not necessarily a matroid. This then raises the question of whether the traditional connection between (infinite) matroid union and intersection still holds. In this paper, we prove the following.

Theorem 1.3. *If M_1 and M_2 are matroids on a common ground set E and $M_1 \vee M_2^*$ is a matroid, then Conjecture 1.1 holds for M_1 and M_2.*

Throughout, M^* denotes the dual of a matroid M.

In [3] we show that the 'largest' class of matroids for which one can have a union theorem is essentially a certain superclass of the finitary matroids called the *nearly finitary* matroids (to be defined next). This, together with Theorem 1.3, enables us to make additional progress on Conjecture 1.1, as set out below.

Nearly finitary matroids are defined as follows [3]. For any matroid M, taking as circuits only the finite circuits of M defines a (finitary) matroid with the same ground set as M. This matroid is called the *finitarization* of M and denoted by M^{fin}.

It is not hard to show that every basis B of M extends to a basis B^{fin} of M^{fin}, and conversely every basis B^{fin} of M^{fin} contains a basis B of M. Whether or not $B^{\text{fin}} \setminus B$ is finite will in general depend on the choices for B and B^{fin}, but given a choice for one of the two, it will no longer depend on the choice for the second one.

We call a matroid M *nearly finitary* if every base of its finitarization contains a base of M such that their difference is finite.

Next, let us look at some examples of nearly finitary matroids. There are three natural extensions to the notion of a finite graphic matroid in an infinite context [5]; each with ground set $E(G)$. The most studied one is the *finite-cycle matroid*, denoted $M_{FC}(G)$, whose circuits are the finite cycles of G. This is a finitary matroid, and hence is also nearly finitary.

The second extension is the *algebraic-cycle matroid*, denoted $M_A(G)$, whose circuits are the finite cycles and double rays of G [5, 4][3].

[3] $M_A(G)$ is not necessarily a matroid for any G; see [8].

Proposition 1.4. *$M_A(G)$ is a nearly finitary matroid if and only if G has only a finite number of vertex-disjoint rays.*

The third extension is the *topological-cycle matroid*, denoted $M_C(G)$[4], whose circuits are the topological cycles of G (Thus $M_C^{\text{fin}}(G) = M_{FC}(G)$ for any finitely separable graph G; see Section 5.2 or [4] for definitions).

Proposition 1.5. *Suppose that G is 2-connected and locally finite. Then, $M_C(G)$ is a nearly finitary matroid if and only if G has only a finite number of vertex-disjoint rays.*

Having introduced nearly finitary matroids, we now state the result of [3].

Theorem 1.6. [Nearly finitary union theorem [3]]
If M_1 and M_2 are nearly finitary matroids, then $M_1 \vee M_2$ is a nearly finitary matroid.

The following is a consequence of Theorem 1.6 and Theorem 1.3.

Corollary 1.7. *Conjecture 1.1 holds for M_1 and M_2 whenever M_1 is nearly finitary and M_2 is the dual of a nearly finitary matroid.*

Aharoni and Ziv [2] proved that the infinite matroid intersection conjecture is true whenever one matroid is finitary and the other is a countable direct sum of finite-rank matroids. Note that Corollary 1.7 does not imply this result of [2] nor is it implied by it.

Proposition 1.5 and Corollary 1.7 can be used to prove the following.

Corollary 1.8. *Suppose that G and H are 2-connected, locally finite graphs with only a finite number of vertex-disjoint rays. Then their finite-cycle matroids $M_{FC}(G)$ and $M_{FC}(H)$ satisfy the intersection conjecture.*

Similar results are true for *the algebraic-cycle matroid*, *the topological-cycle matroid*, and their duals.

This paper is organized as follows. Additional notation, terminology, and basic lemmas are given in Section 2. In Section 3 we prove Theorem 1.2. In Section 4 we prove Theorem 1.3, and in Section 5 we prove Propositions 1.4 and 1.5 and Corollary 1.8.

[4] $M_C(G)$ is a matroid for any G; see [4].

2 Preliminaries

Notation and terminology for graphs are that of [7], and for matroids that of [9, 5].

Throughout, G always denotes a graph where $V(G)$ and $E(G)$ denote its vertex and edge sets, respectively. We write M to denote a matroid and write $E(M)$, $\mathcal{I}(M)$, $\mathcal{B}(M)$, and $\mathcal{C}(M)$ to denote its ground set, independent sets, bases, and circuits, respectively.

We review the definition of a matroid as this is given in [5]. A set system $\mathcal{I}$ is the set of independent sets of a matroid if it satisfies the following *independence axioms*:

(I1) $\emptyset \in \mathcal{I}$.

(I2) $\lceil \mathcal{I} \rceil = \mathcal{I}$, that is, $\mathcal{I}$ is closed under taking subsets.

(I3) Whenever $I, I' \in \mathcal{I}$ with I' maximal and I not maximal, there exists an $x \in I' \setminus I$ such that $I + x \in \mathcal{I}$.

(IM) Whenever $I \subseteq X \subseteq E$ and $I \in \mathcal{I}$, the set $\{I' \in \mathcal{I} \mid I \subseteq I' \subseteq X\}$ has a maximal element.

The following is a well-known fact for finite matroids (see, e.g., [9]), which generalizes easily to infinite matroids.

Lemma 2.1. [5, Lemma 3.11]
Let M be a matroid. Then, $|C \cap C^*| \neq 1$, whenever $C \in \mathcal{C}(M)$ and $C^* \in \mathcal{C}(M^*)$.

We end this section with the definition of exchange chains. For a set $X \subseteq E(M)$, an *X-circuit* is a circuit containing X. For sets $I_1 \in \mathcal{I}(M_1)$ and $I_2 \in \mathcal{I}(M_2)$, and elements $x \in I_1 \cup I_2$ and $y \in E(M_1) \cup E(M_2)$ (possibly in $I_1 \cup I_2$), a tuple $Y = (y_0 = y, \ldots, y_n = x)$ is called an *even* (I_1, I_2, y, x)-*exchange chain* (or *even* (I_1, I_2, y, x)-*chain*) of *length* n if the following terms are satisfied.

(X1) For an even i, there exists a $\{y_i, y_{i+1}\}$-circuit $C_i \subseteq I_1 + y_i$ of M_1.

(X2) For an odd i, there exists a $\{y_i, y_{i+1}\}$-circuit $C_i \subseteq I_2 + y_i$ of M_2.

If $n \geq 1$, then (X1) and (X2) imply that $y_0 \notin I_1$ and that, starting with $y_1 \in I_1 \setminus I_2$, the elements y_i alternate between $I_1 \setminus I_2$ and $I_2 \setminus I_1$; the single exception being y_n which might lie in $I_1 \cap I_2$.

By an *odd exchange chain* (or *odd chain*) we mean an even chain with the words 'even' and 'odd' interchanged in the definition. Consequently, we say *exchange chain* (or *chain*) to refer to either of these notions. Furthermore, a subchain of a chain is also a chain; that is, given an (I_1, I_2, y_0, y_n)-chain $(y_0, \ldots, y_n)$, the tuple $(y_k, \ldots, y_l)$ is an (I_1, I_2, y_k, y_l)-chain for $0 \leq k \leq l \leq n$.

Lemma 2.2. [3, Lemma 4.4]
If there exists an (I_1, I_2, y, x)-chain, then $(I + y) - x \in \mathcal{I}(M_1 \vee M_2)$ where $I := I_1 \cup I_2$. Moreover, if $x \in I_1 \cap I_2$, then $I + y \in \mathcal{I}(M_1 \vee M_2)$.

3 From infinite matroid intersection to the infinite Menger theorem

In this section, we prove Theorem 1.2; asserting that the infinite matroid intersection conjecture implies the infinite Menger theorem.

Given a graph G and $S, T \subseteq V(G)$, a set $X \subseteq V(G)$ is called an *S–T separator* if $G - X$ contains no S–T path. The infinite Menger theorem reads as follows.

Theorem 3.1 (Aharoni and Berger [1])**.** *Let G be a connected graph. Then for any $S, T \subseteq V(G)$ there is a set $\mathcal{L}$ of vertex-disjoint S–T paths and an S–T separator $X \subseteq \bigcup_{P \in \mathcal{L}} V(P)$ satisfying $|X \cap V(P)| = 1$ for each $P \in \mathcal{L}$.*

Infinite matroid union cannot be used in order to obtain the infinite Menger Theorem directly via Theorem 1.3 and Theorem 1.2. Indeed, in [3, Proposition 1.1] we construct a finitary matroid M and a co-finitary matroid N such that their union is not a matroid. Consequently, one cannot apply Theorem 1.3 to the finitary matroids M and N^* in order to obtain Conjecture 1.1 for them. However, it is easy to see that Conjecture 1.1 is true for these particular M and N^*.

Next, we prove Theorem 1.2.

Proof of Theorem 1.2. Let G be a connected graph and let $S, T \subseteq V(G)$ be as in Theorem 3.1. We may assume that $G[S]$ and $G[T]$ are both connected. Indeed, an S–T separator with $G[S]$ and $G[T]$ connected gives rise to an S–T separator when these are not necessarily connected. Abbreviate $E(S) := E(G[S])$ and $E(T) := E(G[T])$, let M be the finite-cycle matroid $M_F(G)$, and put $M_S := M/E(S) - E(T)$ and $M_T := M/E(T) - E(S)$; all three matroids are clearly finitary.

Assuming that the infinite matroid intersection conjecture holds for M_S and M_T, there exists a set $I \in \mathcal{I}(M_S) \cap \mathcal{I}(M_T)$ which admits a partition $I = J_S \cup J_T$ satisfying

$$\mathrm{cl}_{M_S}(J_S) \cup \mathrm{cl}_{M_T}(J_T) = E,$$

where $E = E(M_S) = E(M_T)$. We regard I as a subset of $E(G)$.

For the components of $G[I]$ we observe two useful properties. As I is disjoint from $E(S)$ and $E(T)$, the edges of a cycle in a component of $G[I]$ form a circuit in both, M_S and M_T, contradicting the independence of I in either. Consequently,

$$\text{the components of } G[I] \text{ are trees.} \tag{3}$$

Next, an S-path[5] or a T-path in a component of $G[I]$ gives rise to a circuit of M_S or M_T in I, respectively. Hence,

$$|V(C) \cap S| \leq 1 \text{ and } |V(C) \cap T| \leq 1 \text{ for each component } C \text{ of } G[I]. \tag{4}$$

Let $\mathcal{C}$ denote the components of $G[I]$ meeting both of S and T. Then by (3) and (4) each member of $\mathcal{C}$ contains a unique S–T path and we denote the set of all these paths by $\mathcal{L}$. Clearly, the paths in $\mathcal{L}$ are vertex-disjoint.

In what follows, we find a set X comprised of one vertex from each $P \in \mathcal{L}$ to serve as the required S–T separator. To that end, we show that one may alter the partition $I = J_S \cup J_T$ to yield a partition

$$I = K_S \cup K_T \text{ satisfying } cl_{M_S}(K_S) \cup cl_{M_T}(K_T) = E \text{ and (Y.1-4)}, \tag{5}$$

where (Y.1-4) are as follows.

(Y.1) Each component C of $G[I]$ contains a vertex of $S \cup T$.

(Y.2) Each component C of $G[I]$ meeting S but not T satisfies $E(C) \subseteq K_S$.

(Y.3) Each component C of $G[I]$ meeting T but not S satisfies $E(C) \subseteq K_T$.

(Y.4) Each component C of $G[I]$ meeting both, S and T, contains at most one vertex which at the same time

- (a) lies in S or is incident with an edge of K_S, and
- (b) lies in T or is incident with an edge of K_T.

[5]A non-trivial path meeting $G[S]$ exactly in its end vertices.

Postponing the proof of (5), we first show how to deduce the existence of the required S–T separator from (5). Define a pair of sets of vertices (V_S, V_T) of $V(G)$ by letting V_S consist of those vertices contained in S or incident with an edge of K_S and defining V_T in a similar manner. Then $V_S \cap V_T$ may serve as the required S–T separator. To see this, we verify below that (V_S, V_T) satisfies all of the terms (Z.1-4) stated next.

(Z.1) $V_S \cup V_T = V(G)$;

(Z.2) for every edge e of G either $e \subseteq V_S$ or $e \subseteq V_T$;

(Z.3) every vertex in $V_S \cap V_T$ lies on a path from $\mathcal{L}$; and

(Z.4) every member of $\mathcal{L}$ meets $V_S \cap V_T$ at most once.

To see (Z.1), suppose v is a vertex not in $S \cup T$. As G is connected, such a vertex is incident with some edge $e \notin E(T) \cup E(S)$. The edge e is spanned by K_T or K_S; say K_T. Thus, $K_T + e$ contains a circle containing e or $G[K_T + e]$ has a T-path containing e. In either case v is incident with an edge of K_T and thus in V_T, as desired.

To see (Z.2), let $e \in cl_{M_T}(K_T) \setminus K_T$; so that $K_T + e$ has a circle containing e or $G[K_T + e]$ has T-path containing e; in either case both end vertices of e are in V_T, as desired. The treatment of the case $e \in cl_{M_S}(K_S)$ is similar.

To see (Z.3), let $v \in V_S \cap V_T$; such is in S or is incident with an edge of K_S, and in T or is incident with an edge in K_T. Let C be the component of $G[I]$ containing v. By (Y.1-4), $C \in \mathcal{C}$, i.e. it meets both, S and T and therefore contains an S–T path $P \in \mathcal{L}$. Recall that every edge of C is either in K_S or K_T and consider the last vertex w of a maximal initial segment of P in $C - K_T$. Then w satisfies (Y.4a), as well as (Y.4b), implying $v = w$; so that v lies on a path from $\mathcal{L}$.

To see (Z.4), we restate (Y.4) in terms of V_S and V_T: each component of $\mathcal{C}$ contains at most one vertex of $V_S \cap V_T$. This clearly also holds for the path from $\mathcal{L}$ which is contained in C.

It remains to prove (5). To this end, we show that any component C of $G[I]$ contains a vertex of $S \cup T$. Suppose not. Let e be the first edge on a $V(C)$–S path Q which exists by the connectedness of G. Then $e \notin I$ but without loss of generality we may assume that $e \in \mathrm{cl}_{M_S}(J_S)$. So in $G[I] + e$ there must be a cycle or an S-path. The latter implies that C contains a vertex of S and the former means that Q was not internally disjoint to $V(C)$, yielding contradictions in both cases.

We define the sets K_S and K_T as follows. Let C be a component of $G[I]$.

1. If C meets S but not T, then include its edges into K_S.

2. If C meets T but not S, then include its edges into K_T.

3. Otherwise (C meets both of S and T) there is a path P from $\mathcal{L}$ in C. Denote by v_C the last vertex of a maximal initial segment of P in $C - J_T$. As C is a tree, each component C' of $C - v_C$ is a tree and there is a unique edge e between v_C and C'. For every such component C', include the edges of $C' + e$ in K_S if $e \in J_S$ and in K_T otherwise, i.e. if $e \in J_T$.

Note that, by choice of v_C, either v_C is the last vertex of P or the next edge of P belongs to J_T. This ensures that K_S and K_T satisfy (Y.4). Moreover, they form a partition of I which satisfies (Y.1-3) by construction. It remains to show that $\mathrm{cl}_{M_S}(K_S) \cup \mathrm{cl}_{M_T}(K_T) = E$.

As $K_S \cup K_T = I$, it suffices to show that any $e \in E \setminus I$ is spanned by K_S in M_S or by K_T in M_T. Suppose $e \in \mathrm{cl}_{M_S}(J_S)$, i.e. $J_S + e$ contains a circuit of M_S. Hence, $G[J_S]$ either contains an e-path R or two disjoint e–S paths R_1 and R_2. We show that $E(R) \subseteq K_S$ or $E(R) \subseteq K_T$ in the former case and $E(R_1) \cup E(R_2) \subseteq K_S$ in the latter.

The path R is contained in some component C of $G[I]$. Suppose $C \in \mathcal{C}$ and v_C is an inner vertex of R. By assumption, the edges preceding and succeeding v_C on R are both in J_S and hence the edges of both components of $C - v_C$ which are met by R plus their edges to v_C got included into K_S, showing $E(R) \subseteq K_S$. Otherwise $C \notin \mathcal{C}$ or $C \in \mathcal{C}$ but v_C is no inner vertex of R. In both cases the whole set $E(R)$ got included into K_S or K_T.

We apply the same argument to R_1 and R_2 except for one difference. If $C \notin \mathcal{C}$ or $C \in \mathcal{C}$ but v_C is no inner vertex of R_i, then $E(R_i)$ got included into K_S as R_i meets S.

Although the definitions of K_S and K_T are not symmetrical, a similar argument shows $e \in \mathrm{cl}_{M_S}(K_S) \cup \mathrm{cl}_{M_T}(K_T)$ if e is spanned by J_T in M_T. $\square$

Note that the above proof requires only that Conjecture 1.1 holds for finite-cycle matroids.

4 From infinite matroid union to infinite matroid intersection

In this section, we prove Theorem 1.3.

Proof of Theorem 1.3. Our starting point is the well-known proof from finite matroid theory that matroid union implies a solution to the matroid intersection problem. With that said, let $B_1 \cup B_2^* \in \mathcal{B}(M_1 \vee M_2^*)$ where $B_1 \in \mathcal{B}(M_1)$

and $B_2^* \in \mathcal{B}(M_2^*)$, and let $B_2 = E \setminus B_2^* \in \mathcal{B}(M_2)$. Then, put $I = B_1 \cap B_2$ and note that $I \in \mathcal{I}(M_1) \cap \mathcal{I}(M_2)$. We show that I admits the required partition.

For an element $x \notin B_i$, $i = 1, 2$, we write $C_i(x)$ to denote the fundamental circuit of x into B_i in M_i. For an element $x \notin B_2^*$, we write $C_2^*(x)$ to denote the fundamental circuit of x into B_2^* in M_2^*. Put $X = B_1 \cap B_2^*$, $Y = B_2 \setminus I$, and $Z = B_2^* \setminus X$, see Figure 1.

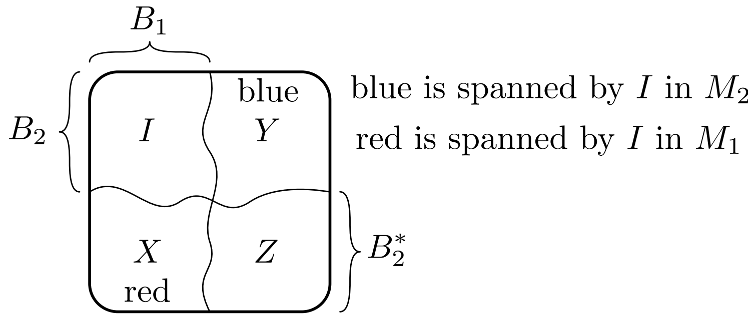

Figure 1: The sets X, Y, and Z and their colorings.

Observe that

$$cl_{M_1}(I) \cup cl_{M_2}(I) = E = I \cup X \cup Y \cup Z. \tag{6}$$

To see (6), note first that

$$X \subseteq cl_{M_2}(I). \tag{7}$$

Clearly, no member of X is spanned by I in M_1. Assume then that $x \in X$ is not spanned by I in M_2 so that there exists a $y \in C_2(x) \cap Y$. Then, $x \in C_2^*(y)$, by Lemma 2.1. Consequently, $B_1 \cup B_2^* \subsetneq B_1 \cup (B_2^* + y - x) \in \mathcal{I}(M_1 \vee M_2^*)$; contradiction to the maximality of $B_1 \cup B_2^*$, implying (7).

By a similar argument, it holds that

$$Y \subseteq cl_{M_1}(I). \tag{8}$$

To see that

$$Z \subseteq cl_{M_1}(I) \cup cl_{M_2}(I), \tag{9}$$

assume, towards contradiction, that some $z \in Z$ is not spanned by I neither in M_1 nor in M_2 so that there exist an $x \in C_1(z) \cap X$ and a $y \in C_2(z) \cap Y$. Then $B_1 - x + z$ and $B_2 - y + z$ are bases and thus $B_1 \cup B_2^* \subsetneq (B_1 - x + z) \cup (B_2^* - z + y)$; contradiction to the maximality of $B_1 \cup B_2^*$. Assertion (6) is proved.

The problem of finding a suitable partition $I = J_1 \cup J_2$ can be phrased as a (directed) graph coloring problem. By (6), each $x \in E \setminus I$ satisfies

$C_1(x) - x \subseteq I$ or $C_2(x) - x \subseteq I$. Define $G = (V, E)$ to be the directed graph whose vertex set is $V = E \setminus I$ and whose edge set is given by

$$E = \{(x, y) : C_1(x) \cap C_2(y) \cap I \neq \emptyset\}. \tag{10}$$

Recall that a *source* is a vertex with no incoming edges and a *sink* is a vertex with no outgoing edges. As $C_1(x)$ does not exist for any $x \in X$ and $C_2(y)$ does not exist for any $y \in Y$, it follows that

$$\text{the members of } X \text{ are sinks and those of } Y \text{ are sources in } G. \tag{11}$$

A 2-coloring of the vertices of G, by say blue and red, is called *divisive* if it satisfies the following:

(D.1) I spans all the blue elements in M_1;

(D.2) I spans all the red elements in M_2; and

(D.3) $J_1 \cap J_2 = \emptyset$ where $J_1 := (\bigcup_{x \text{ blue}} C_1(x)) \cap I$ and $J_2 := (\bigcup_{x \text{ red}} C_2(x)) \cap I$.

Clearly, if G has a divisive coloring, then I admits the required partition.

We show then that G admits a divisive coloring. Color with blue all the sources. These are the vertices that can only be spanned by I in M_1. Color with red all the sinks, that is, all the vertices that can only be spanned by I in M_2. This defines a partial coloring of G in which all members of X are red and those of Y are blue. Such a partial coloring can clearly be extended into a divisive coloring of G provided that

$$G \text{ has no } (y, x)\text{-path with } y \text{ blue and } x \text{ red.} \tag{12}$$

Indeed, given (12) and (11), color all vertices reachable by a path from a blue vertex with the color blue, color all vertices from which a red vertex is reachable by a path with red, and color all remaining vertices with, say, blue. The resulting coloring is divisive.

It remains to prove (12). We show that the existence of a path as in (12) contradicts the following property:
Suppose that M and N are matroids and $B \cup B'$ is maximal in $\mathcal{I}(M \vee N)$. Let $y \notin B \cup B'$ and let $x \in B \cap B'$. Then, (by Lemma 2.2)

$$\textit{there exists no } (B, B', y, x)\textit{-chain;} \tag{13}$$

(in fact, the contradiction in the proofs of (7),(8), and (9) arose from simple instances of such forbidden chains).

Assume, towards contradiction, that P is a (y, x)-path with y blue and x red; the intermediate vertices of such a path are not colored since they are not a sink nor a source. In what follows we use P to construct a $(B_1, B_2^*, y_0, y_{2|P|})$-chain $(y_0, y_1, \ldots, y_{2|P|})$ such that $y_0 \in Y$, $y_{2|P|} \in X$, all odd indexed members of the chain are in $V(P) \cap Z$, and all even indexed elements of the chain other than y_0 and $y_{2|P|}$ are in I. Existence of such a chain would contradict (13).

Definition of y_0. As y is pre-colored blue then either $y \in Y$ or $C_2(y) \cap Y \neq \emptyset$. In the former case set $y_0 = y$ and in the latter choose $y_0 \in C_2(y) \cap Y$.

Definition of $y_{2|P|}$. In a similar manner, x is pre-colored red since either $x \in X$ or $C_1(x) \cap X \neq \emptyset$. In the former case, set $y_{2|P|} = x$ and in the latter case choose $y_{2|P|} \in C_1(x) \cap X$.

The remainder of the chain. Enumerate $V(P) \cap Z = \{y_1, y_3, \ldots, y_{2|P|-1}\}$ where the enumeration is with respect to the order of the vertices defined by P. Next, for an edge $(y_{2i-1}, y_{2i+1}) \in E(P)$, let $y_{2i} \in C_1(y_{2i-1}) \cap C_2(y_{2i+1}) \cap I$; such exists by the assumption that $(y_{2i-1}, y_{2i+1}) \in E$. As $y_{2i+1} \in C_2^*(y_{2i})$ for all relevant i, by Lemma 2.1, the sequence $(y_0, y_1, y_2, \ldots, y_{2|P|})$ is a $(B_1, B_2^*, y_0, y_{2|P|})$-chain in $\mathcal{I}(M_1 \vee M_2^*)$.

This completes our proof of Theorem 1.3. $\square$

Note that in the above proof, we do not use the assumption that $M_1 \vee M_2^*$ is a matroid; in fact, we only need that $\mathcal{I}(M_1 \vee M_2^*)$ has a maximal element.

5 The graphic nearly finitary matroids

In this section we prove Propositions 1.4 and 1.5 yielding a characterization of the graphic nearly finitary matroids.

For a connected graph G, a maximal set of edges containing no finite cycles is called an *ordinary spanning tree.* A maximal set of edges containing no finite cycles nor any double ray is called an *algebraic spanning tree.* These are the bases of $M_F(G)$ and $M_A(G)$, respectively. We postpone the discussion about $M_C(G)$ to Section 5.2.

To prove Propositions 1.4 and 1.5, we require the following theorem of Halin [7, Theorem 8.2.5].

Theorem 5.1 (Halin 1965)**.** *If an infinite graph G contains k disjoint rays for every $k \in \mathbb{N}$, then G contains infinitely many disjoint rays.*

5.1 The nearly finitary algebraic-cycle matroids

The purpose of this subsection is to prove Proposition 1.4.

Proof of Proposition 1.4. Suppose that G has k disjoint rays for every integer k; so that G has a set $\mathcal{R}$ of infinitely many disjoint rays by Theorem 5.1. We show that $M_A(G)$ is not nearly finitary.

The edge set of $\bigcup \mathcal{R} = \bigcup_{R \in \mathcal{R}} R$ is independent in $M_A(G)^{\text{fin}}$ as it induces no finite cycle of G. Therefore there is a base of $M_A(G)^{\text{fin}}$ containing it; such induces an ordinary spanning tree, say T, of G. We show that

$$T - F \text{ contains a double ray for any finite edge set } F \subseteq E(T). \tag{14}$$

This implies that $E(T) \setminus I$ is infinite for every independent set I of $M_A(G)$ and hence $M_A(G)$ is not nearly finitary. To see (14), note that $T - F$ has $|F| + 1$ components for any finite edge set $F \subseteq E(T)$ as T is a tree and successively removing edges always splits one component into two. So one of these components contains infinitely many disjoint rays from $\mathcal{R}$ and consequently a double ray.

Suppose next, that G has at most k disjoint rays for some integer k and let T be an ordinary spanning tree of G, that is, $E(T)$ is maximal in $M_A(G)^{\text{fin}}$. To prove that $M_A(G)$ is nearly finitary, we need to find a finite set $F \subseteq E(T)$ such that $E(T) \setminus F$ is independent in $M_A(G)$, i.e. it induces no double ray of G. Let $\mathcal{R}$ be a maximal set of disjoint rays in T; such exists by assumption and $|\mathcal{R}| \leq k$. As T is a tree and the rays of $\mathcal{R}$ are vertex-disjoint, it is easy to see that T contains a set F of $|\mathcal{R}| - 1$ edges such that $T - F$ has $|\mathcal{R}|$ components which each contain one ray of $\mathcal{R}$. By maximality of $\mathcal{R}$ no component of $T - F$ contains two disjoint rays, or equivalently, a double ray. $\square$

5.2 The nearly finitary topological-cycle matroids

In this section we prove Proposition 1.5 that characterizes the nearly finitary topological-cycle matroids. Prior to that, we first define these matroids. To that end we shall require some additional notation and terminology on which more details can be found in [4].

An *end of* G is an equivalence class of rays, where two rays are *equivalent* if they cannot be separated by a finite edge set. In particular, two rays meeting infinitely often are equivalent. Let the *degree* of an end ω be the size of a maximal set of vertex-disjoint rays belonging to ω, which is well-defined [7]. We say that a double ray *belongs to* an end if the two rays which arise from the removal of one edge from the double ray belong to that end; this does not depend on the choice of the edge. Such a double ray is an example of a *topological cycle*[6]

For a graph G the topological-cycle matroid of G, namely $M_C(G)$, has $E(G)$ as its ground set and its set of circuits consists of the finite and topological cycles. In fact, every infinite circuit of $M_C(G)$ induces at least one double ray; provided that G is locally finite [7].

A graph G has only finitely many disjoint rays if and only if G has only finitely many ends, each with finite degree. Also, note that

$$\text{every end of a 2-connected locally finite graph has degree at least 2.} \tag{15}$$

Indeed, applying Menger's theorem inductively, it is easy to construct in any k-connected graph for any end ω a set of k disjoint rays of ω.

Now we are in a position to start the proof of Proposition 1.5.

Proof of Proposition 1.5. If G has only a finite number of vertex-disjoint rays then $M_A(G)$ is nearly finitary by Proposition 1.4. Since $M_A(G)^{\text{fin}} = M_C(G)^{\text{fin}}$ and $\mathcal{I}(M_A(G)) \subseteq \mathcal{I}(M_C(G))$, we can conclude that $M_C(G)$ is nearly finitary as well.

Now, suppose that G contains k vertex-disjoint rays for every $k \in \mathbb{N}$. If G has an end ω of infinite degree, then there is an infinite set $\mathcal{R}$ of vertex-disjoint rays belonging to ω. As any double ray containing two rays of $\mathcal{R}$ forms a circuit of $M_C(G)$, the argument from the proof of Proposition 1.4 shows that $M_C(G)$ is not nearly finitary.

Assume then that G has no end of infinite degree. There are infinitely many disjoint rays, by Theorem 5.1. Hence, there is a countable set of ends $\Omega = \{\omega_1, \omega_2, \ldots\}$.

We inductively construct a set $\mathcal{R}$ of infinitely many vertex-disjoint double rays, one belonging to each end of Ω. Suppose that for any integer $n \geq 0$ we have constructed a set $\mathcal{R}_n$ of n disjoint double rays, one belonging to each of the ends $\omega_1, \ldots, \omega_n$. Different ends can be separated by finitely many vertices so there is a finite set S of vertices such that $\bigcup \mathcal{R}_n$ has no vertex in the component C of $G - S$ which contains ω_{n+1}. Since ω_{n+1} has degree 2 by (15), there are two disjoint rays from ω_{n+1} in C an thus also a double ray D belonging to ω_{n+1}. Set $\mathcal{R}_{n+1} := \mathcal{R}_n \cup \{D\}$ and $\mathcal{R} := \bigcup_{n\in\mathbb{N}} \mathcal{R}_n$.

As $\bigcup \mathcal{R}$ contains no finite cycle of G, it can be extended to an ordinary spanning tree of G. Removing finitely many edges from this tree clearly leaves an element of $\mathcal{R}$ intact. Hence, the edge set of the resulting graph still contains a circuit of $M_C(G)$. Thus, $M_C(G)$ is not nearly finitary in this case as well. $\square$

[6]Formally, the topological cycles of G are those subgraphs of G which are homeomorphic images of S^1 in the Freudenthal compactification $|G|$ of G. However, the given example is the only type of topological cycle which shall be needed for the proof.

In the following we shall propose a possible common generalization of Propositions 1.4 and 1.5 to all infinite matroids. We call a matroid M *k-nearly finitary* if every base of its finitarization contains a base of M such that their difference has size at most k. Note that saying 'at most k' is not equivalent to saying 'equal to k', consider for example the algebraic-cycle matroid of the infinite ladder. In terms of this new definition, Propositions 1.4 and 1.5 both state for a certain class of infinite matroids that each member of this class is k-nearly finitary for some k. In fact, for all known examples of nearly finitary matroids, there is such a k. This raises the following open question.

Question 5.2. *Is every nearly finitary matroid k-nearly finitary for some k?*

5.3 Graphic matroids and the intersection conjecture

By Corollary 1.7, the intersection conjecture is true for $M_C(G)$ and $M_{FC}(H)$ for any two graphs G and H since the first is co-finitary and the second is finitary. Using also Proposition 1.5, we obtain the following.

Corollary 5.3. *Suppose that G and H are 2-connected, locally finite graphs with only a finite number of vertex-disjoint rays. Then, $M_C(G)$ and $M_C(H)$ satisfy the intersection conjecture.* □

Using Proposition 1.4 instead of Proposition 1.5, we obtain the following.

Corollary 5.4. *Suppose that G and H are graphs with only a finite number of vertex-disjoint rays. Then, $M_A(G)$ and $M_A(H)$ satisfy the intersection conjecture if both are matroids.* □

With a little more work, the same is also true for $M_{FC}(G)$, see Corollary 1.8.

Proof of Corollary 1.8. First we show that $(((M_C(G)^{\text{fin}})^*)^{\text{fin}})^* = M_C(G)$ if G is locally finite. Indeed, then $M_C(G)^{\text{fin}} = M_{FC}(G)$, $M_{FC}(G)^*$ is the matroids whose circuits are the finite and infinite bonds of G, and its finitarization has as its circuits the finite bonds of G. And the dual of this matroid is $M_C(G)$, see [5] for example.

Having showed that $(((M_C(G)^{\text{fin}})^*)^{\text{fin}})^* = M_C(G)$ if G is locally finite, we next show that if $M_C(G)$ is nearly finitary, then so is $M_{FC}(G)^*$. For this let B be a base of $M_{FC}(G)^*$ and B' be a base of $(M_{FC}(G)^*)^{\text{fin}}$. Then $B' \setminus B = (E \setminus B) \setminus (E \setminus B')$. Now $E \setminus B$ is a base of $M_{FC}(G) = M_C(G)^{\text{fin}}$ and by the above $E \setminus B'$ is a base of $M_C(G)$. Since $M_C(G)$ is nearly finitary, $B' \setminus B$ is finite, yielding that $M_{FC}(G)^*$ is nearly finitary.

As $M_{FC}(G)^*$ is nearly finitary and $M_{FC}(H)$ is finitary, $M_{FC}(H)$ and $M_{FC}(G)$ satisfy the intersection conjecture by Corollary 1.7. $\square$

A similar argument shows that if G and H are are 2-connected, locally finite graphs with only a finite number of vertex-disjoint rays, then one can also prove that $M_{FC}(G)^*$ and $M_{FC}(H)^*$ satisfy the intersection conjecture. Similar results are true for $M_C(G)^*$ or $M_A(G)^*$ in place of $M_{FC}(G)^*$.

References

[1] R. Aharoni and E. Berger. Menger's theorem for infinite graphs. *Invent. Math.* **176** (2009), 1–62.

[2] R. Aharoni and R. Ziv. The intersection of two infinite matroids. *J. London Math. Soc. (2)* **58** (1998), 513–525.

[3] E. Aigner-Horev, J. Carmesin, and J. Fröhlich. Infinite matroid union. *arXiv*:1111.0602 [math.CO] (2012).

[4] H. Bruhn and R. Diestel. Infinite matroids in graphs. *arXiv*:1011.4749 [math.CO] (2010).

[5] H. Bruhn, R. Diestel, M. Kriesell, R. Pendavingh, and P. Wollan. Axioms for infinite matroids. *arXiv*:1003.3919 [math.CO] (2010).

[6] R. Christian. Infinite graphs, graph-like spaces and B-matroids. *PhD thesis*, University of Waterloo, Waterloo, 2010.

[7] R. Diestel. *Graph Theory.* Fourth edition. Graduate Texts in Mathematics, 173. Springer, Heidelberg, 2010.

[8] D. A. Higgs. Infinite graphs and matroids. *Recent Progress in Combinatorics; Proceedings of the Third Waterloo Conference on Combinatorics, May 1968* pp. 245–253, Academic Press, New York, 1969.

[9] J. Oxley. *Matroid Theory.* Oxford Science Publications. Oxford University Press, New York, 1992.

[10] A. Schrijver. *Combinatorial optimization. Polyhedra and efficiency. Vol. B. Matroids, trees, stable sets.* Chapters 39–69. Algorithms and Combinatorics, 24, B. Springer-Verlag, Berlin, 2003.

Summary

This cumulative dissertation consists of the three papers, "Linkages in Large Graphs of Bounded Tree-Width" (chapter 1), "Infinite matroid union" (chapter 2), and "On the intersection of infinite matroids" (chapter 3), that all belong to the area of discrete mathematics and involve very large or infinite structures.

The smallest known function f such that every $f(k)$-connected graph is k-linked is $f(k) = 10k$ as shown by Thomas and Wollan in 2005. In "Linkages in Large Graphs of Bounded Tree-Width" we show that for all positive integers k and w there is an integer N such that every $(2k+3)$-connected graph G of tree-width less than w on at least N vertices is k-linked.

The papers "Infinite matroid union" and "On the intersection of infinite matroids" extend the notion of a union of two finite matroids to infinite matroids and apply the derived union theorem.

In the former paper, we show that the union of two infinite matroids need not be a matroid. We establish a union theorem for a superclass of the finitary matroids (matroids with no infinite cycle), called the *nearly finitary* matroids and show that for every matroid that is not nearly finitary and satisfies a certain countability condition there is a finitary matroid such that the union of the two is not a matroid. We use the union theorem for finitary matroids to obtain base covering and base packing results for finitary and co-finitary matroids, respectively.

In the latter paper, we show that Nash-Williams' infinite matroid intersection conjecture implies the infinite Menger theorem. We also establish a link between our union theorem and Nash-Williams' conjecture by the use of *exchange chains*, a technique that we also used in the former paper to show that any union of two matroids satisfies the independence axiom (I3). Finally, we explore the implications of the matroidal framework developed in both papers for cycle matroids of graphs.

Zusammenfassung

Diese kumulative Dissertation besteht aus drei Arbeiten, „Linkages in Large Graphs of Bounded Tree-Width“ (Kapitel 1), „Infinite matroid union“ (Kapitel 2) und „On the intersection of infinite matroids“ (Kapitel 3), die alle zum Gebiet der diskreten Mathematik zu zählen sind und sich mit sehr großen oder unendlichen Strukturen beschäftigen.

Die kleinste bekannte Funktion f, so dass jeder $f(k)$-zusammenhängende Graph k-verbunden ist, ist $f(k) = 10k$ und wurde 2005 von Thomas und Wollan gefunden. In „Linkages in Large Graphs of Bounded Tree-Width“ zeigen wir, dass es für alle natürlichen Zahlen k und w eine natürliche Zahl N gibt, so dass jeder $(2k + 3)$-zusammenhängende Graph G der Baumweite kleiner w auf mindestens N Ecken k-verbunden ist.

Die Arbeiten „Infinite matroid union“ und „On the intersection of infinite matroids“ erweitern den Begriff der Vereinigung zweier endlicher Matroide auf unendliche Matroide und wenden den dafür hergeleiteten Vereinigungssatz an.

In der ersten Matroid-Arbeit zeigen wir, dass die Vereinigung zweier unendlicher Matroide kein Matroid sein muss. Wir beweisen einen Vereinigungssatz für eine Oberklasse der finitären Matroide (solche ohne unendlichen Kreis), die der *fast finitären* Matroide, und konstruieren für jedes Matriod, das nicht fast finitär ist und einer gewissen Abzählbarkeitsbedingung genügt, ein finitäres Matroid, so dass die Vereinigung der beiden kein Matroid ist. Mit dem Vereinigungssatz für finitäre Matroide zeigen wir einen Basisüberdeckungs-Satz für finitäre und einen Basispackungs-Satz für co-finitäre Matroide.

In der zweiten Matroid-Arbeit leiten wir den unendlichen Satz von Menger aus Nash-Williams' Schnitt-Vermutung für unendliche Matroide her. Wir finden außerdem eine Verbindung zwischen unserem Vereinigungssatz und Nash-Williams' Vermutung durch die Anwendung von *Austauschketten*, einer Technik, die wir in unserer ersten Matroid-Arbeit eingeführt haben, um nachzuweisen, dass die Vereinigung zweier Matroide stets dem Unabhängigkeitsaxiom (I3) genügt. Schlussendlich untersuchen wir die Folgen der in beiden Arbeiten entwickelten Resultate für Kreismatroide von Graphen.

Die Entwicklung der Linkage-Arbeit

Die Entwicklung der Arbeit „Linkages in Large Graphs of Bounded Tree-Width“ gliedert sich zeitlich in drei Abschnitte.

1. Das Projekt wird zunächst von Ken-ichi Kawarabayashi, Theodor Müller und Paul Wollan in Tokio im September 2009 bearbeitet. Hier wird die grundlegende Beweisstrategie entwickelt. Das allgemeine Linkage-Problem soll auf ein Linkage-Problem in einem Graphen mit einer langen Wegzerlegung zurückgeführt werden. Das Umleiten der Wege wird kombinatorisch über ein *Token-Game* beschrieben. Mit Hilfe der Resultate von Robertson und Seymour zu *rural societies* (cf. Theorem 6.1) sollen Kreuze in Brücken zwischen Fundamentalwegen gefunden werden, um die Brückenkonfiguration zur Anwendung des Token-Games zu verbessern. Es wird erwartet, dass das Token-Game nur in einfachen Fällen betrachtet werden muss (z.B. dass der Hilfsgraph ein Stern ist). Es gibt die Hoffnung, dass ein Zusammenhang von $2k+2$ ausreicht. Technische Details werden nur oberflächlich diskutiert, da erwartet wird, dass die Umsetzung an vielen Stellen analog zu anderen Resultaten möglich sei (z.B. lange Wegesysteme wie in [1]; *Vortex* verdichten wie in [2]).

2. Von Oktober 2009 bis August 2011 arbeitet Theodor Müller in Hamburg allein weiter. In dieser Zeit fand eine tiefe technische Analyse der Beweisstrategie statt. Dabei offenbaren sich eine Reihe von größeren technischen und konzeptionellen Problemen. Ergebnis dieser Arbeit sind Lösungskriterien für das Token-Game. Dabei werden auch triviale Fundamentalwege berücksichtigt, die zuvor zu Schwierigkeiten geführt haben. Vorläufer der Begriffe *rich* und *rich block* werden entwickelt. Das Verdichten eines solchen Blocks um die Brückenkonfiguration zu verbessern wird für den Fall gelöst, dass der Block nur eine Artikulation des Hilfsgraphen enthält. Es zeigt sich, dass die Beweismethode einen Zusammenhang von $2k+3$ erfordert.

3. Im August 2011 kommen Jan-Oliver Fröhlich und Julian Pott zum Projekt hinzu. Theodor Müller begleitet das Projekt ab diesem Zeitpunkt im Hintergrund und beteiligt sich dann, wenn Probleme auftreten. Der Begriff einer *stable regular decomposition* wird entwickelt. Die Verdichtung eines Blocks mit beliebig vielen Artikulationen wird gelöst. Das Token-Game wird in seiner endgültigen Fassung beschrieben. Die verschiedenen Beweisteile werden organisiert und formalisiert, die dazu benötigten Notationen werden entwickelt und vereinheitlicht. Die enge Verstrickung zwischen den verschiedenen Beweisteilen führt dazu, dass

selbst kleine notwendige technische Änderungen Auswirkungen auf den gesamten Beweis haben, sodass Definitionen und Notationen mehrfach vollständig überarbeitet werden müssen. So wird die endgültige Definition von *rich* erst im November 2013 gefunden. Die Ausarbeitung von technischen Details (z.B. Lemmas 6.4 und 6.5) nimmt viel Zeit in Anspruch. Sowohl konzeptionelle als auch technische Probleme werden häufig zu dritt gemeinsam an der Tafel bearbeitet. Die letztendliche Ausformulierung des Beweises wird von Jan-Oliver Fröhlich vorgenommen. Viele der dabei auftretenden technischen Problemen und Detailfragen werden gemeinsam geklärt. Die Ausarbeitung des Beweises kommt im Februar 2014 zum Abschluss. Im März 2014 verfasst Paul Wollan die Abschnitte 1 und 2 der Arbeit.

Literatur

[1] T. Böhme, K. Kawarabayashi, J. Maharry und B. Mohar. Linear Connectivity Forces Large Complete Bipartite Minors. *J. Combin. Theory B* **99** (2009), 557–582.

[2] R. Diestel, K. Kawarabayashi, T. Müller und P. Wollan. On the excluded minor structure theorem for graphs of large tree-width. *J. Combin. Theory B* **102** (2012), 1189–1210.

Jan-Oliver Fröhlich, Theodor Müller, Julian Pott
Hamburg, 31.03.2014

Jan-Oliver Fröhlich

Research Interests

graph minors, infinite graphs, infinite matroids

Current Occupation

since 2014 postdoc with P. Wollan at Sapienza University of Rome

Education

2014 **PhD in Mathematics**, *University of Hamburg*, Hamburg, Germany.

2009 **Diploma in Mathematics**, *University of Hamburg*, Hamburg, Germany.

2008 **Certificate of Advanced Study in Mathematics**, *University of Cambridge*, Cambridge, United Kingdom.

Publications

Linkages in Large Graphs of Unbounded Tree-Width, *in preparation*, with K. Kawarabayashi, T. Müller, J. Pott, and P. Wollan.

2014 **Linkages in Large Graphs of Bounded Tree-Width**, *preprint*, with K. Kawarabayashi, T. Müller, J. Pott, and P. Wollan, `http://arxiv.org/abs/1402.5549`.

2011 **On the intersection of infinite matroids**, *submitted*, with E. Aigner-Horev and J. Carmesin, `http://arxiv.org/abs/1111.0606`.

2011 **Infinite matroid union**, *submitted*, with E. Aigner-Horev and J. Carmesin, `http://arxiv.org/abs/1111.0602`.

2011 **Linear connectivity forces large complete bipartite minors: An alternative approach**, *J. Combin. Theory Ser. B*, **101** (2011), 502–508, with T. Müller.

Talks

2011 **Elgersburg Workshop 2011**, Elgersburg, Germany.
Title: Taming the apex set

Scholarships

2006–2009 scholar of the Studienstiftung des deutschen Volkes

April 1, 2014