The Erdős-Pósa property for clique minors in highly connected graphs

Reinhard Diestel   Ken-ichi Kawarabayashi[‡]   Paul Wollan [§¶]

Abstract

We prove the existence of a function $f\colon \mathbb{N}^2 \to \mathbb{N}$ such that, for all $p, k \in \mathbb{N}$, every $(k(p-3) + 14p + 14)$-connected graph either has k disjoint K_p minors or contains a set of at most $f(p, k)$ vertices whose deletion kills all its K_p minors. For fixed $p \geq 5$, the connectivity bound of about $k(p-3)$ is smallest possible, up to an additive constant: if we assume less connectivity in terms of k, there will be no such function f.

Key Words: Erdős-Pósa, clique minor, packing, disjoint cycles

1 Introduction

A set of graphs $\mathcal{C}$ has the *Erdős-Pósa property* if there exists a function $f = f(k)$ such that for all $k \geq 1$, any graph G either contains k vertex disjoint subgraphs in $\mathcal{C}$, or there exists a subset of vertices $X \subseteq V(G)$ with $|X| \leq f(k)$ such that every subgraph of G in $\mathcal{C}$ intersects a vertex of X. The name derives from an article of Erdős and Pósa [6] where they show that the set $\mathcal{C}$ of cycles has this property.

Let G and X be graphs. An *extension of X* is a graph that can be contracted to X. An *instance of an X-minor* in G is a subgraph H of G isomorphic to an extension of X. The set $\mathcal{C}$ of cycles can be thought of as the set of extensions of K_3, the complete graph of three vertices. Thus the result of Erdős and Pósa can be reformulated as follows: there exists a function $f(k)$ such that any graph G either contains k disjoint instances of K_3 as a minor, or there exists a subset of vertices $X \subseteq V(G)$ with $|X| \leq f(k)$ such that $G - X$ does not contain K_3 as a minor. For any graph H, let $\mathcal{C}_H$ be the set of extensions of H. Robertson and Seymour [10] have exactly characterized which graphs H have the property that the set $\mathcal{C}_H$ has the Erdős-Pósa property: the set $\mathcal{C}_H$ has the Erdős-Pósa property if and only if H is planar.

The purpose of this article is to prove the following theorem.

Theorem 1.1 *There exists an $\mathbb{N}^2 \to \mathbb{N}$ function f such that, for all $p, k \in \mathbb{N}$, every $(k(p-3) + 14p + 14)$-connected graph G either contains k disjoint instances of a K_p-minor or has a set X of at most $f(p, k)$ vertices such that $G - X$ has no K_p-minor.*

[*]National Institute of Informatics, 2-1-2, Hitotsubashi, Chiyoda-ku, Tokyo, Japan

[†]Research is partly supported by Japan Society for the Promotion of Science, Grant-in-Aid for Scientific Research, by C & C Foundation, by Kayamori Foundation and by Inoue Research Award for Young Scientists.

[‡]Email address: `k_keniti@nii.ac.jp`

[§]University of Rome, "La Sapienza" Department of Computer Science, Via Salaria 113, Rome 00198 Italy

[¶]Email address: `paul.wollan@gmail.com`

We also show that the connectivity bound in Theorem 1.1 is best possible, up to an additive constant, for fixed $p \geq 5$. Indeed for each p we shall find a constant c_p such that for all $k, n \in \mathbb{N}$ there are $(k(p-3) - c_p)$-connected graphs that do not contain k disjoint instances of K_p as a minor but in which no set of at most n vertices kills all their K_p minors. Hence it is not possible to define a function $f(p, k)$ as in Theorem 1.1 that makes the theorem true for all $(k(p-3) - c_p)$-connected graphs.

We will need the following definitions. We write $X \preccurlyeq G$ to express that X is a minor of G. Given an extension H of an X minor in G, a *branch set* the X minor is a maximal subset of vertices of H which is contracted to a single vertex when contracting H to X. By kX we denote the disjoint union of k copies of a graph X. A path starting in $x \in X$ and ending in $y \in Y$ is an X–Y *path* if x is its only vertex in X and y is its only vertex in Y. A set $\mathcal{P}$ of disjoint paths is a *linkage*. If it consists of X–Y paths and these meet all of $X \cup Y$, it is an X–Y *linkage*. (Then $|X| = |\mathcal{P}| = |Y|$.) Two linkages $\mathcal{P}$ and $\mathcal{Q}$ of the same order are *equivalent* if for every $P \in \mathcal{P}$ there exists a (*corresponding*) path $Q \in \mathcal{Q}$ such that P and Q have the same endpoints.

We recall that a *tree decomposition* of a graph G is a pair $(T, \mathcal{W})$ where T is a tree and $\mathcal{W} = \{W_t \subseteq V(G) : t \in V(T)\}$ is a collection of subsets of the vertices of G indexed by the vertices of T. Moreover, the collection of subsets $\mathcal{W}$ satisfies the following:

- $\bigcup_{t \in V(T)} W_t = V(G)$,

- for every edge $e = uv$ in $E(G)$, there exists $t \in V(T)$ such that $v, u \in W_t$, and

- for all $v \in V(G)$, the vertices $\{t \in V(T) : v \in W_t\}$ induce a connected subtree of T.

The *width* of the decomposition $(T, \mathcal{W})$ is $\max_{t \in V(T)} |W_t| - 1$, and the *tree-width* of a graph G is the minimum width of a tree decomposition of G. A *path decomposition* is simply a tree decomposition where the graph T is a path. Given a path decomposition $(P, \mathcal{W})$ where the vertices of P are $v_1, v_2, \ldots, v_k$ and occur in that order on the path, we will often simplify the notation and refer to the path decomposition as $(W_1, W_2, \ldots, W_k)$ where $W_{v_i} =: W_i$ for $1 \leq i \leq k$.

For any further notions not covered here we refer to [5].

The paper is structured as follows. We begin in Section 2 by proving our theorem for graphs of small tree-width. For graphs of large tree-width we shall use a structure theorem or Robertson and Seymour, although we will follow the notation and statement of [4]; this is explained in Section 3. At the end of Section 3 we give a more detailed overview of how the proof then proceeds until the end of Section 7. In Section 8 we give our construction showing that the connectivity bound in Theorem 1.1 is tight.

2 Proof of Theorem 1.1

The proof of Theorem 1.1 proceeds by considering separately the cases of when the tree-width of the graph is large or small. In this, we follow much of the recent work analyzing the existence of clique minors in large graphs. See [1, 2, 3, 7]. The bounded tree-width case is easy:

Theorem 2.1 *For every $w \in \mathbb{N}$ there is a function $f_w \colon \mathbb{N}^2 \to \mathbb{N}$ such that, for all $p, k \in \mathbb{N}$, every graph G of tree-width $< w$ either contains k disjoint instances of a K_p minor or has a set X of at most $f_w(p, k)$ vertices such that $G - X$ has no K_p-minor.*

Proof. For fixed w and p we define $f_w(p,k)$ recursively for $k = 1, 2, \ldots$. Clearly, $f_w(p,1) := 0$ satisfies the theorem for $k = 1$. Given $k \geq 2$, let

$$f_w(p,k) := 2f_w(p, k-1) + w \,.$$

To see that this satisfies the theorem, let G be given, with a tree-decomposition $(T, (V_t)_{t \in T})$ of width $< w$. Direct the edges $t_1 t_2$ of T as follows. Let T_1, T_2 be the components of $T - t_1 t_2$ containing t_1 and t_2, respectively, and put

$$G_1 := G[\bigcup_{t \in T_1} (V_t \setminus V_{t_2})] \quad \text{and} \quad G_2 := G[\bigcup_{t \in T_2} (V_t \setminus V_{t_1})] \,.$$

Direct the edge $t_1 t_2$ towards G_i if G_i has a K_p-minor, thereby giving $t_1 t_2$ either one or both or neither direction.

If every edge of T receives at most one direction, we follow these to a node $t \in T$ such that no edge at t in T is directed away from t. As K_p is connected, this implies that V_t meets every instance of a K_p minor in G [5, Lemma 12.3.1.]. This completes the proof with $X = V_t$, since $|V_t| \leq w \leq f_w(p,k)$ by the choice of our tree-decomposition.

Suppose now that T has an edge $t_1 t_2$ that received both directions. For each $i = 1, 2$ let us ask if G_i has a set X_i of at most $f_w(p, k-1)$ vertices such that $G_i - X_i$ has no K_p-minor. If this is the case for both i, then as earlier there is no K_p-minor in $G - X$ for $X := X_1 \cup X_2 \cup (V_{t_1} \cap V_{t_2})$.

Suppose then that G_1, say, has no such set X_1 of vertices. By the induction hypothesis, G_1 contains $(k-1)$ disjoint instances of a K_p-minor. Since $t_1 t_2$ was also directed towards t_2, there is another such instance in G_2. This gives the desired total of k disjoint instances of a K_p-minor in G. $\qquad\square$

The bulk of the work in proving Theorem 1.1 will be the case of large tree-width:

Theorem 2.2 *For all $p, k \in \mathbb{N}$ there exists $w = w(p,k) \in \mathbb{N}$ such that every $(k(p-3) + 14p + 14)$-connected graph of tree-width at least w contains k disjoint instances of a K_p minor.*

Proof of Theorem 1.1, assuming Theorems 2.1 and 2.2. Given $p, k \in \mathbb{N}$ define $f(p,k) := f_w(p,k)$, where $w = w(p,k)$ is provided by Theorem 2.2 and f_w by Theorem 2.1. Let G be a $(k(p-3) + 14p + 14)$-connected graph. If G has tree-width $< w$, the assertion which Theorem 1.1 makes about G is tantamount to that of Theorem 2.1. If G has tree-width at least w, the assertion follows from Theorem 2.2. $\qquad\blacksquare$

Given Theorem 2.2, one might ask if a stronger statement might be true: whether there exists a constant c such that every sufficiently large $(k(p-3) + cp)$-connected graph contains k disjoint instances of a K_p minor. However, the bound on the connectivity is not sufficient for such a strengthening. Consider the complete bipartite graph $K_{k(p-2)+cp, T}$ for large values of T and some fixed constant c. Such a graph cannot contain k disjoint instances of K_p as a minor (assuming k is chosen to be at least $cp + 1$). However, the graph has tree-width $k(p-2) + cp$, i.e. the tree-width is bounded with respect to k and p.

The proof of Theorem 2.2 will occupy us until the end of Section 7. Let $p, k \in \mathbb{N}$ be given, *and fixed until the end of the proof of Theorem 2.2.* Several parameters defined in the course of the proof will depend implicitly on this choice of p and k.

3

Given positive integers ℓ and n, let us define the ℓ-ladder $L(\ell)$ and the fan $F(\ell, n)$ as follows. Let $P = u_1 \ldots u_\ell$ and $Q = v_1 \ldots v_\ell$ be disjoint paths, and let $L(\ell)$ be obtained from their union by adding all the edges $u_i v_i$. To obtain $F(\ell, n)$ from $L(\ell)$, add n independent vertices $w_1, \ldots, w_n$, and join each of these to all the vertices of Q.

It is easy to see that $F(p, p - 3)$ has a K_p minor: with $p - 3$ two-vertex branch sets of the form $\{v_i, w_i\}$, and three further branch sets $\{v_{p-2}\}$, $\{v_{p-1}\}$, and $\{v_p, u_p, u_{p-1}, u_{p-2}\}$. Consequently, $F(kp, k(p - 3))$ contains k disjoint instances of a K_p minor. It will thus suffice for our proof of Theorem 2.2 to find a $F(kp, k(p-3))$-minor in the graph under consideration.

3 The excluded minor theorem

In this section, we present a structure theorem for graphs with no large clique minor of Robertson and Seymour [13]. We follow the notation and exact statement in [4].

A *vortex* is a pair $V = (G, \Omega)$, where G is a graph and $\Omega =: \Omega(V)$ is a linearly ordered set $(w_1, \ldots, w_n)$ of vertices in G. These vertices are the *society vertices* of the vortex; the number n is its *length*. We do not always distinguish notationally between a vortex and its underlying graph; for example, a *subgraph of* V is just a subgraph of G. Also, we will often use Ω to refer both to the linear order of the vertices $w_1, \ldots, w_n$ as well as the set of vertices $\{w_1, \ldots, w_n\}$.

A path–decomposition $\mathcal{D} = (X_1, \ldots, X_m)$ of G is a *decomposition of* V if $m = n$ and $w_i \in X_i$ for all i. The *depth* of the vortex V is the minimum width of a path–decomposition of G that is a decomposition of V.

The *adhesion* of our decomposition $\mathcal{D}$ is the maximum value of $|X_{i-1} \cap X_i|$, taken over all $1 < i \leq n$. Write $Z_i := (X_{i-1} \cap X_i) \backslash \Omega$, for all $1 < i \leq n$. Then, $\mathcal{D}$ is *linked* if

  i. all these Z_i have the same size;

  ii. there are $|Z_i|$ disjoint Z_i–Z_{i+1} paths in $G[X_i] - \Omega$, for all $1 < i < n$;

  iii. $X_i \cap \Omega = \{w_i, w_{i+1}\}$ for all $i = 1, \ldots, n$, where $w_{n+1} := w_n$.

Note that the union of those Z_i–Z_{i+1} paths is a disjoint union of X_1–X_n paths in G; we call the set of these paths a *linkage* of V with respect to $(X_1, \ldots, X_m)$. We define the *(linked) adhesion* of a vortex to be the minimum adhesion of a (linked) decomposition of that vortex; if it has no linked decomposition, its *linked adhesion* is infinite.

For a positive integer α, a graph G is α–*nearly embeddable* in a surface Σ if there is a subset $A \subseteq V(G)$ with $|A| \leq \alpha$ such that there are natural numbers $\alpha' \leq \alpha$ and $n \geq \alpha'$ for which $G - A$ can be written as the union of $n + 1$ graphs $G_0, \ldots, G_n$ such that the following holds:

  i. For all $1 \leq i \leq n$ and $\Omega_i := V(G_i \cap G_0)$, the pair $(G_i, \Omega_i) =: V_i$ is a vortex, and for $1 \leq i < j \leq n$, $G_i \cap G_j \subseteq G_0$.

  ii. The vortices $V_1, \ldots, V_{\alpha'}$ are disjoint and have adhesion at most α; we denote this set of vortices by $\mathcal{V}$.

  iii. The vortices $V_{\alpha'+1}, \ldots, V_n$ have length at most 3; we denote this set of vortices by $\mathcal{W}$.

iv. There are closed discs in Σ with disjoint interiors $D_1, \ldots, D_n$ and an embedding $\sigma : G_0 \hookrightarrow \Sigma - \bigcup_{i=1}^{n} D_i$ such that $\sigma(G_0) \cap \partial D_i = \sigma(\Omega_i)$ for all i and the generic linear ordering of Ω_i is compatible with the natural cyclic ordering of its image (i.e., coincides with the linear ordering of $\sigma(\Omega_i)$ induced by $[0, 1)$ when ∂D_i is viewed as a suitable homeomorphic copy of $[0, 1]/\{0, 1\}$). For $i = 1, \ldots, n$ we think of the disc D_i as *accommodating* the (unembedded) vortex V_i, and denote D_i as $D(V_i)$.

We call $(\sigma, G_0, A, \mathcal{V}, \mathcal{W})$ an α–*near embedding* of G in Σ.

Let G_0' be the graph resulting from G_0 by joining any two unadjacent vertices $u, v \in G_0$ that lie in a common vortex $V \in \mathcal{W}$; the new edge uv of G_0' will be called a *virtual edge*. By embedding these virtual edges disjointly in the discs $D(V)$ accommodating their vortex V, we extend our embedding $\sigma : G_0 \hookrightarrow \Sigma$ to an embedding $\sigma' : G_0' \hookrightarrow \Sigma$. We shall not normally distinguish G_0' from its image in Σ under σ'.

The more widely known version of the excluded minor theorem of Robertson and Seymour ([12], see also [5]) decomposes a graph not containing a fixed H as a minor into a tree-like structure of α-nearly embeddable graphs, where the value of α depends solely on the graph H. We will need a variation of the structure theorem which ensures both that the vortices are linked and that there is a large grid-like graph embedded in the surface when the graph is assumed to have large tree width.

A vortex (G_i, Ω_i) is *properly attached* to G_0 if, for every pair of distinct vertices $x, y \in \Omega_i$, there is a path P_{xy} in G_i with endvertices x and y and all inner vertices in $G_i - \Omega_i$ and further, for every choice of three distinct vertices $x, y, z \in \Omega_i$, the paths P_{xy} and P_{yz} can be chosen internally disjoint.

The *distance* of two points $x, y \in \Sigma$ is the minimal value of $|G \cap C|$ taken over all curves C in the surface that link x and y and hit the graph in vertices only. The *distance* of two vortices V and W is the minimal distance of a point $v \in D(V)$ and a point $v' \in D(W)$.

When a graph is embedded in a surface, a topological component of the surface minus the graph that is homeomorphic to a disc is a *face*. The *outer cycle* of a 2-connected plane graph is the cycle bounding its infinite face. A cycle C is *flat* if C bounds a disc $D \subseteq \Sigma$. Let $C_1, \ldots, C_n$ be flat cycles that bound discs $D_1, \ldots, D_n$, respectively. The cycles $(C_1, \ldots, C_n)$ are *concentric* if $D_i \supseteq D_{i+1}$ for all $1 \leq i < n$.

For positive integers r, define a graph H_r as follows. Let $P_1, \ldots, P_r$ be r vertex disjoint ('horizontal') paths of length $r - 1$, say $P_i = v_1^i \ldots v_r^i$. Let $V(H_r) = \bigcup_{i=1}^{r} V(P_i)$, and let

$$E(H_r) = \bigcup_{i=1}^{r} E(P_i) \cup \left\{ v_j^i v_j^{i+1} \mid i, j \text{ odd}; \ 1 \leq i < r; \ 1 \leq j \leq r \right\}$$

$$\cup \left\{ v_j^i v_j^{i+1} \mid i, j \text{ even}; \ 1 \leq i < r; \ 1 \leq j \leq r \right\}.$$

The 6-cycles in H_r are its *bricks*. In the natural plane embedding of H_r, these bound its 'finite' faces. The outer cycle of the unique maximal 2-connected subgraph is called the *boundary cycle* of H_r.

Any subdivision $H = TH_r$ of H_r will be called an r–*wall*. The *bricks* and the *boundary cycle* of H are its subgraphs that form subdivisions of the bricks and the boundary cycle of H_r, respectively. The *first n boundary cycles* $C_1, \ldots, C_n$ of H_r are defined inductively: C_n is the outer cycle (in the induced embedding) of the unique maximal 2-connected subgraph $H_r^{-(n-1)}$ of $H_r - (C_1 \cup \ldots \cup C_{n-1})$. An embedding of H in a surface Σ is a *flat* embedding, and H is

5

flat in Σ, if the boundary cycle C of H bounds a disc that contains a vertex of degree 3 of $H - C$. We refer to the disc bounded by C as $\Delta(\Sigma, H)$.

An α–near embedding of a graph G in some surface Σ is β–*rich* if the following statements hold:

i. G_0' contains a flat r–wall H for some $r \geq \beta$.

ii. For every vortex $V \in \mathcal{V}$ there are β disjoint, concentric cycles $(C_1, \ldots, C_\beta)$ in G_0' that bound discs $(D_1, \ldots, D_\beta)$, respectively, the innermost disc D_β contains $\Omega(V)$ and H does not intersect with D_1.

iii. Every two vortices in $\mathcal{V}$ have distance at least β.

iv. Let $V \in \mathcal{V}$ with $\Omega(V) = (w_1, \ldots, w_n)$. Then there is a linked decomposition of V of adhesion at most α and a path P in $V \cup \bigcup \mathcal{W}$ with $V(P \cap G_0) = \Omega(V)$, avoiding all the paths of the linkage of V, and traversing $w_1, \ldots, w_n$ in their order.

v. For every vortex $V \in \mathcal{V}$ the society vertices $\Omega(V)$ are linked in G_0' to the vertices of H of degree 3 by a path system of β disjoint paths and these paths have no inner vertices in H.

vi. All vortices in $\mathcal{W}$ are properly attached to G_0.

Theorem 3.1 *For every graph R, there is an integer α such that for every integer β there is an integer $w = w(R, \beta)$ such that the following holds. Every graph G with $\mathrm{tw}(G) \geq w$ that does not contain R as a minor has an α–near, β–rich embedding in some surface Σ in which R cannot be embedded.*

Here is an outline of how we shall use Theorem 3.1 in our proof of Theorem 2.2. By Euler's formula, a graph embedded in a fixed surface has average degree at most $6 + o(1)$ (in terms of its order). The high connectivity we assumed for our graph G thus implies that, when we apply Theorem 3.1 to it, G cannot be entirely embedded in Σ: when the wall $H \subseteq G_0'$ gets large, the embedded subgraph G_0' of G must have many vertices of degree at most 6. These vertices send their remaining edges outside G_0': to the apex set A, to components of $G_0 - G_0'$, or into the vortices $G_1, \ldots, G_{\alpha'}$.

Distinguishing vertices of large and small degree in G_0' will be crucial to our proof. However, we put the threshold a little higher than 6, at $10p$. We shall first show, in Section 4, that by carefully choosing a subwall H' of H, we can ensure that the vertices of G_0' in $\Delta(\Sigma, H')$ have large degree in G_0', and have no neighbours outside G_0' other than in A. In Sections 5 and 6 we then find a large linkage in G_0' from a cycle deep inside H' to vertices that have small degree in G_0'. These vertices send many edges out of G_0'. If these edges go directly to A or to components of $G_0 - G_0'$ (which in turn sends many edges to A, by the connectivity of G), we can build from this linkage, some cycles in H' through which it passes, and many common neighbours in A of the endvertices of our linkage or of those components, an instance of an $F(kp, k(p-3))$-minor which contains our desired kK_p-minor. Otherwise, most of the endvertices of our linkage send their many edges out of G_0' into vortices, and many into the same vortex. We shall then find our kK_p minor using that vortex (Section 7).

4 Isolating a subwall in a disc with all degrees large

Our aim in this section is to show that when we apply Theorem 3.1 to our highly connected graph G, we can choose a subwall H' of the wall H so that the vertices of G'_0 in $\Delta(\Sigma, H)$ have large degree in G'_0, and have no neighbours outside G'_0 other than in A.

Lemma 4.1 *Let $\alpha \in \mathbb{N}$ be as provided by Theorem 3.1 for $R = kK_p$. For every $r \in \mathbb{N}$ there exists $w \in \mathbb{N}$ such that every $(k(p-3)+14p+14)$-connected graph $G \not\succeq kK_p$ of tree-width at least w admits an α-near β-rich embedding for some $\beta \geq r$ such that there exists an r-wall H' contained in $G'_0 \cap \Delta(\Sigma, H)$ with the property that every vertex in $\Delta(\Sigma, H')$ has degree at least $10p$ in G'_0 and has no neighbour in $G - A$ outside G'_0.*

Proof.   Let r be given. We will choose $\beta = \beta(r)$ below; it must be sufficiently large to guarantee the β wall H in an α-near β-rich embedding contains enough disjoint r-walls so that if none of these can serve as H' for our lemma, we can combine them all to find a kK_p minor. Given such a β, the existence of w is then implied by Theorem 3.1. Let G be a $(k(p-3)+14p+14)$-connected graph with an α-near β-rich embedding in Σ. Choose the α-near embedding so that $|G'_0|$ is minimum. This implies that for every subwall H' of H the graph $G'_0 \cap \Delta(\Sigma, H')$ is connected: any component other than that containing H' could be included in V_i for some $V \in \mathcal{W}$, decreasing $|G'_0|$.

Consider a component C of $G_0 - G'_0$, and pick a vertex $v \in C$. Then C is separated from G'_0 in $G - A$ by the at most 3 vertices in G'_0. Since G is $(k(p-3)+14p+14)$-connected, this means that C has at least $k(p-3)$ distinct neighbours in A. Let G' be obtained from G by contracting every component C of $G_0 - G'_0$ to one vertex; for every vertex $v \in C$ we denote this new vertex contracted from C as v'.

Call a vertex u of G'_0 in $\Delta(\Sigma, H)$ *bad* if it has degree $< 10p$ in G'_0 or has a neighbour in $(G-A) - G'_0$. If u has a neighbour v in $(G-A) - G'_0$, then v must lie in $G_0 - G'_0$; recall that, by definition a β-rich α-near embedding, the disc $\Delta(\Sigma, H)$ contains no vertex from any vortex $V \in \mathcal{V}$. In G', the contracted vertex v' has $k(p-3)$ neighbours in A. Similarly if u has degree < 10 in G'_0 but no neighbour in $(G-A) - G'_0$, then u itself has more than $k(p-3)$ neighbours in A, by the connectivity assumed for G.

By making β large enough in terms of r and ℓ (see below), we can find in H an instance of an $L(\ell)$-minor (an ℓ-ladder) in which every branch set induces a subgraph in H containing an r-wall, and these r-walls H_i are sufficiently spaced out in $\Delta(\Sigma, H)$ that the discs $\Delta(\Sigma, H_i)$ are disjoint and not joined by edges of G'_0. In particular, for any vortex $V \in \mathcal{W}$, the corresponding vertices $\Omega(V)$ meet at most one of these $\Delta(\Sigma, H_i)$. If one of these discs $\Delta(\Sigma, H_i)$ contains no bad vertex, our lemma is proved with $H' := H_i$. So assume that each of them contains a bad vertex. Let $H_1, \ldots, H_\ell$ be the r-walls from the branch sets of the 'top' row of our ℓ-ladder minor, and put $\Delta_i = \Delta(\Sigma, H_i)$ for $i = 1, \ldots, \ell$. For each i, pick a bad vertex $u_i \in \Delta_i$. If u_i has a neighbour v_i in $(G-A) - G'_0$, its neighbour v'_i in G' has (in G) at least $k(p-3)$ neighbours in A, and these v'_i are distinct for different i. Let G'' be obtained from G' by contracting the edge $u_i v'_i$, and call the contracted vertex w_i. If u_i has no neighbour in $(G-A) - G'_0$, then u_i itself has $k(p-3)$ neighbours in A; let us rename these u_i as w_i.

For each $i = 1, \ldots, \ell$, the vertex w_i has, in G', a set A_i of $k(p-3)$ neighbours in A. We now choose ℓ large enough that for kp values of i, say those in I, the sets A_i coincide. (Notice that ℓ depends only on α, k and p, all of which are constant.) Let A' denote this common set A_i for all $i \in I$. Together with A' and the vertices v'_i with $i \in I$, our instance of an

$L(\ell)$-minor in H' contains an instance of an $F(kp, k(p-3))$-minor in G': the $k(p-3)$ vertices in A' form singleton branch sets, their neighbouring branch sets are sets $V(G'_0) \cap \Delta_i$ for $i \in I$, plus v'_i as appropriate (recall that these sets are connected by the minimality of $|G'_0|$), and the remaining branch sets found in our ladder $L(\ell)$. Thus, $kK_p \preccurlyeq F(kp, k(p-3)) \preccurlyeq G' \preccurlyeq G$, contradicting our choice of G. □

For easier reference later, let us summarize as a formal hypothesis the properties ensured by Lemma 4.1 along with the aspects of a β-rich embedding we will need as we go forward. We will be able to ensure these properties as long as the graph we are interested has sufficiently large tree width. Let Σ and $\alpha \in \mathbb{N}$ be as provided by Theorem 3.1 for $R = kK_p$ applied to the graph G be a graph. Let $r > 0$ an integer.

Hypothesis H(G, r): The graph G is $(k(p-3)+14p+14)$-connected graph and has no kK_p minor. The graph G has an α-near embedding satisfying the following properties:

i. There is a flat r-wall H in G'_0.

ii. Every vertex $v \in G'_0 \cap \Delta(\Sigma, H)$ has degree at least $10p$ in $G'_0 \cap \Delta(\Sigma, H)$ and for every vortex $V \in \mathcal{W}$, the vertices $\Omega(V)$ are disjoint from $G'_0 \cap \Delta(\Sigma, H)$.

iii. Let $V \in \mathcal{V}$ with $\Omega(V) = (w_1, \ldots, w_n)$. Then, there is a linked decomposition $(X_1, \ldots, X_n)$ of V of adhesion at most α and there is a path P in $V \cup \bigcup \mathcal{W}$ with $V(P \cap G_0) = \Omega(V)$, the path P is disjoint to all paths of the linkage of V and traverses $w_1, \ldots, w_n$ in their linear order.

iv. All vortices in $\mathcal{W}$ are properly attached to G_0.

Lemma 4.1 says that, for every $r \in \mathbb{N}$, every $(k(p-3)+14p+14)$-connected graph $G \not\succcurlyeq kK_p$ of large enough tree-width satisfies Hypothesis H(G, r). Note that if G satisfies H(G, r) then it also satisfies H(G, r') for every $r' \leq r$: just take an r'-wall H' inside the given r-wall H.

5   Optimizing linkages

In this section we prove three lemmas about linkages, which may also be of use elsewhere.

An X–Y linkage $\mathcal{P}$ in a graph G is *singular* if $V(\bigcup \mathcal{P}) = V(G)$ and G does not contain any other X–Y linkage.

Lemma 5.1 *If a graph G contains a singular linkage $\mathcal{P}$, then G has path-width at most $|\mathcal{P}|$.*

Proof. Let $\mathcal{P}$ be a singular X–Y linkage in G. Applying induction on $|G|$, we show that G has a path-decomposition $(X_0, \ldots, X_n)$ of width at most $|\mathcal{P}|$ such that $X \subseteq X_0$. Suppose first that every $x \in X$ has a neighbour $y(x)$ in G that is not its neighbour on the path $P(x) \in \mathcal{P}$ containing x. Then $y(x) \notin P(x)$ by the uniqueness of $\mathcal{P}$. The digraph on $\mathcal{P}$ obtained by joining for every $x \in X$ the 'vertex' $P(x)$ to the 'vertex' $P(y(x))$ contains a directed cycle D. Let us replace in $\mathcal{P}$ for each $x \in X$ with $P(x) \in D$ the path $P(x)$ by the X–Y path that starts in x, jumps to $y(x)$, and then continues along $P(y(x))$. Since every 'vertex' of D has in- and outdegree both 1 there, this yields an X–Y linkage with the same endpoints as $\mathcal{P}$ but

different from $\mathcal{P}$. This contradicts our assumption that $\mathcal{P}$ is singular. Thus, there exists an $x \in X$ without any neighbours in G other than (possibly) its neighbour on $P(x)$. Consider this x.

If $P(x)$ is trivial, then x is isolated in G and $x \in X \cap Y$. By induction, $G - x$ has a path-decomposition $(X_1, \ldots, X_n)$ of width at most $|\mathcal{P}| - 1$ with $X \setminus \{x\} \subseteq X_1$. Add $X_0 := X$ to obtain the desired path-decomposition of G. If $P(x)$ is not trivial, let x' be its second vertex, and replace x in X by x' to obtain X'. By induction, $G - x$ has a path-decomposition $(X_1, \ldots, X_n)$ of width at most $|\mathcal{P}|$ with $X' \subseteq X_1$. Add $X_0 := X \cup \{x'\}$ to obtain the desired path-decomposition of G. $\qquad\square$

Our next lemma will help us re-route segments of an X–Y linkage $\mathcal{P}$ in G through a subgraph $H \subseteq G$, which may or may not intersect $\bigcup \mathcal{P}$. Let $\mathcal{Q}$ be a set of disjoint paths that start in H, have no further vertices in H, and end in $\bigcup \mathcal{P}$. (They may have earlier vertices on $\mathcal{P}$.) The $(\mathcal{Q}, H)$-*segment* of a path $P \in \mathcal{P}$ is the unique maximal subpath of P that starts and ends in a vertex of $\bigcup \mathcal{Q} \cup H$; this subpath may be trivial, or even empty. We call $\mathcal{Q}$ an H–$\mathcal{P}$ *comb* if the set of endvertices of $(\mathcal{Q}, H)$-segments of paths in $\mathcal{P}$ equals the set of final vertices of paths in $\mathcal{Q}$.

Lemma 5.2 *Let t be an integer, let $\mathcal{P}$ be an X–Y linkage in a graph G, and let $H \subseteq G$. If G contains t disjoint H–$(X \cup Y)$ paths, then G contains an H–$\mathcal{P}$ comb consisting of at least t paths.*

Proof. Let $\mathcal{Q}$ be a set of as many disjoint H–$(X \cup Y)$ paths as possible, chosen with the least possible number of edges not in $\bigcup \mathcal{P}$. By the maximality of $\mathcal{Q}$, every endvertex v of a $(\mathcal{Q}, H)$-segment of a path $P \in \mathcal{P}$ lies on a path $Q \in \mathcal{Q}$. By our choice of $\mathcal{Q}$, the final segment vQ of Q then lies in P. Deleting the final segments $\mathring{v}Q$ after v for each such endvertex of a $(\mathcal{Q}, H)$-segment turns $\mathcal{Q}$ into an H–$\mathcal{P}$ comb. $\qquad\square$

While it is not typically true that a subset of a comb will again be a comb, the following is true. We omit the straightforward proof.

Observation 5.3 *Let $\mathcal{P}$ be a linkage and H a subgraph in a graph G. Let $\mathcal{R}$ be an $H - \mathcal{P}$ comb. Then for any sublinkage $\mathcal{P}'$ of $\mathcal{P}$, the linkage*

$$\mathcal{R}' := \{R \in \mathcal{R} : \text{ there exists a } (\mathcal{R}, H)\text{-segment in } \mathcal{P}' \text{ sharing an endpoint with } R\}$$

is a $H - \mathcal{P}'$ comb.

We finally turn to linkages in graphs that are, for the most part, embedded in a cylinder. Let $C_1, \ldots, C_s$ be disjoint cycles. A linkage $\mathcal{P}$ is *orthogonal* to $C_1, \ldots, C_s$ if for all $P \in \mathcal{P}$, $V(P) \cap V(C_i) \neq \emptyset$ for all $1 \leq i \leq s$ and P intersects the cycles $C_1, C_2, \ldots, C_s$ in that order when traversing P from one endpoint to the other. Moreover, each of the graphs $P \cap C_i$ is a path (possibly consisting of a single vertex). The next lemma is a weaker version of Theorem 10.1 of [3]. We include its proof for completeness.

Lemma 5.4 *Let s, s', and t be positive integers with $s \geq s' + t$. Let G' be a graph embedded in the plane and let $(C_1, \ldots, C_s)$ be concentric cycles in G'. Let G'' be another graph, with $V(G') \cap V(G'') \subseteq V(C_1)$. Assume that $G' \cup G''$ contains an X–Y linkage $\mathcal{P} = \{P_1, \ldots, P_t\}$ with $X \subseteq C_s$ and $Y \subseteq C_1$. Then there exist concentric cycles $(C_1', \ldots, C_{s'}')$ in G', a set $X' \subseteq V(C_s')$, and an X'–Y linkage $\mathcal{P}'$ in $G' \cup G''$ such that $\mathcal{P}'$ is orthogonal to $C_1', \ldots, C_{s'}'$.*

Proof. Assume the lemma is false, and let G', G'', $\mathcal{P}$, and $(C_1, \ldots, C_s)$ form a counterexample containing a minimal number of edges. To simplify the notation, we let $G = G' \cup G''$. By minimality, it follows that the graph $G = \bigcup_1^s C_i \cup \mathcal{P}$. Also, for all $P \in \mathcal{P}$ and for all $1 \leq i \leq s$, every component of $P \cap C_i$ is a single vertex. If $P \cap C_i$ had a component that was a non-trivial path containing an edge e, then G'/e would form a counterexample with fewer edges. Similarly, we conclude that $V(G) = V(\mathcal{P})$.

Note that no subpath $Q \subseteq \mathcal{P} \cap G'$ that is internally disjoint from $\bigcup_1^s C_i$ has both endpoints contained in C_j for some $1 \leq j \leq s$. There are two cases to consider. If $Q \subseteq \Delta(C_s)$, we violate our choice of a minimal counterexample by restricting the $\mathcal{P}$ path containing Q to a subpath from Y to $V(C_s)$ avoiding the edges of Q. If $Q \not\subseteq \Delta(C_s)$, we could reroute C_j through the path Q to find s concentric cycles in G' and again contradict our choice of a counterexample containing a minimal number of edges. We claim:

$$\textit{The graph } G \textit{ consists of a singular linkage.} \tag{1}$$

To see that the claim is true, observe that $E(\mathcal{P})$ is disjoint from $E\left(\bigcup_1^s C_i\right)$. It follows that if there exists a linkage $\overline{\mathcal{P}}$ from X to Y distinct from $\mathcal{P}$, then at least one of the edges of $\mathcal{P}$ is not contained in $\overline{\mathcal{P}}$. We conclude that the subgraph $\bigcup_1^s C_i \cup \overline{\mathcal{P}}$ forms a counterexample to the claim with fewer edges, a contradiction. This proves (1).

A *local peak* of the linkage $\mathcal{P}$ is a subpath $Q \subseteq \mathcal{P}$ such that Q has both endpoints on C_j for some $j > 1$ and every internal vertex of $Q \cap \left(\bigcup_{i \neq j} V(C_i)\right) \subseteq V(C_{j-1})$. As we have seen above, it must then be the case that $V(Q) \cap V(C_{j-1}) \neq \emptyset$ when $j > 1$.

We claim the following.

$$\textit{For all } j > 1, \textit{ there does not exist a local peak with endpoints in } C_j. \tag{2}$$

Fix Q to be a local peak with endpoints in C_j with Q chosen over all such local peaks so that j is maximal. Assume Q is a subpath of $P \in \mathcal{P}$. Let the endpoints of Q be x and y. Lest we re-route P through C_j and find a counter-example containing fewer edges, there exists a component $P' \in \mathcal{P}$ intersecting the subpath of C_j linking x and y. By planarity, P' either contains a subpath internally disjoint from the union of the C_i with both endpoints in C_s, or P' contains a subpath forming a local peak with endpoints in C_{j-1}. Either is a contradiction to our choice of a minimal counterexample. This proves (2).

An immediate consequence of (1) and (2) is the following. For every $P \in \mathcal{P}$, let x be the endpoint of P in X and let y be the vertex of $V(C_1) \cap V(P)$ closest to x on P. Define the path $\overline{P}$ be the subpath xPy of P. The path $\overline{P}$ is orthogonal to the cycles $C_1, \ldots, C_s$. In fact, $\overline{P} \cap C_i$ is a single vertex for each $1 \leq i \leq s$. The final claim will complete the proof.

$$\textit{For all } P \in \mathcal{P}, \textit{ the path } P - \overline{P} \textit{ does not intersect } C_{t+1}. \tag{3}$$

To see (3) is true, fix $P \in \mathcal{P}$ such that $(P - \overline{P}) \cap C_{t+1} \neq \emptyset$. It follows now from (2) that $P - \overline{P}$ contains a subpath Q with one endpoint in C_{t+1} and one endpoint in C_1 such that Q is orthogonal to the cycles $C_{t+1}, C_t \ldots, C_1$. By the planarity of G', we see that G contains a subgraph isomorphic to the subdivision of the $(t+1) \times (t+1)$ grid. This contradicts (1) and Lemma 5.1, proving (3).

We conclude that $\mathcal{P}$ is orthogonal to the s' disjoint cycles $C_s, C_{s-1}, \ldots, C_{t+1}$. This contradicts our choice of G, and the lemma is proven. $\qquad\square$

6 Linking the wall to a vortex

Consider a graph G satisfying Hypothesis $H(G, r)$. Our first aim in this section is to find a large linkage from a cycle deep inside H to vertices of small degree in G'_0. By Lemma 5.4 we shall be able to assume that this linkage is orthogonal to a pair of cycles C and C'. If the many of the last vertices of our linkage send many edges to A, or an edge to a component of $G_0 - G'_0$ (which in turn sends many edges to A, by the connectivity of G), we shall be able to convert the cycles C and C', the linkage, and those neighbours into an $F(kp, k(p-3))$-minor, completing the proof. If not, then most of those last vertices send many edges into vortices. As we have only a bounded number of vortices, many send their edges to the same vortex. That case we shall treat in Section 7.

Lemma 6.1 *For all positive integers t and s there exists an integer $R = R(s, t)$ such for every graph G satisfying Hypothesis $H(G, r)$ with $r \geq R$ there are t disjoint X–Y paths in G'_0, where X is the vertex set of the s'th boundary cycle C_s of H, and $Y := \{v \in V(G'_0) : d_{G'_0}(v) < 10p\}$.*

Proof. If the desired paths do not exist then, by Menger's theorem, G'_0 has a separation (A, B) of order less than t with $X \subseteq A$ and $Y \subseteq B$. By the choice of Y, every vertex in $A \setminus B$ has degree at least $10p$ in G'_0. The sum of all these degrees is at least $10p |A \setminus B|$, so $G'_0[A]$ has at least $5p |A \setminus B|$ edges. As $|A| \geq |X| \geq r - 4s$, and Σ is determined by our constants p and k, choosing R sufficiently large in terms of s and t yields

$$5p |A \setminus B| \geq 5p (|A| - t) > 3 |A| - 3 \chi(\Sigma),$$

which is the maximum number of edges a graph of order $|A|$ embedded in Σ can have (by Euler's formula). As $G'_0[A]$ is such a graph, this is a contradiction. $\square$

Our next lemma says that by rerouting the paths if necessary we can make the linkage from Lemma 6.1 orthogonal to two concentric cycles. Recall that the wall H in Hypothesis $H(G, r)$ is flat; we think of the topological disc $\Delta(\Sigma, H) \subseteq \Sigma$, which contains H and is bounded by its outer cycle C_1, as a disc in $\mathbb{R}^2$.

Lemma 6.2 *Let t be an integer. Let G be a graph satisfying Hypothesis $H(G, r)$ for some r large enough that H has boundary cycles $C_1, \ldots, C_{t+2}$. Suppose further that G'_0 contains an X–Y linkage $\mathcal{P}$ of order t, where $X \subseteq V(C_{t+2})$ and $Y \subseteq V(G'_0) \setminus \Delta(\Sigma, H)$. Then $\bigcup \mathcal{P} \cup C_1 \cup \ldots \cup C_{t+2} \subseteq G'_0$ contains disjoint cycles C'_1, C'_2 in $G'_0 \cap \Delta(\Sigma, H)$, and an X'–Y linkage orthogonal to C'_1, C'_2 with $X' \subseteq V(C_{t+2})$.*

Proof. Let Y' be the set of the last vertices in $\Delta(C_1)$ of paths in $\mathcal{P}$; this is a subset of $V(C_1)$. Let $\mathcal{P}'$ be the set of X–Y' paths contained in the paths in $\mathcal{P}$ (one in each). Let G' be the union of all the cycles $C_1, \ldots, C_{t+2}$ and the subpaths in $\Delta(C_1)$ of paths in $\mathcal{P}'$. Then G' is planar and $(C_1, \ldots, C_{t+2})$ form a set of concentric cycles. Let G'' be the union of the remaining segments of paths in $\mathcal{P}'$; then $G' \cup G'' = \bigcup \mathcal{P}' \cup C_1 \cup \ldots \cup C_{t+2}$, and $V(G') \cap V(G'') \subseteq V(C_1)$. Applying Lemma 5.4 to the linkage $\mathcal{P}'$ in $G' \cup G''$, we obtain a two disjoint cycles C'_1 and C'_2 contained in $\Delta(C_1)$ and a set $X' \subseteq (C'_2)$ such that the cycles are orthogonal to an X'–Y' linkage $\mathcal{P}''$ in $G' \cup G''$. Append to this linkage the Y'–Y paths contained in the paths from $\mathcal{P}$ (which meet $G' \cup G''$ only in Y', by the choice of Y') to obtain the desired linkage for the lemma. $\square$

11

Note that the linkage obtained in Lemma 6.2 has the same order as $\mathcal{P}$, since the target set Y remained unchanged. The proof of the lemma could clearly be modified to provide a set of s cycles orthogonal to the linkage, for arbitrary s, rather than just two, but we shall only need two in the following arguments.

We return now to the linkage provided by Lemma 6.1. Using Lemma 6.2, we show that all but a bounded number of the paths of that linkage lie in G (that is, contain no virtual edges) and end in vortices.

Lemma 6.3 *Let $t \geq 3kp\binom{\alpha}{k(p-3)}$, and let G be a graph satisfying Hypothesis $\mathrm{H}(G,r)$, for some r large enough that the $(t+2)$th boundary cycle C_{t+2} of H exists. Then G_0' contains no X–$(Y \cup Z)$ linkage of order t such that $X \subseteq V(C_{t+2})$, the set Z contains vertices of $\bigcup_{V \in \mathcal{W}} \Omega(V)$, and $Y \subseteq \{v \in V(G_0') : d_{G_0'}(v) < 10p\} \setminus \bigcup_{V \in \mathcal{V} \cup \mathcal{W}} \Omega(V)$.*

Proof. Suppose there is an X–$(Y \cup Z)$ linkage in G_0' as stated. Since $Y \cup Z \subseteq V(G_0') \setminus \Delta(\Sigma, H)$, Lemma 6.2 provides us with an X'–$(Y \cup Z)$ linkage $\mathcal{P}$ in G_0' that is orthogonal to two cycles C_1', C_2' contained in $\Delta(\Sigma, H)$, where again $X' \subseteq V(C_{t+2})$. Since $\mathcal{P}$ has the same target set $Y \cup Z$ as the original linkage, it also has the same order $t \geq 3kp\binom{\alpha}{k(p-3)}$.

Each path in $\mathcal{P}$ contains at most one vertex contained in a vortex, since this will be its last vertex. Also, for every vortex $V \in \mathcal{W}$, $\Omega(V)$ intersects at most 3 paths in $\mathcal{P}$. We find a subset $\mathcal{P}'$ of $\mathcal{P}$ of order $kp\binom{\alpha}{k(p-3)}$ and for every $P \in \mathcal{P}'$ such that P intersects some vortex, we assign a vortex $V(P)$ such that for $P, Q \in \mathcal{P}'$, $V(P) \neq V(Q)$. We can construct the subset $\mathcal{P}'$ and the vortex assignments greedily - we begin considering $\mathcal{P}$, and as long some path P intersects an unassigned vortex V, we set $V(P) := V$ and delete any other path Q with both $Q \cap \Omega(V) \neq \emptyset$ and $V(Q)$ undefined.

By definition of Y and the connectivity of G, every last vertex $y \in Y \setminus Z$ of a path $P \in \mathcal{P}'$ has a set A_P of $k(p-3)$ distinct neighbours in A. By definition of Z, every last vertex $z \in Z$ of a path $P \in \mathcal{P}'$ sends an edge to some component C of $V(P) - G_0'$. In $G - A$, a set of at most three vertices of G_0' (which includes z) separates C from the rest of G_0'. Hence by the connectivity of G, the component C has a set A_P of $k(p-3)$ distinct neighbours in A. For every such z, contract the component C on to z. (By definition of $\mathcal{P}'$ and the assignment $V(P)$, these C are distinct, and hence disjoint, for different z.) In the resulting minor G' of G, the vertex z is adjacent to every vertex in A_P (for the $P \in \mathcal{P}'$ ending in z).

Since $|\mathcal{P}'| = kp\binom{\alpha}{k(p-3)}$, there is a subset $\mathcal{P}''$ of $\mathcal{P}'$ of order kp such that for all the paths $P \in \mathcal{P}''$ their sets A_P conincide; let us write A' for this subset of A of order $k(p-3)$.

Of each path $P \in \mathcal{P}''$ let us keep only its segment P' between C_2' and C_1', contracting the final segment of P that follows its vertex v_P in C_1' on to v_P. In the minor G'' of G' obtained by all these contractions, the final vertices v_P of the paths P' with $P \in \mathcal{P}''$ are adjacent to all the $k(p-3)$ vertices in A'. The cycles C_1' and C_2', the kp paths P' with $P \in \mathcal{P}''$, and the edges between the vertices v_P and A' together contain a subdivided fan $F(kp, k(p-3))$. Thus,

$$kK_p \preccurlyeq F(kp, k(p-3)) \preccurlyeq G'' \preccurlyeq G' \preccurlyeq G \,,$$

a contradiction. $\square$

Since $b \leq \alpha$, Lemma 6.3 implies that of paths from the linkage of Lemma 6.1 some unbounded number end in the same W_i. These have small degree in G_0'

12

7 Proof of Theorem 2.2

Let r be the integer $R(s,t)$ provided by Lemma 6.1 for

$$t = 2\alpha \left(kp \binom{2\alpha}{k(p-3)} + k \left(\binom{p}{2} + 1 \right) \binom{\alpha}{p} \right) + 3kp \binom{\alpha}{k(p-3)}$$

and $s = t + 2$. Let w be large enough that, by Lemma 4.1, every $(k(p-3) + 14p + 14)$-connected graph $G \not\succeq kK_p$ of tree-width at least w contains an r-wall H such that (G, H) satisfies Hypothesis $H(G, r)$, for this r.

For our proof of Theorem 2.2, let G be a $(k(p-3) + 14p + 14)$-connected graph of tree-width at least w; we have to show that $G \succeq kK_p$. Suppose not. Then (G, H) satisfies Hypothesis $H(G, r)$ for the value of r defined above, by our choice of w.

Let $C_1, C_2, \ldots, C_{t+2}$ be the first $t+2$ boundary cycles of H. By Lemma 6.1, there are t disjoint paths in G_0' from $V(C_{t+2})$ to vertices of degree $< 10p$ in G_0'. By Lemma 6.3, all but at most $3kp\binom{\alpha}{k(p-3)}$ of these paths intersect exactly one vortex V at their endpoints, and furthermore, this vortex V is among the α' vortices of $\mathcal{V}$. Since $\alpha' \leq \alpha$, at least $1/\alpha$ of these paths end in the same vortex, say $V_a = (G_a, \Omega_a)$. These paths, then, form an X–Y linkage $\mathcal{P}$ in G (i.e. the linkage does not contain any of the virtual edges of G_0') of order

$$|\mathcal{P}| \geq 2 \left(kp \binom{2\alpha}{k(p-3)} + k \left(\binom{p}{2} + 1 \right) \binom{\alpha}{p} \right), \tag{4}$$

with $X \subseteq V(C_{t+2})$ and $Y \subseteq \Omega_a =: \{w_1, \ldots, w_m\}$.

Let us add to the graph G_a all the vertices from A (together with the edges they send to G_a), putting them in every part of its vortex decomposition. This does not affect our assumption that this decomposition is linked, since every vertex in A becomes a trivial path in the linkage through G_a. The new (induced) subgraph G_a of G has a path decomposition $(U_1, \ldots, U_m)$ with the following properties (where $U_i^+ := U_i \cap U_{i+1} =: U_{i+1}^-$):

- $A \subseteq U_i$ for all $i = 1, \ldots, m$;

- $U_i \cap \Omega_a = \{w_{i-1}, w_i\}$ for all $i = 1, \ldots, m$ with $w_0 := w_1$;

- all the sets U_i^+ and U_i^- have the same order $(\leq 2\alpha)$;

- $G_a - \Omega_a$ contains a $(U_1^+ \setminus \{w_1\})$–$(U_m^- \setminus \{w_{m-1}\})$ linkage $\mathcal{Q}$.

For each $i = 0, \ldots, m$, let $H_i = G[U_i \cup U_{i+1}]$ (putting $U_0 = \{w_1\}$ and $U_{m+1} = \{w_m\}$). The set $U_i^- \cup U_{i+1}^+ \cup \{w_i\}$ of size at most $4\alpha + 1$ separates H_i from the rest of G (put $U_0^- = U_{m+1}^+ = \emptyset$). Let $\mathcal{Q}_i$ be the set of the segments in H_i of paths in $\mathcal{Q}$. These are U_i^-–U_{i+1}^+ paths, one for each $Q \in \mathcal{Q}$, when $1 < i < m$. We write $\mathcal{T}_i$ for the set of trivial paths in $\mathcal{Q}_i$; when $1 < i < m$, this is the set

$$\mathcal{T}_i = \left\{ \{v\} \mid v \in U_i^- \cap U_{i+1}^+ \right\} \subseteq \mathcal{Q}_i.$$

Note that $\mathcal{T}_i$ contains every path $\{v\}$ with $v \in A$, and that $|\bigcup \mathcal{T}_i| \leq |\mathcal{Q}| < 2\alpha$.

Deleting at most half the paths in $\mathcal{P}$, we can ensure that for the remaining linkage $\mathcal{P}' \subseteq \mathcal{P}$ there is no $i < m$ such that both w_i and w_{i+1} are endpoints of a path in $\mathcal{P}'$. Let $I_1 \subset \{1, \ldots, m\}$ be the set of those i for which w_i is the final vertex of a path in $\mathcal{P}'$.

For each $i \in I_1$, let J_i denote the component of $H_i - U_i^- - U_{i+1}^+$ containing w_i. Note that $J_i \cap \bigcup \mathcal{T}_i = \emptyset$ for each i, and that the J_i are disjoint for different $i \in I_1$. Let $I_2 \subseteq I_1$ be the set of those $i \in I_1$ for which J_i has at least $k(p-3)$ neighbours in $\bigcup \mathcal{T}_i$, and put $I_3 := I_1 \setminus I_2$. Let us show that

$$|I_3| \geq k \left(\binom{p}{2} + 1 \right) \binom{\alpha}{p}. \tag{5}$$

Suppose not; then $|I_2| \geq kp\binom{2\alpha}{k(p-3)}$, by (4). For each $i \in I_2$, the at least $k(p-3)$ neighbours of J_i in $\bigcup \mathcal{T}_i$ lie on different paths in $\mathcal{Q}$. Since $|\mathcal{Q}| \leq 2\alpha$, there is a set of $k(p-3)$ paths Q in $\mathcal{Q}$ and a set $I \subseteq I_2$ of order kp such that for each of those Q and every $i \in I$ we have $Q \cap H_i \in \mathcal{T}_i$ and the unique vertex in this graph sends an edge to J_i. Contract each of these Q to one vertex, and contract each J_i with $i \in I$ on to its vertex w_i. Then each of these kp vertices w_i is adjacent to those $k(p-3)$ vertices contracted from paths in $\mathcal{Q}$. Together with the kp paths in $\mathcal{P}$ ending in these w_i and the cycles $C_1, \ldots, C_{t+2}$ in our wall H, we obtain a fan $F(kp, k(p-3))$ as in the proof of Lemma 6.3, contradicting our assumption that $G \not\succeq kK_p$. This proves (5).

Let $\mathcal{P}''$ be the set of paths in $\mathcal{P}'$ ending in some w_i with $i \in I_3$. For every $i \in I_3$, the graph J_i has at most $k(p-3) - 1$ neighbours in $\bigcup \mathcal{T}_i$. Our plan now is to find some fixed paths $Q^1, \ldots, Q^p \in \mathcal{Q}$ and many indices $i \in I_3$, one for every edge in kK_p, such that for each of these i the segments $Q_i^j := Q^j \cap H_i$ are non-trivial and we can connect two of them by a path through J_i. (This will require some re-routing of $\mathcal{Q}_i$ inside H_i.) Dividing the linkage $(Q^1, \ldots, Q^p)$ into k chunks kept well apart by the $k-1$ subgraphs H_i between them (in which all these paths have non-trivial segments; it is here only that we need the non-triviality of segments), and contracting the p paths in each chunk to p vertices, we shall thus obtain our desired kK_p minor.

Let us begin by choosing the segments $Q_i^1, \ldots, Q_i^p$ locally for each $i \in I_3$, allowing the choice of $Q^1, \ldots, Q^p$ to depend on i. It will be easy later to find enough i for which these choices agree. Let us prove the following:

> For every $i \in I_3$ there are paths $Q^1, \ldots, Q^p \in \mathcal{Q}$ with $Q_i^1, \ldots, Q_i^p \in \mathcal{Q}_i \setminus \mathcal{T}_i$ such
> that for every choice of $1 \leq j < \ell \leq p$ there is a linkage $(\hat{Q}_i^1, \ldots, \hat{Q}_i^p)$ in H_i
> equivalent to $(Q_i^1, \ldots, Q_i^p)$ for which $J_i - (\hat{Q}_i^1 \cup \ldots \cup \hat{Q}_i^p)$ contains a path $R_i^{j,\ell}$ $\qquad$ (6)
> from a vertex adjacent to $\hat{Q}_i^j$ to a vertex adjacent to $\hat{Q}_i^\ell$.

To prove (6), let $i \in I_3$ be given. Note that if any vertex v of J_i sends $p+1$ edges to $U_i^- \setminus (\{w_{i-1}\} \cup \mathcal{T}_i)$ or to $U_{i+1}^+ \setminus (\{w_{i+1}\} \cup \mathcal{T}_i)$, the proof of (6) is immediate with $R_i^{j,\ell} = \{v\}$: since v lies on at most one of the $p+1$ non-trivial paths in $\mathcal{Q}$ to which it sends an edge, we can find p such paths avoiding v, no re-routing being necessary. So let us assume that this is not the case.

Consider the graph $J_i - w_i$. As $i \in I_3$, the vertex w_i has fewer than $k(p-3)$ neighbours in $\bigcup \mathcal{T}_i$, fewer than $10p$ neighbours in G_0' (by definition of $\mathcal{P}$), and at most $2p$ neighbours in $(U_i^- \cup U_{i+1}^+) \setminus (\{w_{i-1}, w_{i+1}\} \cup \mathcal{T}_i)$. As w_i has degree at least $k(p-3) + 14p + 14$ in G, the graph $J_i - w_i$ is non-empty. By the same argument,

$$\delta(J_i - w_i) \geq \left(k(p-3) + 14p + 14 \right) - \left(k(p-3) - 1 \right) - 2p - 3 = 12(p+1). \tag{7}$$

By Mader's theorem [5, Thm. 1.4.3] and the main result from [15] (which says that $2s$-connected graphs of average degree at least $10s$ are s-linked), (7) implies that $J_i - w_i$ has

a $(p + 1)$-linked subgraph H'_i. In particular, $|H'_i| \geq 2p + 2$. Let Z_i consist of the vertices w_{i-1}, w_i, w_{i+1} and the neighbours of J_i in $\bigcup \mathcal{T}_i$. As $i \in I_3$ we have $|Z_i| \leq k(p - 3) + 2$, so $G - Z_i$ is still $2p$-connected. Since $H'_i \subseteq J_i - w_i$, the graph H'_i has no vertex in Z_i. By Menger's theorem, there are $2p$ disjoint paths in $G - Z_i$ from H'_i to our wall H. By definition of J_i, their first vertices outside J_i lie in $U_i^- \cup U_{i+1}^+$ (recall that this set and w_i together separate H_i from H in G), and hence on a path in $\mathcal{Q}_i \setminus \mathcal{T}_i$. By Lemma 5.2, there exists an H'_i–$(\mathcal{Q}_i \setminus \mathcal{T}_i)$ comb of at least $2p$ paths. Each path in $\mathcal{Q}$ meets at most two of them. Observation 5.3 implies that we can find p paths $Q^1, \ldots, Q^p \in \mathcal{Q}$ such that $Q_i^1, \ldots, Q_i^p \in \mathcal{Q}_i \setminus \mathcal{T}_i$ (as in (6)) together with an H'_i–$\{Q_i^1, \ldots, Q_i^p\}$ subcomb $\mathcal{R}$ meeting all of $Q^1, \ldots, Q^p$. Let $\bar{Q}_i^q$ denote the $(\mathcal{R}, H'_i)$-segment of Q_i^q, for each $q = 1, \ldots, p$; these segments are non-empty, but they may be trivial.

We now define the paths $\hat{Q}_i^1, \ldots, \hat{Q}_i^p$. For all q whose $\bar{Q}_i^q$ is trivial we let $\hat{Q}_i^q = Q_i^q$. For those q whose $\bar{Q}_i^q$ is non-trivial, we let $h_1^q \in H'_i$ be the starting vertex of the path $R_1^q \in \mathcal{R}$ that ends on the first vertex of $\bar{Q}_i^q$, and let $h_2^q \in H'_i$ be the starting vertex of the path $R_2^q \in \mathcal{R}$ that ends on the last vertex of $\bar{Q}_i^q$. Our aim is to link h_1^q to h_2^q in H'_i for each q, but we must define $R_i^{j,\ell}$ at the same time. If $\bar{Q}_i^j$ is trivial, let $r^j \in H'_i$ be the starting vertex of the unique path $R^j \in \mathcal{R}$ that ends on $\bar{Q}_i^j$. If $\bar{Q}_i^j$ is non-trivial, let $r^j \in H'_i$ be a neighbour of h_1^j in $H'_i - \bigcup \mathcal{R}$; such a neighbour exists, since H'_i, being $(p + 1)$-linked, is $(2p + 1)$-connected [5, Ex. 3.22]. Define r^ℓ analogously. Now choose a linkage in H'_i consisting of a path $R = r^j \ldots r^\ell$ and paths $R^q = h_1^q \ldots h_2^q$ for all those q such that $\bar{Q}_i^q$ is non-trivial. For these q, let $\hat{Q}_i^q$ be obtained from Q_i^q by replacing $\bar{Q}_i^q$ with $R_1^q \cup R^q \cup R_2^q$. If both $\bar{Q}_i^j$ and $\bar{Q}_i^\ell$ are trivial, let $R_i^{j,\ell}$ be the interior of the path $R^j \cup R \cup R^\ell$. If $\bar{Q}_i^j$ is trivial but $\bar{Q}_i^\ell$ is not, let $R_i^{j,\ell}$ be the path $R^j \cup R$ minus its first vertex. If $\bar{Q}_i^\ell$ is trivial but $\bar{Q}_i^j$ is not, let $R_i^{j,\ell}$ be the path $R \cup R^\ell$ minus its last vertex. If neither $\bar{Q}_i^j$ nor $\bar{Q}_i^\ell$ is trivial, let $R_i^{j,\ell}$ be the path R. This completes the proof of (6).

By (5), we can find a set $I_4 \subseteq I_3$ of $k \left(\binom{p}{2} + 1 \right)$ indices i in I_3 for which the choice of paths $Q^1, \ldots, Q^p$ in (6) coincides. (Recall that these paths are always chosen from the original vortex linkage of order $\leq \alpha$, since the trivial paths $\{v\}$ with $v \in A$ which we added later lie in every $\mathcal{T}_i$.) For notational reasons only, let $\hat{p} := \binom{p}{2}$. Divide I_4 into k segments

$$\left(i_1^1, \ldots, i_{\hat{p}}^1, i^1 \right), \ \ldots, \ \left(i_1^k, \ldots, i_{\hat{p}}^k, i^k \right)$$

of length $\binom{p}{2} + 1$. For every upper index $n = 1, \ldots, k$ contract in each of $Q^1, \ldots, Q^p$ the segment from $H_{i_1^n}$ to $H_{i_{\hat{p}}^n}$ (inclusive) to a vertex, and make these vertices into a K_p minor using the paths $R_i^{j,\ell}$ from (6) for subdivided edges, one for each $i = i_1^n, \ldots, i_{\hat{p}}^n$. Note that the k instances of a K_p minor thus obtained are disjoint, because they are 'buffered' by the unused segments of the paths $Q^1, \ldots, Q^p$ in H_{i^n} for $n = 1, \ldots, k - 1$.

8 Tightness of the connectivity bound

The goal of this section will be to provide a construction of a graph $G_{n,k,p}$ for all integers $p \geq 5$, $k \geq p$, and $n \geq 1$, such that the graph $G_{n,k,p}$ does not contain k disjoint instances of K_p as a minor, nor does the graph $G_{n,k,p}$ contain a subset X of vertices with $|X| \leq n$ such that $G - X$ does not contain K_p as a minor. Moreover, we will construct such a graph $G_{n,k,p}$ that is $(k(p - 3) - \frac{(p-3)(p-4)}{2} - 6)$-connected. This will imply that the connectivity bound obtained in Theorem 1.1 is best possible for all fixed p, $p \geq 5$, up to an additive constant.

For the remainder of this section, we fix $p \geq 5$. Let Σ be an orientable surface of minimum genus in which K_p embeds. The Euler genus of Σ is at most $\frac{(p-3)(p-4)}{6} + 1$ (see [9]).

We will use the following facts (see [9] for details):

Lemma 8.1 *There are at most $\frac{(p-3)(p-4)}{6} + 1$ disjoint instances of K_5-minors in a graph which is embedded in the surface Σ. Moreover, suppose there are connected subgraphs $B_1, \ldots, B_q$ in a graph embedded in the surface Σ, such that each B_i contains a K_5-minor. Assume there is a vertex v such that $v \in V(B_i)$ for each i and $(V(B_i) - \{v\}) \cap (V(B_j) - \{v\}) = \emptyset$ for $i \neq j$. Then $q \leq \frac{(p-3)(p-4)}{6} + 1$.*

Lemma 8.1 can be generalized as follows (again, see [9] for details):

Lemma 8.2 *Suppose there are q disjoint minors isomorphic to $K_{l_1}, K_{l_2}, \ldots, K_{l_q}$ ($l_i \geq 5$ for $i = 1, \ldots, q$), respectively, in a graph G that is embedded in the surface Σ. Then $\Sigma_{i=1}^{q} \left\lceil \frac{(l_i-3)(l_i-4)}{6} \right\rceil \leq \frac{(p-3)(p-4)}{6} + 1$. Suppose there are connected graphs $B_1, \ldots, B_q$ in a graph that is embedded into the surface Σ, such that each B_i contains a K_{l_i}-minor (with $l_i \geq 5$ for $i = 1, \ldots, q$), and there is a vertex v such that $v \in V(B_i)$ for each i and $(V(B_i) - \{v\}) \cap (V(B_j) - \{v\}) = \emptyset$ for $i \neq j$. Then $\Sigma_{i=1}^{q} \left\lceil \frac{(l_i-3)(l_i-4)}{6} \right\rceil \leq \frac{(p-3)(p-4)}{6} + 1$.*

We are almost ready to construct the graph $G(n, k, p)$. We first recall that the *face-width* of a graph embedded in a surface is the minimum number of times a non-contractable loop intersects the embedded graph taken over all possible non-contractable loops. The following observation follows immediately from the definition of face-width.

Observation 8.3 *Let G be a graph embedded in a surface Γ with face-width k. Let X be a set of t vertices in G. Then $G - X$ is embedded in Γ with face width at least $k - t$.*

For a further discussion of face-width, we refer to [9]. We will need the following result.

Theorem 8.4 ([11]) *Let $t \geq 5$ be a positive integer and let Γ be a surface in which K_t can be embedded. Then there exists a value $r = r(\Gamma, t)$ such that every graph embedded in Γ with face-width r contains K_t as a minor.*

Fix r to be the value given by Theorem 8.4 to ensure a graph embedded in Σ contains K_p as a minor. We first construct a graph G' which is embedded in the surface Σ, with the following properties:

1. The face-width of G' embedded in Σ is at least $n + r$.

2. There is a cycle C in G' which bounds a disk D in Σ, and the set of vertices on the outer boundary of the disk D is defined by $V(D)$. We assume that no vertex, except for the vertex set $V(D)$, exists inside the disk D.

3. For each vertex v outside the disk D, there are at least $k(p-3) - \frac{(p-3)(p-4)}{2} - 6$ internally disjoint paths from v to $V(D)$ in G'.

4. G' is 3-connected, and hence each vertex in $V(D)$ has degree at least 3.

16

A graph G' with the desired embedding is known to exist [8]; we outline such a construction. We begin with a 3-connected graph H allowing a *closed 2-cell embedding* in Σ, in other words, a 3-connected graph H which embeds in Σ so that the topological closure of every facial region is homeomorphic to the closed disk. Consider the following operation for a fixed facial region F. The region F is bounded by a cycle C_F in H. We subdivide every edge of C and add a new vertex embedded in the region F adjacent to every vertex on the subdivided cycle C. The resulting graph is 3-connected and the new embedding is a closed 2-cell embedding as well. Note that if we perform this operation on every facial region, the resulting graph will be embedded in Σ with face width at least twice that of the original embedding. Thus by repeatedly performing the operation, we find a 3-connected graph H_1 along with a closed 2-cell embedding in Σ satisfying 1 above.

Given the embedded graph H_1, let H^* be the dual graph with vertex set equal to the set of facial regions and two facial regions are adjacent in H^* if their boundary cycles share an edge. Note that by the 3-connectivity of H_1, the graph H^* is a simple connected graph. Let T be a spanning tree of H^*, and fix a root R of the tree T. Let $F \in V(T) - R$ be a facial region of H_1 forming a leaf in T. Let C_F be the boundary cycle of F, and let e_F be the edge of H_1 shared with the neighboring facial region in T. We subdivide the edge e_F sufficiently many times to add $k(p-3) - \frac{(p-3)(p-4)}{2} - 6$ neighbors in the subdivided e_F for every vertex of $C_F - e_F$. Given that the region is homeomorphic to the disc, it is clear that we can add the edges maintaining the embedding in Σ. Moreover, we maintain 3-connectivity of the graph. In the resulting graph, every vertex of $C_F - e_F$ will have the desired large degree. We repeatedly delete the leaf F from the tree T and apply the same process to a leaf of $V(T) - F$ until only the vertex R remains. Let G' be the resulting graph. We claim G' satisfies $2 - 4$ above with the disc D being the boundary cycle of the facial region R. The properties 2 and 4 follow easily from the construction. To see that we satisfy 3 as well, pick a vertex of $v \in V(G') \setminus V(D)$. We can find the desired paths from v to $V(D)$ by looking at the path of facial regions in T connecting v to $V(D)$. At each facial region along the path, a given vertex has in fact $k(p-3) - \frac{(p-3)(p-4)}{2} - 6$ neighbors on the next facial region. This completes our outline of the construction of G'.

We now define $G = G(n, k, p)$ as follows. Let Z be a set of $k(p-3) - \frac{(p-3)(p-4)}{2} - 6$ vertices. The vertex set of G will be $Z \cup V(G')$, and the edge set will the union of the edges of G' along with every possible edge of the form zd for all $z \in Z$ and $d \in D$. We will see that the graph G satisfies the desired properties.

We first claim that G is $(k(p-3) - \frac{(p-3)(p-4)}{2} - 6)$-connected. Assume there exists a cutset $X \subseteq V(G)$ dividing the graph into at least two connected pieces with $X \leq k(p-3) - \frac{(p-3)(p-4)}{2} - 7$. Let u and v be two vertices such that u and v are in distinct components of $G - X$. There exists at least one element $z \in Z$ contained in $G - X$, and so it follows that $V(D) \setminus X$ is contained in a single component of $G - X$. Given that there exist $k(p-3) - \frac{(p-3)(p-4)}{2} - 6$ internally disjoint paths from each of v and u to $V(D)$, it follows that $V(D) \setminus X$, u, and v are all contained in the same component of $G - X$, contrary to our choice of u and v.

We now observe that there is no vertex set X of order at most n in G such that $G - X$ does not contain a K_p-minor. By Observation 8.3, for any vertex set X of order n, the graph $G - X$ has face-width at least r. It follows that $G - X$ contains K_p as a minor by Theorem 8.4.

As a final step, we now prove that G cannot contain k disjoint instances of K_p-minors

when $k \geq p$. Suppose, to reach a contradiction, that G contains pairwise disjoint subgraphs $H_1, \ldots, H_k$, each of which contains K_p as a minor. Recall that Z is the set of vertices adjacent every vertex of D, and $|Z| = k(p-3) - \frac{(p-3)(p-4)}{2} - 6$. We are now interested in all of the instances $H_1, \ldots, H_k$ that contain at most $p-4$ vertices of Z. We fix the value t, and possibly re-number the subgraphs H_i for $1 \leq i \leq k$ such that H_i contains at most $p-4$ vertices of Z if and only if $1 \leq i \leq t$. We let l_i be defined to be $|V(H_i) \cap Z|$ for $1 \leq i \leq t$. It follows immediately that:

$$
\begin{aligned}
|Z| &= \sum_1^t l_i + \sum_{i=t+1}^k |V(H_i) \cap Z| \\
&\geq \sum_1^t l_i + (k-t)(p-3) \\
&\geq \sum_1^t l_i + k(p-3) - tp + 3t
\end{aligned}
$$

If we combine the resulting inequality with the bound on $|Z|$, we conclude that

$$
\sum_1^t (p - l_i - 3) \geq \frac{(p-3)(p-4)}{2} + 6.
$$

We now define a new graph $\overline{G}$ to be the graph G' embedded in Σ with an additional vertex x attached to every vertex of D. It is clear that the graph $\overline{G}$ embeds in Σ as well. We also define $\overline{H}_i$ for $1 \leq i \leq t$ to be the subgraph of $\overline{G}$ formed by $H_i \cap G'$ and the vertex x. Observe that $\overline{H}_i$ contains a $p - l_i + 1$ clique minor. This follows as every branch set of H_i which does not intersect Z remains a connected branch set of $\overline{H}_i$, and we form one additional branch set consisting of the union of the remaining branch sets of the clique minor in H_i along with the vertex x. Note that by our choice of H_i, $p - l_i + 1 \geq 5$ for $1 \leq i \leq t$.

We now apply Lemma 8.2 to the subgraphs $\overline{H}_i$, $1 \leq i \leq t$, of the graph $\overline{G}$. It follows that:

$$
\begin{aligned}
\frac{(p-3)(p-4)}{6} + 1 &\geq \sum_1^t \left\lceil \frac{(p - l_i - 2)(p - l_i - 3)}{6} \right\rceil \\
&\geq \sum_1^t \left\lceil \frac{1}{3}(p - l_i - 3) \right\rceil \\
&\geq \sum_1^t \frac{1}{3}(p - l_i - 3).
\end{aligned}
$$

However, given our lower bound on $\sum_1^t (p - l_i - 3)$, we now arrive at a contradiction.

This completes the proof that there exists a graph $G_{n,k,p}$ which is $(k(p-3) + \frac{(p-3)(p-4)}{2} - 6)$-connected graph such that for all integers $n \geq 1$, $p \geq 5$, $k \geq p$, the graph $G_{n,k,p}$ does not contain k disjoint instances of K_p as a minor, nor does it contain a subset X of vertices with $|X| \leq n$ such that $G(n,k,p) - X$ does not contain K_p as a minor.

References

[1] T. Böhme, K. Kawarabayashi, J. Maharry and B. Mohar, Linear connectivity forces large complete bipartite graph minors, *J. Combin. Theory Ser. B*, **99** (2009), 557–582.

[2] K. Kawarabayashi, S. Norine, R. Thomas, P. Wollan, K_6 minors in 6-connected graphs of bounded tree width, *submitted*.

[3] K. Kawarabayashi, S. Norine, R. Thomas, P. Wollan, K_6 minors in large 6-connected graphs, *submitted*.

[4] R. Diestel, K. Kawarabayashi, Th. Müller, P. Wollan, On the excluded-minor structure theorem for graphs of large tree-width, `http://arxiv.org/abs/0910.0946`.

[5] R. Diestel, *Graph Theory* (3rd edition), Springer-Verlag, 2005,
Electronic edition available at: `http://diestel-graph-theory.com/GrTh.html`.

[6] P. Erdős and L. Pósa, On the maximal number of disjoint circuits of a graph, *Publ. Math. Debrecen* **9** (1962), 3–12.

[7] J.-O. Fröhlich and T. Müller, Linear Connectivity Forces Large Complete Bipartite Minors: the Patch for the Large Tree-Width Case, `http://arxiv.org/abs/0906.2568`.

[8] B. Mohar, *personal communication*.

[9] B. Mohar and C. Thomassen, Graphs on Surfaces, Johns Hopkins University Press, Baltimore, MD, 2001.

[10] N. Robertson and P. Seymour, Graph Minors V: Excluding a planar graph, *J. Combin. Theory, Ser. B* **41**, (1986) pp.92 - 114.

[11] N. Robertson and P. D. Seymour, Graph Minors. VII. Disjoint paths on a surface, *J. Combin. Theory, Ser. B*, **45** (1988), 212–254.

[12] N. Robertson and P. D. Seymour, Graph Minors. XVI. Excluding a non-planar graph *J. Combin. Theory, Ser. B* **89** (2003) 43-76.

[13] N. Robertson and P. Seymour, Graph Minors. XVII. Taming a vortex, *J. Combin. Theory, Ser. B* **77**, (1999) pp. 162 - 210.

[14] N. Robertson and P. Seymour, Graph Minors. XXI. Graphs with unique linkages, *J. Combin. Theory, Ser. B*, **99**, 583–616.

[15] R. Thomas and P. Wollan, An improved linear edge bound for graph linkages, *Europ. J. Combinatorics,* **26** (2005) 309 - 324.