Erdős-Pósa property for labelled minors: 2-connected minors

Henning Bruhn* Felix Joos† Oliver Schaudt

Abstract

In the 1960s, Erdős and Pósa proved that there is a packing-covering duality for cycles in graphs. As part of the graph minor project, Robertson and Seymour greatly extended this: there is such a duality for H-expansions in graphs if and only if H is a planar graph (this includes the previous result for $H = K_3$). We consider vertex labelled graphs and minors and provide such a characterisation for 2-connected labelled graphs H.

1 Introduction

The most satisfactory optimisation results are arguably the ones that also provide a certificate that the optimum is attained. An example is Menger's theorem stating that the maximum number of disjoint paths between two vertex sets is achieved if there is a separator of the same size. More generally this is captured by the min-flow/max-cut theorem or by the duality principle of linear programming.

Not always, however, concise certificates for optimality are known or do even exist. In such a case, an approximate certificate may be available. There are a few classic examples for this. One is the triangle removal lemma due to Ruzsa and Szemerédi [14] (for every $\epsilon \in (0,1)$, there is a $\delta > 0$ such that every graph on n vertices contains either δn^3 triangles or ϵn^2 edges whose deleting makes the graph triangle-free) and its generalisations. The importance of removal lemmas is for example demonstrated by its various applications in number theory, discrete geometry, graph theory and computer science [2].

Another example is a theorem due to Erdős and Pósa [4], which also (including its generalisations) has several applications in graph theory and computer science: every graph G that does not contain k disjoint cycles, admits a vertex set of size $O(k \log k)$ that meets every cycle. More generally, we say that a family of graphs $\mathcal{H}$ has the *Erdős-Pósa property* if there exists a function $f : \mathbb{N} \to \mathbb{R}_+$ such that for every graph G and every integer k, there exist k disjoint subgraphs in G that are isomorphic to graphs in $\mathcal{H}$, or G contains a vertex set X of size $|X| \leq f(k)$ such that every subgraph of G isomorphic to a graph in $\mathcal{H}$ meets X. Thus, the class of cycles has the Erdős-Pósa property.

The Erdős-Pósa property has been investigated for numerous graph classes (see [10] for a recent survey). One of the most striking results is the following due

*Partially supported by DFG, grant no. BR 5449/1-1

†Partially supported by DFG, grant no. JO 1457/1-1

to Robertson and Seymour that is a by-product of their graph minor project. It provides another characterisation of planar graphs, which in fact does not involve any topological arguments. (Essentially, a graph is an *H-expansion* if it can be turned into H by a series of edge contractions; see the next section for a formal definition.)

Theorem 1 (Robertson and Seymour [11])**.** *Let H be a graph. The family of H-expansions has the Erdős-Pósa property if and only if H is planar.*

Observe that this includes the class of cycles (set $H = K_3$).

There are further extensions of the theorem of Erdős and Pósa. Suppose we specify a set of labelled vertices S in a graph G and now we ask for cycles that contain at least one vertex from S (such cycles are also known as S-cycles). Kakimura, Kawarabayashi and Marx [6] proved that S-cycles also have the Erdős-Pósa property (see [1, 9] for further extensions). Clearly, this is a generalisation because we may that $S = V(G)$. In [5], Huynh, Joos and Wollan extended this to cycles with two labels.

We characterise all labelled 2-connected graphs H such that the class of labelled H-expansions has the Erdős-Pósa property. For simplicity, let us assume for now that every vertex has at most one label and we define a (sub)graph to be *simply-labelled* if all vertices with a label have the same one.

Theorem 2. *Let H be a labelled 2-connected graph such that each vertex carries at most one label. Then the labelled H-expansions have the Erdős-Pósa property if and only if there is an embedding of H in the plane such that the boundary C of the outer face contains all labelled vertices, and there are two simply-labelled subpaths $P, Q \subseteq C$ that cover all of $V(C)$.*

We have actually not yet specified what a labelled H-expansion is. There are several ways to define labelled expansions. We choose a definition such that the resulting labelled minor relation is transitive and we also generalise the results about labelled cycles. As the precise definition is a bit technical, we defer it to Section 2. We note that, with a slightly stronger notion of labelled subdivisions, Liu [7] proved a half-integral Erdős-Pósa type result for labelled subdivisions.

Theorem 2 has a number of applications. It implies the result of Kakimura, Kawarabayashi and Marx that S-cycles have the Erdős-Pósa property as well as the result due to Huynh, Joos and Wollan that the same is true for cycles with two labels. Moreover, more complicated variants of cycles with labelled vertices are covered. For instance, the theorem shows that, given a set S, the family of cycles that each contain, say, at least 42 vertices from S has the Erdős-Pósa property. Instead of S-cycles, we could also consider S-K_4-subdivisions, that is, subdivisions of K_4 that each contain at least one vertex from S. As a consequence of our theorem, the set of these has the Erdős-Pósa property, too. Similar statements involving two labels are also covered.

Our main theorem requires the graph H to be 2-connected. This is necessary: if H is not 2-connected then the conclusion of the theorem becomes false; in particular, there are simply-labelled graphs H such that all labelled vertices belong to the boundary of a single face but H-expansions do not have the Erdős-Pósa property. We investigate the Erdős-Pósa property for unconnected and merely 1-connected graphs H in a follow-up paper in which we heavily rely on the results of this paper.

2 Labelled graphs and minors

In this section we introduce several definitions concerning labelled graphs, minors, expansions, walls, and tangles. All definitions not involving labels are standard and commonly used in the literature. Most of our notation is standard and in accordance with Diestel [3].

We start with expansions and minors without labels. For a graph H, a pair (X, π) of a graph X and a mapping $\pi : V(H) \cup E(H) \to V(X) \cup E(X)$ is an *H-expansion* if

(i) $\{\pi(u)\}_{u \in V(H)}$ is a partition of $V(X)$ into vertex-disjoint induced subgraphs of X such that $\pi(u)$ is a tree for all $u \in V(H)$; and

(ii) for every two distinct $u, v \in V(H)$ if u and v are adjacent in H there is exactly one $\pi(u)$–$\pi(v)$ edge in X, the edge $\pi(uv)$, and if u and v are not adjacent there is no such edge.

Often we omit π and simply say that X is an H-expansion. If a graph G contains an H-expansion X as a subgraph, we say that H is a *minor* of G. Note that for every vertex u of H the induced subgraph $X[\pi(u)]$ together with all edges $\pi(uv)$ for $v \in N_H(u)$ forms a tree, which we denote by T^π_u. We refer to $\pi(u)$ as the *branch set* of u.

2.1 Labelled graphs

Let us now formally introduce labelled graphs and labelled expansions. We call a graph G a *labelled graph* if some of its vertices are marked with one or more labels from some alphabet Σ. Formally, G is endowed with a function $\ell : V(G) \to \mathcal{P}(\Sigma)$, and we say that a vertex v *is labelled with* $\alpha \in \Sigma$ if $\alpha \in \ell(v)$. Note that a vertex may have several labels or none at all. We also write that a graph G is *Σ-labelled*.

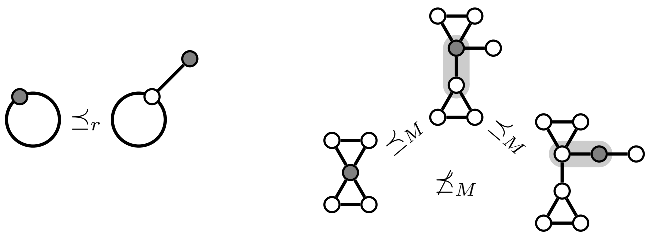

Figure 1: Labelled graphs with two different minor relations. Labelled vertices in grey. Left: an S-cycle as a rooted minor. Right: transitivity fails for naive labelled minor relation

What should it mean that some (labelled) graph has some other graph H as a labelled minor, or equivalently, contains a labelled H-expansion? A natural labelled minor relation has been explored before: Wollan [15] and Marx, Seymour and Wollan [8] treat *rooted minors*, minors with a single label. In this setting, a vertex in a minor is labelled as soon as its branch set contains a labelled vertex (a *root*). While this definition bears its own merit, it does not capture all structures we want to express. In particular, it does not capture S-cycles:

if a graph contains an S-cycle as a rooted minor, then it does not necessarily contain an S-cycle as a subgraph; see Figure 1. Our notion of a labelled minor will be designed to capture S-cycles, as well as *long* S-cycles, S-cycles of length at least a fixed length ℓ. These are known to have the Erdős-Pósa property [1].

The problem with rooted minors, at least in view of S-cycles, is that a branch set may send out an appendix to pick up a labelled vertex, where this appendix is unnecessary for the (unlabelled) minor relation. At first sight, the following variant of the definition fixes this issue: say a vertex v in a minor is labelled as soon as its branch set contains a labelled vertex and that labelled vertex lies on a path between two edges in the expansion that connect that branch set to the branch sets of other vertices. This definition, however, leads to a labelled minor relation that is not transitive, which is clearly problematic (see Fig. 1) also because labelled H-expansions do not necessarily contain a labelled H'-expansion for all subgraphs H' of H. Our notion of a labelled minor is slightly different but transitive and hence also closed under taking subgraphs.

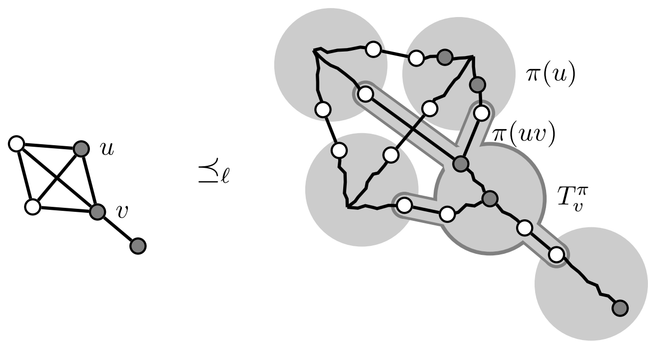

Figure 2: A labelled expansion (labelled vertices in grey)

Fix some alphabet Σ and let H be a Σ-labelled graph. A pair (X, π) of a labelled graph X and a mapping $\pi : V(H) \cup E(H) \to V(X) \cup E(X)$ is a *labelled H-expansion* if

(i) (X, π) is an H-expansion; and

(ii) if $v \in V(H)$ is labelled with α then every non-trivial, if T_v^π is not an isolated vertex, leaf-to-leaf path in T_v^π contains a vertex contained in $\pi(u)$ that is labelled with α.

Observe that T_v^π may only be an isolated vertex if v is an isolated vertex. Intuitively, the definition says that if u and v are neighbors of some vertex w in H, the direct path from $\pi(u)$ to $\pi(v)$ through $\pi(w)$ contains vertices of every label in $\ell(v)$.

Again, if the mapping π is clear from the context, we may simply call X itself a labelled H-expansion. If a labelled graph G contains a labelled H-expansion as a subgraph, H is a *labelled* minor of G. We write $H \preceq_\ell G$ for short.

Let us first convince ourselves that this definition yields a transitive minor relation. To this end, we say that a labelled H-expansion (X, π) is *minimal* if for all $u \in V(H)$ the following holds:

- If $d_H(u) \geq 2$, then every leaf of $\pi(v)$ is contained in some $\pi(uv)$ for some $v \in N_H(u)$; and

- if $d_H(u) \leq 1$, then $\pi(u)$ is a path and if $d_H(u) = 1$, then an endvertex of this path is contained in $\pi(uv)$ where v is the unique neighbour of u.

It is easy to see that every labelled H-expansion (X, π) contains a minimal H-expansion (X', π') such that $\pi'(uv) = \pi(uv)$ for all $uv \in E(H)$ and $\pi'(u)$ is a subtree of $\pi(u)$ for all $u \in V(H)$.

Lemma 3. *Let A, B, C be labelled graphs such that $A \preceq_\ell B$ and $B \preceq_\ell C$. Then also $A \preceq_\ell C$.*

Proof. Observe first that whenever a graph G contains a labelled H-expansion of a graph H, then G also contains a labelled H'-expansion for any subgraph H' of H.

Hence we may assume that (B, β) is a minimal labelled A-expansion, and that (C, γ) is a minimal labelled B-expansion. Define

$$\pi(a) = \bigcup_{b \in V(\beta(a))} \gamma(b) \cup \bigcup_{bb' \in E(\beta(a))} \gamma(bb')$$

for every $a \in V(A)$, and set $\pi(aa') = \gamma(\beta(aa'))$ for all $aa' \in E(A)$. Forgetting the labels, it is a standard task to check that (C, π) is an A-expansion. Thus, it remains to verify condition (ii) in the definition of labelled expansions.

For this, let $a \in V(A)$ be labelled with α, and let $P = u \ldots v$ be a leaf-to-leaf path in T_a^π. Since B and C are minimal, u (resp. v) either does not belong to $\pi(a)$ or $d_A(a) \leq 1$. Observe that $T_a^\pi = \bigcup_{b \in \beta(a)} T_b^\gamma$. Then P defines a leaf-to-leaf path P' in T_a^β (if $d_A(a) \leq 1$, then $P' = T_a^\beta$). The path P' contains a vertex $b^* \in \beta(a)$ that is labelled with α as (B, β) is a labelled A-expansion. The path $Q = P \cap T_{b^*}^\gamma$ is, in $T_{b^*}^\gamma$, a leaf-to-leaf path as well. Since b^* is labelled with α it follows that Q contains a vertex c^* in $\gamma(b^*)$ that is labelled with α as well. Since $c^* \in V(\gamma(b^*)) \subseteq V(\pi(a))$ we have found a vertex in $\pi(a)$ on P that is labelled with α, as desired. $\square$

The definition of a labelled graph or expansion allows for vertices to receive two or more labels, and this is necessary for the labelled minor relation to make sense. However, our main result, Theorem 2, requires the vertices in the graph H to have at most one label. This is mostly because we favour main theorems with simple statements. Allowing doubly-labelled vertices in H complicates matters somewhat. While we can (and will) handle these complications, the resulting statement becomes more complex, and less attractive (see Theorem 13).

3 Tangles

The concept of a tangle plays a key role in this paper. We start with the definition and explain how a minimal counterexample for the Erdős-Pósa property of a certain family of graphs naturally yields a tangle. We then introduce walls and recall how tangles are linked to walls. In Section 3.5, we introduce linkages and in Section 3.6, we state the key tool for our proof.

3.1 Definition

An ordered pair (A, B) of edge-disjoint subgraphs of G that partition $E(G)$ is a *separation*. The *order* of the separation is $|V(A) \cap V(B)|$.

A *tangle* of order r in a graph G is a set $\mathcal{T}$ of tuples (A, B) so that the following assertions hold.

(T1) Every tuple $(A, B) \in \mathcal{T}$ is a separation of order less than r.

(T2) For all separations (A, B) of G of order less than r, exactly one of (A, B) and (B, A) lies in $\mathcal{T}$.

(T3) $V(A) \neq V(G)$ for all $(A, B) \in \mathcal{T}$.

(T4) $A_1 \cup A_2 \cup A_3 \neq G$ for all $(A_1, B_1), (A_2, B_2), (A_3, B_3) \in \mathcal{T}$.

For any tangle $\mathcal{T}$ and $(A, B) \in \mathcal{T}$, we refer to A as the *$\mathcal{T}$-small* side of the separation (A, B). Suppose $\mathcal{T}$ has order $r \geq 3$ and let X be a vertex set of size at most $r - 2$. Then $G - X$ contains a unique block U such that $V(U) \cup X$ is not contained in any $\mathcal{T}$-small side of a separation in $\mathcal{T}$. We call the block U the *$\mathcal{T}$-large block* of $G - X$.

3.2 Tangles and the Erdős-Pósa property

The concept of tangles goes very well together with the Erdős-Pósa property. To see this we first introduce the notion of a minimal counterexample. Suppose $\mathcal{H}$ is a family of graphs, G is a graph, and $k \in \mathbb{N}$. We say that G is *$\mathcal{H}$-free* if no subgraph of G lies in $\mathcal{H}$. We say the pair (G, k) is a *minimal counterexample* to the function $f : \mathbb{N} \to \mathbb{R}_+$ being an Erdős-Pósa function for the family $\mathcal{H}$ if the following statements hold.

(MC1) The graph G does contain neither k disjoint copies of graphs in $\mathcal{H}$ nor a set $X \subseteq V(G)$ of size at most $f(k)$ such that $G - X$ is $\mathcal{H}$-free.

(MC2) For every $k' < k$, the graph G contains k' disjoint copies of graphs in $\mathcal{H}$ or a set $X \subseteq V(G)$ of size at most $f(k')$ such that $G - X$ is $\mathcal{H}$-free.

We extend this definition to the labelled case in a straightforward way: $\mathcal{H}$ is a labelled family of graphs, G is a labelled graph, and k is minimal such that G does contain neither k disjoint copies of labelled graphs in $\mathcal{H}$ nor a set $X \subseteq V(G)$ of size at most $f(k)$ such that $G - X$ is $\mathcal{H}$-free.

The following lemma shows that every minimal counterexample has a somewhat canonical tangle which indicates where the copies of the graphs in $\mathcal{H}$ lie. Essentially the same lemma was proven by Wollan in [16]. We include the short proof for completeness.

Lemma 4. *Suppose $\mathcal{H}$ is a family of connected labelled graphs. Suppose (G, k) is a minimal counterexample to the function $f : \mathbb{N} \to \mathbb{R}_+$ being an Erdős-Pósa function for $\mathcal{H}$. Suppose that $t \leq \min\{f(k) - 2f(k-1), f(k)/3\}$. Let $\mathcal{T}$ be the collection of all separations (A, B) of order less than t such that B contains a subgraph that lies in $\mathcal{H}$. Then $\mathcal{T}$ is a tangle.*

Proof. To verify that $\mathcal{T}$ is a tangle, we only need to check (T2)–(T4). Let (A, B) be a separation of G of order less than t. We claim that one of $A - B$ and $B - A$ contains a graph of $\mathcal{H}$. If not, set $X = V(A \cap B)$ and observe that $G - X$ is $\mathcal{H}$-free, which is impossible as $|X| < t < f(k)$.

Next, suppose that both A and B contain a copy of a graph in $\mathcal{H}$. Then, neither of $A - B$ and $B - A$ can contain $k - 1$ copies of graphs in $\mathcal{H}$ as (G, k)

is a counterexample. Hence there are a sets $X_A \subseteq V(A)$, $X_B \subseteq V(B)$, each of size at most $f(k-1)$, such that both $A - (V(B) \cup X_A)$ and $B - (V(A) \cup X_B)$ are $\mathcal{H}$-free. But then $G - (X_A \cup X_B \cup (V(A) \cap V(B))$ is $\mathcal{H}$-free (recall that the graphs in $\mathcal{H}$ are connected), which is impossible as

$$|X_A \cup X_B \cup (V(A) \cap V(B))| \leq 2f(k-1) + t \leq f(k).$$

Therefore, (T2) holds. For (T3), observe that $B - A = \emptyset$ if $V(A) = V(G)$, which clearly implies that $B - A$ cannot contain any graph from $\mathcal{H}$.

Finally, suppose there are three separations (A_1, B_1), (A_2, B_2), $(A_3, B_3) \in \mathcal{T}$ such that $A_1 \cup A_2 \cup A_3 = G$. Let $X = \bigcup_{i \in [3]} (V(A_i) \cap V(B_i))$, and observe that $|X| \leq 3t \leq f(k)$. Then, any graph in $\mathcal{H}$ that is disjoint from X must lie in $\bigcap_{i=1}^{3} B_i - A_i = \emptyset$. (Again, we use here that the graphs in $\mathcal{H}$ are connected.) Thus, $G - X$ is $\mathcal{H}$-free, which is again a contradiction. Therefore, (T4) holds and $\mathcal{T}$ is a tangle. $\square$

3.3 Walls

Let $[r]$ denote the set $\{1, \ldots, r\}$. The $r \times s$-*grid*, $r, s \geq 2$, is the graph on the vertex set $[r] \times [s]$ where a vertex (i, j) is adjacent to a vertex (i', j') if and only if $|i - i'| + |j - j'| = 1$. An *elementary r-wall* is the graph obtained from the $2(r+1) \times (r+1)$-grid by deleting all edges of the form $(2i-1, 2j-1)(2i-1, 2j)$, where $i \in [r+1]$ and $j \in [\lceil r/2 \rceil]$, and also all edges of the form $(2i, 2j)(2i, 2j+1)$, where $i \in [r+1]$ and $j \in [\lfloor (r-1)/2 \rfloor]$, and then deleting the two vertices of degree 1. An elementary 8-wall is depicted in Figure 3 (where we assume that first coordinate increases from left to right and the second coordinate increases from bottom to top).

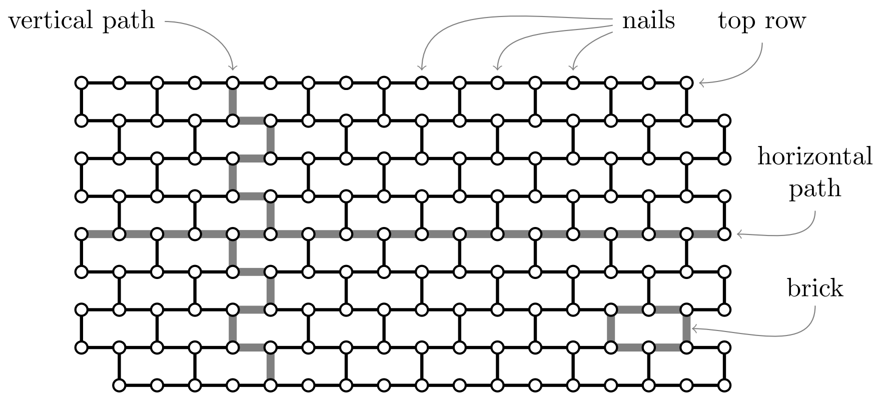

Figure 3: An elementary 8-wall

An *r-wall* or simply a *wall* is a subdivision W of an elementary r-wall Z. In Z we define the path $P^{(h)}_{j-1}$ for $j \in [r+1]$ as the path on vertices ij for $i \in [2(r+1)]$ (where we note that $P^{(h)}_0$ as well as $P^{(h)}_r$ are missing the first or last of these vertices as these are not present in Z.) The paths $P^{(h)}_0, \ldots, P^{(h)}_r$, which are pairwise disjoint, are the *horizontal paths* of Z. There are also $r+1$ pairwise disjoint $P^{(h)}_0$–$P^{(h)}_r$-paths in Z, the *vertical paths* $P^{(v)}_0, \ldots, P^{(v)}_r$ of Z.

The path $P_r^{(h)}$ is also called the *top row* of Z. The vertices of degree 2 in the top row are the *nails* of Z. Any 6-cycle in Z is a *brick* of Z.

We keep using the same concepts for walls as for elementary walls. That is, we will talk about vertical and horizontal paths of W, and mean the paths that arise from subdividing the corresponding paths in the elementary wall. A bit of care has to be applied when it comes to nails, as there are several choices of vertices in W that correspond to the (uniquely defined) nails in Z. But here, if necessary, we assume that the wall W comes with a fixed choice of nails, which allows us to speak about *the* nails of W.

Let $s \leq t$. An *s-subwall* W' of a t-wall W is subgraph of W that is an s-wall and such that every horizontal (vertical) path of W' is a subpath of a unique horizontal (vertical) path of W.

3.4 Tangles and Walls

We collect more facts about tangles and walls. For more details and proofs see Robertson and Seymour [11].

Let $\mathcal{T}$ be a tangle of order r, and let $s \leq r$. Let $\mathcal{T}'$ be the subset of those $(A, B) \in \mathcal{T}$ that are separations of order less than s. Then $\mathcal{T}'$ is again a tangle, the *truncation of $\mathcal{T}$ to order s*.

Let $\mathcal{T}$ be a tangle of order r in a graph H and assume that H is a minor of a graph G. We define a tangle $\mathcal{T}_H$ in G *induced by* H as follows. Let (C, D) be a separation in G of order less than r, and let C_H be the induced subgraph of H on all vertices whose branch set in G intersects C, and define D_H in the analogous way. Then every edge in H lies in C_H or in D_H as otherwise there would be an edge in G between $C - D$ and $D - C$. Moreover, since every branch set that meets C as well as D also contains a vertex in $C \cap D$, it follows that $|V(C_H \cap D_H)| \leq |V(C \cap D)|$. Thus, if we split up the common edges of C_H and D_H we obtain a separation (C_H, D_H) of H of order less than r. Therefore, either $(C_H, D_H) \in \mathcal{T}$ or $(D_H, C_H) \in \mathcal{T}$ and we then put (C, D) resp. (D, C) into $\mathcal{T}_H$. That $\mathcal{T}_H$ is indeed a tangle was shown by Robertson and Seymour [12].

Beside the tangle induced by the copies of a certain family $\mathcal{H}$ of graphs in a minimal counterexample for the Erdős-Pósa property, we consider two further tangles.

Lemma 5 (Robertson and Seymour [12]). *Suppose $n \geq 3$, $t = \lceil \frac{2n}{3} \rceil$, and $\mathcal{T}$ is the set of all $(t-1)$-separations (A, B) of K_n such that $V(B) = V(K_n)$. Then $\mathcal{T}$ is a tangle.*

For a K_t-expansion π, we refer to $\mathcal{T}_\pi$ as the tangle induced by the tangle in K_t that is described in Lemma 5.

Lemma 6 (Robertson and Seymour [12]). *Suppose $t \geq 2$ and W is a t-wall. Let $\mathcal{T}_W$ be the set of all t-separations (A, B) of W such that B contains an entire horizontal path. Then $\mathcal{T}_W$ is a tangle of order $t + 1$.*

We also need the converse direction, namely that a tangle of large order forces the existence of a large wall.

Theorem 7 (Robertson and Seymour [12]). *For every positive integer t, there is an integer $T(t)$ such that if G is a graph that has a tangle $\mathcal{T}$ of order $T(t)$, then there is a t-wall W in G such that $\mathcal{T}_W$ is a truncation of $\mathcal{T}$.*

3.5 Linkages

Let G be a graph, $k \in \mathbb{N}$, and let A, B be subgraphs, or vertex sets, of G. An *A–B-path* is a path from some $a \in A$ to some $b \in B$ that is internally disjoint from $A \cup B$. Moreover, an *A-path* is an A–A-path with at least one edge; if the path consist of a single edge, then this edge must not lie in A.

Let W be a wall with nails N. A *W-linkage* $\mathcal{L}$ of order k, or simply a *linkage*, is a set of k disjoint W-paths with first and last vertices in N. The top row of W defines a linear order $\leq$ (in fact two; we pick one) on the nails. Consider two paths P, Q in $\mathcal{L}$, and let the endvertices of P be $p_1 < p_2$, and let the endvertices of Q be $q_1 < q_2$. By symmetry, we may assume that $p_1 < q_1$. Then P and Q are *in series* if $p_2 < q_1$; they are *nested* if $p_1 < q_1 < q_2 < p_2$; and they are *crossing* if $p_1 < q_1 < p_2 < q_2$; see Figure 4. The linkage $\mathcal{L}$ is *in series*, *nested*, or *crossing* if all paths in $\mathcal{L}$ are mutually in series, nested, or crossing. We call $\mathcal{L}$ *pure* if it is in series, nested, or crossing.

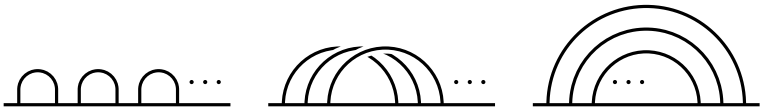

(a) in series (b) crossing (c) nested

Figure 4: The three types of pure linkages

Assume W to be contained in a Σ-labelled graph, and let $\alpha \in \Sigma$. A W-linkage $\mathcal{L}$ is called *α-clean*[1] if

- $\mathcal{L}$ is pure, and
- every path in $\mathcal{L}$ contains a vertex of label α.

Moreover, let $(\mathcal{P}, \mathcal{Q})$ be a partition of a W-linkage $\mathcal{P} \cup \mathcal{Q}$. We call $(\mathcal{P}, \mathcal{Q})$ a pair of *(α, β)-clean* W-linkages if

- $\mathcal{P}$ is α-clean and if $\mathcal{Q}$ is β-clean,
- $|\mathcal{P}| = |\mathcal{Q}|$, and
- for all $P, P' \in \mathcal{P}$ and $Q \in \mathcal{Q}$ with endvertices $p_1 < p_2$, $p'_1 < p'_2$ and $q_1 < q_2$, we have $q_1, q_2 \notin [p_1, p'_1] \cup [p_2, p'_2]$. Here, $[p_1, p'_1]$ is the set of all nails v with $p_1 \leq v \leq p'_1$, and $[p_2, p'_2]$ is defined similarly.

3.6 Flat walls

In their so-called *flat wall theorem* Robertson and Seymour [13] proved that every graph with a huge wall contains a large clique-minor or a large *flat* wall, a wall that lies in a *nearly* planar part of the graph. Huynh, Joos, and Wollan [5] extended the theorem to graphs whose edges are labelled with elements from two groups. We present below a version of the theorem that is adapted to labelled

[1] We adapt here a notion introduced by Huynh et al. [5] to the labelled setting. To keep notation simple, we have slightly weakened it.

graphs.[2] For our purposes it is not important that the wall is flat, so we simply drop the condition.

We need a little bit more notation before we can state our main tool, the result of Huynh et al. We define a sort of doubly-labelled expansion of a complete graph. For technical reasons, we weaken the definition of an expansion slightly. Let π be a mapping from $V(K_n) \cup E(K_n)$ into some graph, and let α, β be two labels. We say π is a (α, β)-*thoroughly labelled* (pseudo) K_n-expansion if

- $\pi(x)$ is a tree for every vertex x of K_n,
- $\pi(xy)$ is a set of at most two edges joining $\pi(x)$ and $\pi(y)$, and
- for every $\gamma \in \{\alpha, \beta\}$ and every triple x, y, z of vertices of K_n, there exist $e_{ab} \in \pi(ab)$ for each $ab \in \{xy, xz, yz\}$ such that $\pi(x) \cup \pi(y) \cup \pi(z) \cup e_{xy} \cup e_{xz} \cup e_{yz}$ contains a vertex with label γ.

Although, technically, these pseudo expansions are not expansions in the strict sense we defined earlier, we will simply call them (α, β)-thoroughly labelled K_n-expansions, which is already long enough.

For walls we have an analogous concept. A wall W is *thoroughly α-labelled* if every brick contains a vertex with label α, and the wall is *thoroughly (α, β)-labelled* if every brick contains a vertex with label α and a vertex with label β.

Theorem 8 (Huynh, Joos, and Wollan [5])**.** *For every $t \in \mathbb{N}$, there exists an integer t' such that if G is an (α, β)-labelled graph that contains a t'-wall W then one of the following statements holds.*

(i) *There is an (α, β)-thoroughly labelled K_t-expansion π in G such that $\mathcal{T}_\pi$ is a restriction of $\mathcal{T}_W$.*

(ii) *There is a $100t$-wall W_0 such that $\mathcal{T}_{W_0}$ is a restriction of $\mathcal{T}$ and*

  (a) *W_0 is (α, β)-thoroughly labelled,*

  (b) *for some $\gamma \in \{\alpha, \beta\}$, the wall W_0 is γ-thoroughly labelled and has an $(\{\alpha, \beta\} \setminus \gamma)$-clean W_0-linkage of size t, or*

  (c) *W_0 has a pair of (α, β)-clean W_0-linkages of size t.*

(iii) *For some $\gamma \in \{\alpha, \beta\}$, there is a set Z such that $|Z| < t'$ and the unique $\mathcal{T}_W$-large block of $G - Z$ does not contain any vertex labelled with γ.*

4 Necessity

In this section we show that all labelled graphs H such that the class of all H-expansion has the Erdős-Pósa property must have at least the properties stated in Theorem 2. We split the proof in several lemmas establishing gradually more properties of such H.

Lemma 9. *Let H be a labelled graph such that the labelled H-expansions have the Erdős-Pósa property. Then there is an embedding of H in the plane such that all its labelled vertices are on the boundary of the outer face.*

[2]For both groups we choose $(\mathbb{Z}, +)$. For an arbitrary ordering $e_1, e_2, \dots$ of the edges of G, we assign to e_i the group value 2^i in the j-th coordinate for $j \in [2]$ if one of the endpoints of e_i is labelled with the j-th labelled and otherwise 0.

Proof. First, we observe that we may assume H to be planar. Indeed, by Theorem 1, non-planar graphs do not enjoy the (ordinary) Erdős-Pósa property. Then, if we label every vertex in any graph G with all the labels of H, the labelled H-expansions do not have the Erdős-Pósa property for the same reasons as in the unlabelled case.

We thus assume that H is a planar labelled graph that, however, does not have any embedding in the plane such that all its labelled vertices are on the boundary of the outer face. Observe that, in particular, H must have a component with that property. Choose a minimum number ℓ such that there is an embedding of H in the plane in which the labelled vertices are contained in the union of ℓ face boundaries. By assumption, $\ell \geq 2$.

Let $R \in \mathbb{N}$ be sufficiently large, in a sense that will be made precise later in the proof. Moreover, let Σ be the alphabet containing all labels of H. Consider a plane $\ell R \times \ell R$-grid, and pick ℓ mutually disjoint cycles $C_1, \dots C_\ell$, each of length at least R (roughly $R/4 \times R/4$ squares), so that each has distance at least $R/4$ from the outer face and so that each two are at a distance of at least $R/4$ from each other. Let G be the graph obtained by deleting the vertices in the interior of each C_i, and labelling every vertex in $\bigcup_{i=1}^{\ell} V(C_i)$ with all labels in Σ.

In what follows we see that every labelled H-expansion separates the interiors of the cycles C_i from each other. Then it will be easy to deduce that there are no two disjoint labelled H-expansions. The fact that we choose R large enough ensures that every hitting set has to be large (as its size grows with R).

Since R is chosen to be large enough, G contains a labelled H-expansion. Indeed, for sufficiently large R the graph G contains an unlabelled H-expansion such that every labelled vertex of H maps to a branch set whose vertices of degree at least 3 are all contained in the same C_i. Such an unlabelled H-expansion is also a labelled H-expansion.

By increasing R, we can force the minimum size of a hitting set for labelled H-expansions to be arbitrarily large. Thus, to finish the proof it suffices to show that G does not contain any two disjoint labelled H-expansions.

Let H' be some labelled H-expansion in G. Denoting the interior faces of the cycles $C_1, \dots, C_\ell$ by $F_1, \dots, F_\ell$, we see that H' has a face $F'_i \supseteq F_i$ for each $i \in [\ell]$. The faces $F'_1, \dots, F'_\ell$ are pairwise distinct: as the face boundaries of $F'_1, \dots, F'_\ell$ contain all the labelled vertices of G in H', it follows from the minimality of ℓ that no two of these faces coincide.

Next, suppose there is a second labelled H-expansion H'' in G that is disjoint from H'. Again, the minimality of ℓ implies that H'' has a component that contains a vertex from C_1 as well as a vertex from C_2 (after relabelling $C_1, \dots, C_\ell$). In particular, H'' contains a path P that starts in a vertex of C_1 and ends in a vertex of C_2. Then, however, P starts in F'_1 or in its boundary, and ends in F'_2 or in its boundary. As $F'_1 \neq F'_2$ it follows that $P \subseteq H''$ meets H', which shows that H' and H'' are not disjoint. □

Recall that a labelled graph is simply-labelled if each labelled vertex only one label and all labelled vertices have the same label. We may use this notion for subgraphs of labelled graphs, too.

Let H be a labelled planar graph H that has an embedding in the plane in which all labelled vertices are on the boundary of the outer face. Define the *label homogeneity* of H as the smallest integer s such that for every sufficiently large integer n there is a labelling of the vertices in the top row P of an elementary n-wall W such that H is a labelled minor of W and such that there are s simply-labelled subpaths of P that cover all labelled vertices of P.

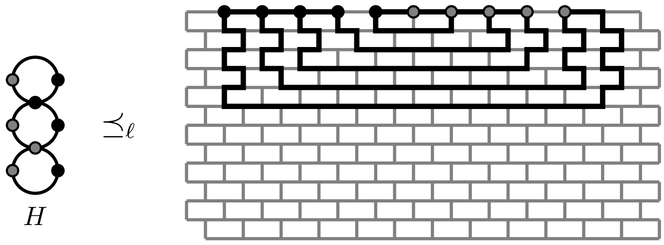

Figure 5: A graph H of label homogeneity 2

Lemma 10. *Let H be a connected labelled graph that has an embedding in the plane in which all labelled vertices are on the boundary of the outer face. If the labelled H-expansions have the Erdős-Pósa property, then H has label homogeneity at most 2.*

In the proof we will consider two grids, each on a vertex set indexed by a set $[n] \times [n]$, that is, on a vertex set $\{v_{ij} : (i,j) \in [n] \times [n]\}$. In both cases we assume that the vertices are chosen in such a way that v_{ij} is adjacent to $v_{i'j'}$ if and only if $i = i'$ and $|j - j'| = 1$, or if $|i - i'| = 1$ and $j = j'$. The vertices v_{jn} for $j \in [n]$ are the vertices of the *top row* of the grid.

Proof of Lemma 10. Suppose that H has label homogeneity $\ell \geq 3$. Then, there is a labelled $r \times r$ grid G' for some sufficiently large r such that G' contains a labelled H-expansion where all labelled vertices of G' are contained in the top row, and such that there are ℓ disjoint simply-labelled subpaths $P'_1, \dots, P'_\ell$ of the top row that cover all its vertices. Let $\{v_{ij} : i,j \in [r]\}$ be the vertex set of G'.

Suppose that f is an Erdős-Pósa function for labelled H-expansions. We enlarge G' to an $r' \times r'$-grid G for $r' = rs = r \cdot 3(f(2) + 1)$ with vertex set $\{w_{ij} : (i,j) \in [r'] \times [r']\}$. We say that w_{ij} has *pre-image* v_{pq} if $i - (p-1)s \in [s]$ and $j - (q-1)s \in [s]$. We label a vertex $w_{jr'}$ in the top row of G with label α if its pre-image v_{pr} is labelled with α in G'.

Let X be a set of at most $f(2)$ vertices in G. Let us convince ourselves that $G - X$ still contains a labelled H-expansion. For every $q \in [r]$, there is a $j \in [r']$ such that none of the $(j-1)$th, the jth or the $(j+1)$th column meets X, and such that $w_{j-1,r'}$, $w_{jr'}$ and $w_{j+1,r'}$ have v_{qr} as pre-image; let J be the set of these j, one for each $q \in [r]$. In a similar way, there are r rows of G, with index set I, that are disjoint from X. In particular, the union of the rows with index in I and the columns with index in J define a subgraph F of $G - X$ that contains a subdivision of an $r \times r$-grid. Let i_1 be the largest integer in I. We modify F by adding for every $j \in J$ the path $w_{j-1,r'} w_{jr'} w_{j+1,r'}$ together with the three vertical paths from these vertices to w_{j-1,i_1}, w_{j,i_1}, and w_{j+1,i_1} respectively. Call the obtained graph F' and observe that the labelled grid G' is a labelled minor of F'. Due to Lemma 3, H is a labelled minor of F'. Since F' is disjoint from X, we see that no set of at most $f(2)$ vertices meets every labelled H-expansion.

Therefore, G must contain two disjoint labelled H-expansions, H_1 and H_2 say. By construction, the vertices of the top row of G can be covered by ℓ disjoint simply-labelled paths $Q_1, \dots, Q_\ell$. By definition of the label homogeneity, each of the two the H-expansions needs to contain at least one vertex from each of the paths $Q_1, \dots, Q_\ell$. Let C_G be the boundary of the outer face of G.

Starting with the plane graph $C_G \cup H_1 \cup H_2$ we add a vertex x drawn in the outer face of C_G and make it adjacent to a vertex from each of Q_1, Q_2, Q_3. (Recall that $\ell \geq 3$.) The resulting graph K is planar. On the other hand, we see that K has a $K_{3,3}$-minor by contracting each of $Q_1, Q_2, Q_3, H_1 - (Q_1 \cup Q_2 \cup Q_3), H_2 - (Q_1 \cup Q_2 \cup Q_3)$ to a single vertex, a contradiction. This completes the proof. $\square$

In Lemma 12 we give a characterisation of the labelled graphs of label homogeneity at most 2. To simplify its proof we use the following definition together with Lemma 11. For a positive integer h we define a graph $W(h)$ as follows. Start with an elementary $2h^2$-wall W_1, and let $n_1, \dots, n_{2h^2}$ be the set of nails (in the order they appear in the top horizontal path). We add to W_1 a set of $2h$ further vertices $a_1, \dots, a_h$, $b_1, \dots, b_h$, and for each $i \in [h]$ we make a_i adjacent to each of $n_{(i-1)h+1}, \dots, n_{ih}$, while we make b_i adjacent to each of $n_{h^2+(i-1)h+1}, \dots, n_{h^2+ih}$. The graph $W(h)$ is (α, β)-*labelled* if each vertex $a_1, \dots, a_h$ is labelled with α and each of $b_1, \dots, b_h$ is labelled with β.

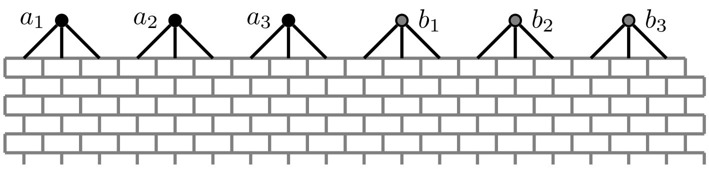

Figure 6: The graph $W(3)$

Lemma 11. *Let H be a labelled graph. Then H has label homogeneity at most 2 if and only if H is labelled with at most two labels, say α and β, and there is an h such that H is a labelled minor of the (α, β)-labelled graph $W(h)$.*

Proof. One direction is easy: if H has label homogeneity at most 2 then it must be labelled with at most two labels, α and β, say, and there is a t such that H is a labelled minor of the elementary $2t$-wall W' in which the first t nails are labelled with α and the other t nails with β. As obviously $W' \preceq_\ell W(t)$ it follows that also $H \preceq_\ell W(t)$.

For the other direction, let h be such that $H \preceq_\ell W(h)$. Let W_0 be an elementary $(2h^2+2)$-wall, and let $n_0, \dots, n_{2h^2+1}$ be its nails (in the order as they appear in the top row). Label the nails $n_0, \dots, n_{h^2}$ with α, and label the other nails with β. We claim that $W(h) \preceq_\ell W_0$.

To see this, denote by Q the top row of W_0, and denote by n_i^- the predecessor of n_i on Q for each i. We define branch sets A_j, B_j for $j \in [h]$ as follows. Set $A_j = n^-_{(j-1)h+1} Q n_{jh-1}$ and $B_j = n^-_{(j-1)h+h^2+1} Q n_{jh+h^2-1}$. Taking in W_0 the sets A_j, B_j as branch sets, as well as all the vertices in $W_0 - Q$ as singleton branch sets, we obtain a labelled $W(h)$-expansion, which means that $W(h)$, and thus also H, is a labelled minor of W_0. As the labels of W_0 can be covered by two simply-labelled subpaths of Q, it follows that H has label homogeneity at most 2. $\square$

Lemma 12. *Let H be a 2-connected graph. Then H has label homogeneity at most 2 if and only if there is an embedding of H in the plane such that the boundary C of the outer face contains all labelled vertices, and there are two internally disjoint subpaths $P, Q \subseteq C$ that together cover all of $V(C)$ and we can associate a label α with P and a label β with Q such that*

- *$P - Q$ is simply-labelled with α and $Q - P$ is simply-labelled with β; and*
- *for all $v \in V(P \cap Q)$, we have $d_H(v) = 2$ and v is labelled with $\{\alpha, \beta\}$.*

Proof. If H has an embedding in the plane as stated above, there is an h such that H is a labelled minor of $W(h)$. By Lemma 11, it follows that H has label homogeneity at most 2.

If, on the other hand, H has label homogeneity at most 2, then there is an h such that H is a labelled minor of an elementary $2h$-wall W, in which the first h nails are labelled with α and the other h nails are labelled with β. Let (H', π) be a minimal labelled H-expansion in W.

If either α or β are not used at vertices of H, the statement of the lemma clearly holds, as any labelled minor of W has all its labels on the boundary of the same face.

We may, therefore, assume that some vertex in H is labelled with α and some vertex is labelled with β. By contracting the branch sets of H', we obtain a planar embedding of H. Let C' be the boundary of the outer face of H', and let C be the boundary of the outer face of the embedding of H. Since H is 2-connected, C is a cycle and since H' is a minimal H-expansion, C' is also a cycle. In fact, C is obtained from C' by contracting all branch sets. As all labelled vertices of W are contained in the top row and C' is the boundary of the outer face, every labelled vertex of H' must be on C'. Hence H has an embedding in the plane such that the boundary C of the outer face contains all labelled vertices.

Consider the nails of W ordered from left to right, say $n_1, \ldots, n_{2h}$, and let n_i be the leftmost nail contained in C'. Following C' in clockwise fashion we obtain a sequence $(n_i = n_{i_1}, n_{i_2}, \ldots, n_{i_r})$ of all labelled vertices on C'. Due to planarity, we have that $i_j < i_{j+1}$ for each $j \in [r-1]$. By definition of W, there is some $j \in [r]$ such that n_{i_j} is the rightmost nail labelled α.

We observe that there are at most two vertices in H that are labelled with $\{\alpha, \beta\}$ because every branch set of such a vertex must contain either $\{n_{i_j}, n_{i_{j+1}}\}$ or $\{n_{i_r}, n_{i_1}\}$. Suppose $u \in V(H)$ is labelled with $\{\alpha, \beta\}$ and it contains both n_{i_j} and $n_{i_{j+1}}$ (the argument for n_{i_1} and n_{i_r} is similar). For a contradiction, assume that $d_H(u) \geq 3$. Note that T_u^π contains a vertex x of degree at least 3 on C'. Observe that x can be neither inside nor outside $n_{i_j} C' n_{i_{j+1}}$ as in both cases there is a leaf-to-leaf path in T_u^π that either contains no vertex labelled α or no vertex labelled β.

Now it is not hard to construct the paths P and Q as in the statement. $\square$

With Lemma 12 we can see that neither of the graphs in Figure 7 has label homogeneity at most 2, which in light of the other results in this section means that the expansions of neither of the graphs have the Erdős-Pósa property.

Figure 7: Two graphs of label homogeneity larger than 2

5 Erdős-Pósa property for 2-connected H

In this section, we prove a slightly stronger version of our main result, Theorem 2.

Theorem 13. *Let H be a labelled 2-connected graph. Then the labelled H-expansions have the Erdős-Pósa property if and only if there is an embedding of H in the plane such that the boundary C of the outer face contains all labelled vertices, and there are two internally disjoint subpaths $P, Q \subseteq C$ that cover all of $V(C)$ and we can associate a label α with P and a label β with Q such that*

- *$P - Q$ is simply-labelled with α and $Q - P$ is simply-labelled with β; and*
- *for all $v \in V(P \cap Q)$, we have $d_H(v) = 2$ and v is labelled with $\{\alpha, \beta\}$.*

Note that Theorem 13 clearly implies Theorem 2. The proof closely follows the different outcomes of Theorem 8.

For two labels α, β we write $K_n^{\alpha,\beta}$ for the complete graph on n vertices in which every vertex is labelled with $\{\alpha, \beta\}$.

Lemma 14. *Let $t \geq 3$, and let α, β be labels. Then every (α, β)-thoroughly labelled K_{6t^2-5t}-expansion contains a labelled $K_t^{\alpha,\beta}$-expansion.*

Proof. Let K be the complete graph on the vertex set

$$c^1, \ldots, c^t, v_1^{ij}, \ldots, v_6^{ij} \text{ for all distinct } i, j \in [t]$$

of $6t^2 - 5t$ distinct vertices. Let (X, π) be a (α, β)-thoroughly labelled K-expansion.

Consider arbitrary distinct indices $i, j \in [t]$. By definition, there is a cycle C in

$$\pi(v_1^{ij}) \cup \pi(v_1^{ij} v_2^{ij}) \cup \pi(v_2^{ij}) \cup \pi(v_2^{ij} v_3^{ij}) \cup \pi(v_3^{ij}) \cup \pi(v_3^{ij} v_1^{ij})$$

that contains a vertex with label α. By renaming the vertices $v_1^{ij}, v_2^{ij}, v_3^{ij}$ if necessary we may assume that there is a $\pi(v_1^{ij})$–$\pi(v_3^{ij})$ path P_1 in C that contains a vertex in $\pi(v_2^{ij})$ with label α. With an analogous argument, we may assume that there is a $\pi(v_4^{ij})$–$\pi(v_6^{ij})$ path P_2 contained in

$$\pi(v_4^{ij}) \cup \pi(v_4^{ij} v_5^{ij}) \cup \pi(v_5^{ij}) \cup \pi(v_5^{ij} v_6^{ij}) \cup \pi(v_6^{ij})$$

that contains a vertex in $\pi(v_5^{ij})$ of label β. Using an edge in $\pi(v_3^{ij} v_4^{ij})$, as well as an edge in $\pi(c^i v_1^{ij})$, we can find a $\pi(c^i)$–$\pi(v_6^{ij})$ path Q^{ij} that contains both a vertex with label α and a vertex with label β in its interior and that itself is contained in the induced graph on $\pi(c^i) \cup \bigcup_{\ell=1}^{6} \pi(v_\ell^{ij})$.

Having constructed all such paths Q^{ij}, let $ab \in \pi(v_6^{ij} v_6^{ji})$ with $a \in \pi(v_6^{ij})$. Set $\pi'(ij) = ab$, and define a tree T_i' by taking the union of all paths Q^{ij}, a–Q^{ij} paths in $\pi(v_6^{ij})$ together with a minimal subtree of $\pi(c^i)$ so as to result in a tree. Set $\pi'(i) = V(T_i')$, and observe that π' defines a K_t-expansion Y that is contained in X. In Y the trees $T_i^{\pi'}$ (recall the definition of a labelled expansion) consist of T_i' together with all edges $\pi'(ij)$. Every leaf-to-leaf path in $T_i^{\pi'}$ passes through $\pi(c^i)$ and then contains Q^{ij} and Q^{ik} for two j, k. Consequently, every leaf-to-leaf path contains a vertex with label α and a vertex with label β that lies in $\pi'(i)$. Therefore, (Y, π') is a labelled $K_t^{\alpha,\beta}$-expansion. $\square$

Lemma 15. *Let H be an (α, β)-labelled graph. For every k (and H), there is a t such that every (α, β)-thoroughly labelled K_t-expansion contains k disjoint labelled H-expansions.*

Proof. Set $h = k|V(H)|$, and observe that there are k disjoint labelled minors of H in $K_h^{\alpha,\beta}$. Set $t = 6h^2 - 5h$, and apply Lemma 14 in order to find $K_h^{\alpha,\beta}$ as a labelled minor in any (α, β)-thoroughly labelled K_t-expansion. $\square$

Lemma 16. *Suppose $t \geq 9r$. If W is a $(t + 1)$-wall with an α-clean linkage of size $2r$, there is a t-subwall W' of W with an α-clean linkage of size r that is in series. Moreover, if W has an (α, β)-clean pair of linkages of size $2r$, there is a t-subwall W' of W with an (α, β)-clean pair of linkages that are both in series and of size r.*

Proof. We prove the second statement since the first one follows in the same way. Let $(\mathcal{P}, \mathcal{Q})$ be an (α, β)-clean pair of linkages of size $2r$, and let R be the top row of W. For each nail u, let S_u be the path contained in W from u to the upper right corner and then to the lower right corner of its brick in W; see Figure 8. Let $\mathcal{P} = \{P_1, \ldots, P_{2r}\}$ be the paths in $\mathcal{P}$ and denote their left endvertices by $p_1, \ldots, p_{2r}$, and the corresponding right endvertices by $p_1', \ldots, p_{2r}'$. Assume $p_1 < \ldots < p_{2r}$ where the ordering is from left to right in the top row R of W. Moreover, let $\mathcal{Q} = \{Q_1, \ldots, Q_{2r}\}$ with left endvertices $q_1, \ldots, q_{2r}$ and right endvertices $q_1', \ldots, q_{2r}'$ be ordered in the same way.

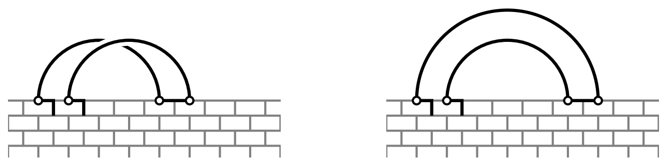

Figure 8: How to turn a crossing or nested linkage into one that is in series

Consider the paths

$$P_i' = S_{p_{2i-1}} p_{2i-1} P_{2i-1} p_{2i-1}' R p_{2i}' P_{2i} p_{2i} S_{p_{2i}}$$

for each $i \in [r]$ if $\mathcal{P}$ is crossing or nested (see Figure 8), otherwise let $P_i' = S_{p_{2i}} p_{2i} P_{2i} p_{2i}' S_{p_{2i}}$. We define Q_i' analogously. Note that for each $i, j \in [r]$ the paths P_i' and Q_j' are pairwise disjoint — this is due to the last condition in the definition of a pair of clean (α, β)-linkages. Let W' be the t-subwall obtained from W by deleting the top row and leftmost column, let $\mathcal{P}' = \{P_i' : i \in [r]\}$, and $\mathcal{Q}' = \{Q'_i : i \in [r]\}$. Note that the pair $(\mathcal{P}', \mathcal{Q}')$ is (α, β)-clean for W', which completes the proof. $\square$

Lemma 17. *Let $r \geq 4t$, and let W be a $100r$-wall that is either thoroughly (α, β)-labelled, or that is thoroughly α-labelled and has a β-clean linkage of size r. Then W contains a $100t$-subwall W' that has an (α, β)-clean pair of linkages of size t such that both linkages are in series.*

Proof. First, by Lemma 16, if W has a β-clean linkage of size $2r$ (rather than being thoroughly (α, β)-labelled) then it also has such a linkage of size $r \geq 2t$ that is in series — at the price of reducing the size of the wall by 1.

Pick a vertical path P of W such that each of the two components W_1, W_2 of $W - P$ contains at least $49r$ of the vertical paths of W. We may assume that if W has a β-clean linkage (which then is in series), then at least half of the paths of the linkage have both endvertices in W_1. That is, W_1 has a β-clean linkage of size t. Also, for each $i \in [2]$, let W'_i be obtained by W_i by deleting the first two horizontal paths.

Let $B_1, \ldots, B_t$ be a choice of r (vertex-)disjoint bricks from the top row of W_2. Let Q_3 be the third horizontal path of W_2 from the top; that is, the top path of W'_2. There are $2t$ disjoint Q_3–$\bigcup_{i=1}^{t} B_i$ paths $R_1, \ldots, R_{2t}$ such that R_{2i-1} and R_{2i} end in B_i for each $i \in [t]$. Since each brick of W_2 contains a vertex labelled with α as W is thoroughly α-labelled, for each i, one of the two paths in B_i between the endvertices of R_{2i-1} and R_{2i} contains a vertex of label α. Denote this subpath by S_i. Hence $(R_{2i-1} \cup S_i \cup R_{2i})_{i \in [t]}$ is an α-clean linkage in series of W'_2 of size t.

If W is thoroughly (α, β)-labelled we repeat this procedure in W_1 with the label β. If W has a β-clean linkage, then, by prolonging the linkage through the wall to W'_1, we obtain a β-clean linkage of W'_1 of size at least t. In both cases, by using the horizontal paths that link W'_1 and W'_2 we find a $100t$-wall W' as a subwall with an (α, β)-clean pair of linkages of size at least t. Moreover, the linkages are in series. $\square$

Note that if W is a wall with an (α, β)-clean pair of linkages which is in series, the union of these two linkages is itself a linkage that is in series. This follows from the definition of a (α, β)-clean pair. Hence, we may simply say that an (α, β)-clean pair of linkages is in series if both linkages are in series.

Lemma 18. *Let H be an (α, β)-labelled graph that has an embedding in the plane such that all labelled vertices lie in the boundary of the outer face, and assume H to have label homogeneity at most 2. For every k there is a t such that the following holds: whenever W is a wall of size at least $20t$ that has an in series (α, β)-clean pair of linkages of size t, then W together with the linkage contains k disjoint labelled H-expansions.*

Proof. For a positive integer t', let $U_{t'}$ be a labelled graph consisting of an elementary $2t'$-wall where the first t' nails are labelled α and the remaining ones are labelled β. As H has label homogeneity 2, there is a t' such that H is a labelled minor of $U_{t'}$. We now fix such a t' and simply write U instead of $U_{t'}$. We will find k disjoint labelled U-expansions, that then contain k disjoint labelled H-expansions.

We set $t = 5kt'$. Let $(\mathcal{P}, \mathcal{Q})$ be an (α, β)-clean pair of linkages of the wall W of size t that is in series. As $\mathcal{P}$ and $\mathcal{Q}$ are in series, all paths in $\mathcal{P}$ connect to the top row of W left of all paths in $\mathcal{Q}$ or vice versa. In particular, W, which has size at least $20t = 100kt'$, together with $\mathcal{P}$ and $\mathcal{Q}$ contains as a labelled minor a $10kt'$-wall W' in which the first $5kt'$ nails are labelled α and the remaining $5kt'$ nails are labelled β. We claim that

$$W' \textit{ contains } k \textit{ disjoint labelled } U\textit{-expansions.} \tag{1}$$

The claim is proved by induction on k. For $k = 1$, (1) holds as U is a labelled minor of W'. Suppose now that $k > 1$. Let W'' be a subwall of W' of size $10(k-1)t'$ that contains exactly $5(k-1)t'$ α-labelled and and exactly $5(k-1)t'$ β-labelled nails of W' such that the horizontal and vertical paths of W' that W'' meets are contiguous. By induction, W'' contains $k-1$ labelled U-expansions. The graph $\tilde{W} = W' - W''$ contains the leftmost and rightmost $5t' - 1$ vertical paths of W', as well as the $5t'$ bottommost horizontal paths. Then, $\tilde{W}$ contains U as a labelled minor. In total, we have found k disjoint labelled U-expansions. This proves (1) and the lemma. □

Lemma 19. *Let H be an (α, β)-labelled graph that has an embedding in the plane such that all labelled vertices lie in the boundary of the outer face, and assume H to have label homogeneity at most 2. For every k there is a t such that the following holds: if a graph G consists of a wall W of size at least $100t$ such that*

(a) *W is (α, β)-thoroughly labelled,*

(b) *for some $\gamma \in \{\alpha, \beta\}$, the wall W is γ-thoroughly labelled and has an $(\{\alpha, \beta\} \setminus \gamma)$-clean linkage of size t, or*

(c) *W has a pair of (α, β)-clean linkages of size t,*

then G contains k disjoint labelled H-expansions.

Proof. For a given k, let s be as the t in the statement of Lemma 18, and set $t = 4s$. Then, with Lemma 17, we may assume that W has size $100s$ and comes with a (α, β)-clean pair of linkages of size s that are in series. Lemma 18 now yields the k disjoint labelled H-expansions. □

We can now prove our main result.

Proof of Theorem 13. Necessity follows from Lemmas 9, 10 and 12.

It remains to prove sufficiency. For this, let H be a 2-connected (α, β)-labelled graph that has an embedding as in the statement. To proceed to the difficult case, we assume that H contains vertices of both labels, α and β. (If not, set $\alpha = \beta$.)

Suppose the theorem is false. Then there is a largest $k < \infty$ such that there are values $f(2), \ldots, f(k-1)$ such that for all $k' < k$ every graph G either contains k' disjoint labelled H-expansions or a vertex set X of size $|X| \le f(k')$ that meets every H-expansion.

Fix numbers $t_1 \gg t_2 \gg t_3 \gg k$, where we make precise what that means below. Moreover, choose $f(k)$ such that $t_1 \le \min\{f(k) - 2f(k-1), f(k)/3\}$ and complete f to a function $f : \mathbb{N} \to \mathbb{N}$.

By the choice of k we may pick a minimal counterexample (G, k) to f being an Erdős-Pósa function for the family of labelled H-expansions. Let $\mathcal{T}$ be the tangle as defined in Lemma 4 with t_1 playing the role of t which, by Lemma 4, has size at least t_1.

We assume t_1 to be chosen large enough such that Theorem 7 yields a t_2-wall W_1 whose induced tangle $\mathcal{T}_{W_1}$ is a truncation of $\mathcal{T}$. Next, we assume t_2 to be large enough such that t_2 and t_3 can play the roles of t' and t in Theorem 8.

We now go through the different outcomes of Theorem 8. For outcome (i), we apply Lemma 15, where we choose t_3 large enough to yield k disjoint H-expansions). For outcome (ii), we apply Lemma 19, where again we assume that t_3 is large enough to ensure k disjoint H-expansions.

Finally, we observe that the outcome (iii) may not occur. Indeed, recall that (iii) yields a label $\gamma \in \{\alpha, \beta\}$ and a set Z such that $|Z| < t_2$ and the unique $\mathcal{T}_{W_1}$-large block of $G - Z$ does not contain any vertex labelled with γ. Since H is 2-connected by assumption, any labelled H-expansion in $G - Z$ is edge-disjoint from the $\mathcal{T}_{W_1}$-large block of $G - Z$. As $|Z| < t_2 \leq f(k)$ and (G, k) is a counterexample, however, $G - Z$ must contain some labelled H-expansion. Thus, there is a separation (A, B) of order at most $|Z| + 1 < t_1$ such that B contains the unique $\mathcal{T}_{W_1}$-large block of $G - Z$. Hence, $(A, B) \in \mathcal{T}$, but A contains a labelled H-expansion (hence $(B, A) \in \mathcal{T}$), which is a contradiction to (T2). This completes the proof. $\square$

References

[1] H. Bruhn, F. Joos, and O. Schaudt, *Long cycles through prescribed vertices have the Erdős-Pósa property*, J. Graph Theory **87** (2018), 275–284.

[2] D. Conlon and J. Fox, *Graph removal lemmas*, Surveys in combinatorics 2013, London Math. Soc. Lecture Note Ser., vol. 409, Cambridge Univ. Press, Cambridge, 2013, pp. 1–49.

[3] R. Diestel, *Graph theory* (5th edition), Springer-Verlag, 2017.

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

[5] T. Huynh, F. Joos, and P. Wollan, *A unified Erdős-Pósa theorem for constrained cycles*, to appear in Combinatorica.

[6] N. Kakimura, K. Kawarabayashi, and D. Marx, *Packing cycles through prescribed vertices*, J. Combin. Theory Ser. B **101** (2011), 378–381.

[7] C.-H. Liu, *Packing topological minors half-integrally*, arXiv:1707.07221 (2017).

[8] D. Marx, P.D. Seymour, and P. Wollan, *Rooted grid minors*, J. Combin. Theory Ser. B **122** (2017), 428 – 437.

[9] M. Pontecorvi and P. Wollan, *Disjoint cycles intersecting a set of vertices*, J. Combin. Theory Ser. B **102** (2012), 1134–1141.

[10] J.-F. Raymond and D. M. Thilikos, *Recent techniques and results on the Erdős-Pósa property*, Discrete Appl. Math. **231** (2017), 25–43.

[11] N. Robertson and P. Seymour, *Graph minors. V. Excluding a planar graph*, J. Combin. Theory Ser. B **41** (1986), 92–114.

[12] N. Robertson and P.D. Seymour, *Graph minors. X. Obstructions to tree-decomposition*, J. Combin. Theory Ser. B **52** (1991), 153 – 190.

[13] N. Robertson and P.D. Seymour, *Graph minors. XIII. The disjoint paths problem*, J. Combin. Theory Ser. B **63** (1995), 65 – 110.

[14] I. Z. Ruzsa and E. Szemerédi, *Triple systems with no six points carrying three triangles*, Combinatorics (Proc. Fifth Hungarian Colloq., Keszthely, 1976), Vol. II, Colloq. Math. Soc. János Bolyai, vol. 18, North-Holland, Amsterdam-New York, 1978, pp. 939–945.

[15] P. Wollan, *Extremal functions for rooted minors*, J. Graph Theory **58** (2008), 159–178.

[16] P. Wollan, *Packing cycles with modularity constraints*, Combinatorica **31** (2011), 95–126.

Version May 1, 2018

Henning Bruhn `<henning.Bruhn@uni-ulm.de>`
Institut für Optimierung und Operations Research
Universität Ulm
Germany

Felix Joos `<f.joos@bham.ac.uk>`
School of Mathematics
University of Birmingham
United Kingdom

Oliver Schaudt `<schaudt@mathc.rwth-aachen.de>`
Department of Mathematics
RWTH Aachen University
Germany