An edge variant of the Erdős–Pósa property

Jean-Florent Raymond, Ignasi Sau, Dimitrios M. Thilikos

An edge variant of the Erdős–Pósa property☆

CrossMark

Jean-Florent Raymond [a,b,c,*], Ignasi Sau [b], Dimitrios M. Thilikos [b,d]

[a] Faculty of Mathematics, Informatics and Mechanics, University of Warsaw, Warsaw, Poland
[b] AlGCo Project Team, CNRS, LIRMM, Montpellier, France
[c] University of Montpellier, Montpellier, France
[d] Department of Mathematics, National and Kapodistrian University of Athens, Athens, Greece

ARTICLE INFO

ABSTRACT

For every $r \in \mathbb{N}$, we denote by θ_r the multigraph with two vertices and r parallel edges. Given a graph G, we say that a subgraph H of G is *a model of* θ_r *in* G if H contains θ_r as a contraction. We prove that the following edge variant of the Erdős–Pósa property holds for every $r \geqslant 2$: if G is a graph and k is a positive integer, then either G contains a packing of k mutually edge-disjoint models of θ_r, or it contains a set S of $f_r(k)$ edges such that $G \setminus S$ has no θ_r-model, for both $f_r(k) = O(k^2 r^3 \text{ polylog } kr)$ and $f_r(k) = O(k^4 r^2 \text{ polylog } kr)$.

1. Introduction

Typically, an Erdős–Pósa property reveals relations between covering and packing invariants in combinatorial structures. The origin of the study of such properties comes from the Erdős–Pósa Theorem [5], stating that there is a function $f : \mathbb{N} \to \mathbb{N}$ such that for every $k \in \mathbb{N}$ and for every graph G, either G contains k vertex-disjoint cycles, or there is a set X of $f(k)$ vertices in G meeting all cycles of G. In particular, Erdős and Pósa proved this result for $f(k) = O(k \cdot \log k)$.

An interesting line of research aims at extending Erdős–Pósa Theorem for packings and coverings of more general graph structures. In this direction, we say that a graph class $\mathcal{G}$ *satisfies the Erdős–Pósa property* if there exists a function $f_{\mathcal{G}} : \mathbb{N} \to \mathbb{N}$ such that, for every graph G and every positive integer k, either G contains k mutually vertex-disjoint subgraphs, each isomorphic to a graph in $\mathcal{G}$, or it contains a set S of $f_{\mathcal{G}}(k)$ vertices meeting every subgraph of G that is isomorphic to a graph in $\mathcal{G}$. When this property holds for a class $\mathcal{G}$, we call the function $f_{\mathcal{G}}$ *the gap of the Erdős–Pósa property for the class* $\mathcal{G}$. In this sense, the classic Erdős–Pósa Theorem says that the class containing all cycles satisfies the Erdős–Pósa property with gap $O(k \cdot \log k)$.

Given a graph J, we denote by $\mathcal{M}(J)$ the set of all graphs containing J as a contraction. Robertson and Seymour proved the following proposition, which in particular can be seen as an extension of the Erdős–Pósa Theorem.

Proposition 1. *Let J be a graph. The class $\mathcal{M}(J)$ satisfies the Erdős–Pósa property if and only if J is planar.*

Fig. 1. The graph θ_5.

A proof of Proposition 1 appeared for the first time in [18]. Another proof can be found in Diestel's monograph [4, Corollary 12.4.10 and Exercise 40 of Chapter 12]. In view of Proposition 1, it is natural to try to derive good estimations of the gap function $f_{\mathcal{M}(J)}$ in the case where J is a planar graph. In this direction, the recent breakthrough results of Chekuri and Chuzhoy imply that $f_{\mathcal{M}(J)}(k) = k \cdot \mathrm{polylog}\, k$ [2] when J is a planar graph and, even more, that $f_{\mathcal{M}(J)} = (k + |V(J)|)^{O(1)}$ [3]. Before this, the best known estimation of the gap for planar graphs was exponential, namely $f_{\mathcal{M}(J)}(k) = 2^{O(k \log k)}$, and could be deduced from [14] using the proof of [18]. Moreover, some improved polynomial gaps have been proven for particular instantiations of the graph J (see [6–9,15,16]). Another direction is to add restrictions on the graphs G that we consider, which usually allows to optimize the gap $f_{\mathcal{M}(J)}$. In this direction, it is known that $f_{\mathcal{M}(J)} = O(k)$ in the case where graphs are restricted to some non-trivial minor-closed class [10].

We consider the edge counterpart of the Erdős–Pósa property, where packings are edge-disjoint (instead of vertex-disjoint) and coverings contain edges instead of vertices. We say that a graph class $\mathcal{G}$ satisfies *the edge variant of the Erdős–Pósa property* if there exists a function $f_{\mathcal{G}}$ such that, for every graph G and every positive integer k, either G contains k mutually edge-disjoint subgraphs, each isomorphic to a graph in $\mathcal{G}$, or it contains a set X of $f_{\mathcal{G}}(k)$ edges meeting every subgraph of G that is isomorphic to a graph in $\mathcal{G}$. Recently, the edge variant of the Erdős–Pósa property was proved in [12] for 4-edge-connected graphs in the case where $\mathcal{G}$ contains all odd cycles.

In this paper we concentrate on the case where $\mathcal{G} = \mathcal{M}(J)$ for some graph J. We find it an interesting question whether an edge-analogue of Proposition 1 exists or not. To our knowledge, the only case for which $\mathcal{M}(J)$ satisfies the edge variant of the Erdős–Pósa property is when $J = K_3$, *i.e.* when the class of graphs $\mathcal{G}$ contains all cycles. This result is the edge-counterpart of the Erdős–Pósa Theorem and appears as a (hard) exercise in [4, Exercise 23 of Chapter 7]. For every $r \geqslant 2$, let θ_r be the graph containing two vertices and r multiple edges between them (see Fig. 1). The results of this paper can be stated as follows:

Theorem 1. *The edge variant of the Erdős–Pósa property holds for $\mathcal{M}(\theta_r)$ with gap $f_{\mathcal{M}(\theta_r)}$, with*

$$f_{\mathcal{M}(\theta_r)}(k) = O(k^2 r^3 \,\mathrm{polylog}\, kr) \quad \text{and} \quad f_{\mathcal{M}(\theta_r)}(k) = O(k^4 r^2 \,\mathrm{polylog}\, kr).$$

Theorem 1 is the edge-counterpart of the main result of [9]. The proof is presented in Section 3 and contains three main ingredients. The first is a reduction of the problem to graphs of bounded degree, presented in Section 3.1. The second is an application of recent results of [2] to obtain bounds on the treewidth of the graphs we consider (Section 3.2) and the last is an extension of the techniques in [10] fitting our needs, which is presented in Section 3.3. Section 2 contains definitions and preliminary results and Section 4 discusses further research about the problem investigated in this paper.

2. Definitions and preliminaries

For any graph G, $V(G)$ (respectively $E(G)$) denotes the *set of vertices* (respectively *edges*) of G. Even when dealing with multigraphs (*i.e.* graphs where more than one edge is allowed between two vertices) we will use the term *graph*. A graph G' is a *subgraph* of a graph G if $V(G') \subseteq V(G)$ and $E(G') \subseteq E(G)$, and we denote this by $G' \subseteq G$. If X is a subset of $V(G)$ (respectively $E(G)$), we denote by $G[X]$ the *subgraph of G induced by X*, *i.e.* the graph with vertex set X (respectively $\cup_{e \in X} e$) and edge set $\{\{x, y\} \in E(G), \ x \in X \text{ and } y \in X\}$ (respectively X). If S is a subset of vertices or edges of a graph G, the graph $G \setminus S$ is the graph obtained from G after the removal of the elements of S. For every vertex $v \in V(G)$ the neighborhood of v in G, denoted by $N_G(v)$, is the subset of vertices that are adjacent to v, and its size is called the *degree* of v in G, written $\deg_G(v)$. The maximum degree $\Delta(G)$ of a graph G is the maximum value taken by $\deg_G$ over $V(G)$. Given a non-negative integer k, a triple (V_1, S, V_2) is called a *k-separation triple* of a graph G if $|S| \leqslant k$ and $\{V_1, S, V_2\}$ is a partition of $V(G)$ such that there is no edge between a vertex of V_1 and a vertex of V_2. Unless otherwise stated, logarithms are binary. For any two integers a, b such that $a \leqslant b$, the notation $[\![a, b]\!]$ stands for the set of integers $\{a, a + 1, \dots, b\}$. In a tree T, rooted at a vertex $r \in V(T)$, a vertex $u \in V(T)$ is said to be a *descendant* of a vertex $v \neq u$ if the path in T from r to u contains v. The set of descendants of v is denoted by $\mathrm{desc}_T(v)$. A graph is *biconnected* if the removal of any vertex leaves the graph connected, and a biconnected component of a graph is a maximal biconnected subgraph.

Minors and models. In a graph G, a *contraction* of an edge $e = \{u, v\} \in E(G)$ is the operation that removes e from G and identifies the vertices u and v. In this paper, we keep multiple edges that may appear between two vertices after a contraction (for instance, contracting an edge in a triangle gives a graph with two vertices connected by two edges). For any graph J, let $\mathcal{M}(J)$ denote the class of *contraction models* (models for short) of J, *i.e.* the class of graphs that can be contracted to J. We say that a graph J is *minor* of a graph G (denoted by $J \leqslant_m G$) if a subgraph of G is a model of J (J-*model* for short), or, equivalently, if J can be obtained from G by a series of vertex deletions, edge deletions, and edge contractions.

Packings and coverings. Let G and J be graphs. We denote by $\mathbf{pack}_J^{\mathbf{v}}(G)$ the maximum number of vertex-disjoint models of J in G and by $\mathbf{cover}_J^{\mathbf{v}}(G)$ the minimum size of a subset $S \subseteq \mathrm{V}(G)$ (called J-*vertex-hitting set*) that meets the vertex sets of all models of J in G. These invariants are widely studied in the context of the classic Erdős–Pósa property.

Similarly, we write $\mathbf{pack}_J^{\mathbf{e}}(G)$ for the maximum number of edge-disjoint models of J in G and $\mathbf{cover}_J^{\mathbf{e}}(G)$ for the minimum size of a subset $S \subseteq \mathrm{E}(G)$ (called J-*edge-hitting set*) that meets the edge sets of all models of J in G. Obviously, for every two graphs G and J, the following inequality holds:

$$\mathbf{pack}_J^{\mathbf{e}}(G) \leqslant \mathbf{cover}_J^{\mathbf{e}}(G).$$

A graph J is said to satisfy the *(vertex-)Erdős–Pósa property for minors* (*vertex-Erdős–Pósa property* for short) if there is a function $f_J : \mathbb{N} \to \mathbb{N}$, called *vertex-Erdős–Pósa gap* of J, such that for every graph G, the following holds:

$$\mathbf{cover}_J^{\mathbf{v}}(G) \leqslant f_J(\mathbf{pack}_J^{\mathbf{v}}(G)).$$

The research of this paper is motivated by the course of detecting graphs J for which there is a function $h_J : \mathbb{N} \to \mathbb{N}$ satisfying the following inequality for every graph G:

$$\mathbf{cover}_J^{\mathbf{e}}(G) \leqslant h_J(\mathbf{pack}_J^{\mathbf{e}}(G)). \tag{1}$$

Such graphs are said to satisfy the *edge variant of the Erdős–Pósa property for minors* (or, in short, the *edge-Erdős–Pósa property*) and the function h_J is called the *gap* of the edge-Erdős–Pósa property for J (*edge-Erdős–Pósa gap* for short). This definition is an edge-counterpart to the existing Erdős–Pósa property and (vertex-)Erdős–Pósa gap.

Treewidth. A *tree decomposition* of a graph G is a pair $(T, \mathcal{V})$ where T is a tree and $\mathcal{V}$ a family $(V_t)_{t \in \mathrm{V}(T)}$ of subsets of $\mathrm{V}(G)$ (called *bags*) indexed by the vertices of T and such that

(i) $\bigcup_{t \in \mathrm{V}(T)} V_t = \mathrm{V}(G)$;
(ii) for every edge e of G there is an element of $\mathcal{V}$ containing both endpoints of e; and
(iii) for every $v \in \mathrm{V}(G)$, the subgraph of T induced by $\{t \in \mathrm{V}(T) \mid v \in V_t\}$ is connected.

The *width* of a tree decomposition T is defined as $\max_{t \in \mathrm{V}(T)} |V_t| - 1$ (that is, the maximum size of a bag minus one). The *treewidth* of G, written $\mathbf{tw}(G)$, is the minimum width of any of its tree decompositions.

A tree decomposition $(T, \mathcal{V})$ of a graph G is said to be a *nice* tree decomposition if

(i) every vertex of T has degree at most 3;
(ii) T is rooted at one of its vertices r whose bag is empty ($V_r = \emptyset$); and
(iii) every vertex t of T is
  • either a *base node*, i.e. a leaf of T whose bag is empty ($V_t = \emptyset$) and different from the root;
  • or an *introduce node*, i.e. a vertex with only one child t' such that $V_t = V_{t'} \cup \{u\}$ for some $u \in \mathrm{V}(G)$;
  • or a *forget node*, i.e. a vertex with only one child t' such that $V_{t'} = V_t \cup \{u\}$ for some $u \in \mathrm{V}(G)$;
  • or a *join node*, i.e. a vertex with two children t_1 and t_2 such that $V_t = V_{t_1} = V_{t_2}$.

It is known that every graph has an optimal tree decomposition which is nice [13].

The graph θ_r and the Erdős–Pósa property. The vertex-Erdős–Pósa property of θ_r received some attention, in particular in [7,9,11]. For instance the main result of [9] is the following estimation of the vertex-Erdős–Pósa gap for θ_r.

Proposition 2 ([9])**.** *For every positive integer r, θ_r has the vertex-Erdős–Pósa property with gap $O(k^2)$.*

However in this estimation the dependency in terms of r is hidden in the multiplicative constant of the Big-O notation. By a careful analysis of the size of a θ_r-hitting set presented in [9] (c.f. Lemma 6), the estimation of the gap of Proposition 2 can be made quadratic in both k and r. From this, we can derive an $O(k^3 r^3)$ edge-Erdős–Pósa gap for θ_r (Corollary 2) by using our Lemma 7 that makes possible to translate a θ_r-vertex-hitting set into a θ_r-edge-hitting set.

However, Theorem 1 gives better estimations of this gap, either in k or in r.

Patterns in graphs of big treewidth. In the following section, we will use several propositions asserting that every graph G of treewidth at least c_H contains some fixed graph H as a minor, where the constant c_H depends on H. For instance, we will show in Lemma 6 a simple relation between the constant $c_{k \cdot \theta_r}$ and the vertex-Erdős–Pósa gap for θ_r. These propositions are stated thereafter.

Proposition 3 ([17, Lemma 3.2])**.** *For every integer $r \geqslant 1$ and graph G, if $\mathbf{tw}(G) \geqslant 2r - 1$ then G contains a θ_r-model.*

Proposition 4 ([9, Lemma 1], see also [1])**.** *Let k and r be two positive integers. For every graph G, if $\mathbf{tw}(G) \geqslant 2k^2 r^2$ then G contains at least k vertex-disjoint models of θ_r.*

Proposition 5 ([2, Theorem 1.1])**.** *There is a function $f_{\text{Proposition 5}}(t) = O(\mathrm{polylog}\, t)$ such that, for every graph G and every positive integers h and p, if $hp^2 \leqslant \frac{\mathbf{tw}(G)}{f_{\text{Proposition 5}}(\mathbf{tw}(G))}$, there is a partition $G_1, \ldots, G_h$ of G into vertex-disjoint subgraphs such that $\mathbf{tw}(G_i) \geqslant p$ for each $i \in [\![1, h]\!]$.*

Proposition 6 ([2, Theorem 1.2])**.** *There is a function $f_{\text{Proposition 6}}(t) = O(\mathrm{polylog}\, t)$ such that, for every graph G and every positive integers h and p, if $h^3 p \leqslant \frac{\mathbf{tw}(G)}{f_{\text{Proposition 6}}(\mathbf{tw}(G))}$ then there is a partition $G_1, \ldots, G_h$ of G into vertex-disjoint subgraphs such that $\mathbf{tw}(G_i) \geqslant p$ for each $i \in [\![1, h]\!]$.*

Fig. 2. A marked tree T with $\mu(T) = 5$ and a 3-good partition (T_1, T_2) of root v. Marked vertices appear in white.

3. The edge-Erdős–Pósa property for graphs θ_r

3.1. Bounding the degree

In the sequel, we deal with graphs in which some vertices are marked. If G is a graph and $m\colon V(G) \to \{0, 1\}$ is a function, we say that (G, m) is a graph *marked* by m. A vertex v of G such that $m(v) = 1$ is said to be *marked*. We denote by μ the function that, given a graph, returns its number of marked vertices. We now define an r-good partition. Given a positive integer r, a marked tree (T, m) is said to have an *r-good partition of root v* if there is a pair $((T_1, m_1), (T_2, m_2))$ of marked trees such that:

(i) T_1 and T_2 are subtrees of T such that $(E(T_1), E(T_2))$ is a partition of $E(T)$;
(ii) $r \leqslant \mu((T_1, m_1)) \leqslant 2r$;
(iii) $v \in V(T_2)$; and
(iv) every vertex that is marked in (T, m) is either marked in (T_1, m_1) or marked in (T_2, m_2), but not in both. In other words, for every $u \in V(T)$,
  - if $v \in V(T_1) \cap V(T_2)$ then $m(v) = 1 \Leftrightarrow m_1(v) = 1$ or $m_2(v) = 1$ but not both;
  - otherwise, let $i \in \{1, 2\}$ be the integer such that $v \in V(T_i)$. Then we have $m(v) = m_i(v)$.

We remark that because of (iv), $\mu(T) = \mu(T_1) + \mu(T_2)$. If for every $v \in V(T)$, (T, m) has an r-good partition of root v, then T is said to have an *r-good partition*. Examples of a marked tree and of a good partition are given in Fig. 2.

Lemma 1. *For every integer $r > 0$ and every marked tree (T, m), if $\mu(T) \geqslant 2r$ then (T, m) has an r-good partition.*

Proof. Let $r > 0$ be an integer. We prove this lemma by induction on the size of the tree.

Base case: $|V(T)| = 0$. Since $2r \geqslant 2 > |V(T)|$, T does not have $2r$ marked vertices and we are done.

Induction step: Assume that for every integer $n' < n$, every marked tree (T', m') on n' vertices and satisfying $\mu((T', m')) \geqslant 2r$ has an r-good partition (induction hypothesis).

Let us prove that every marked tree on n vertices has a r-good partition if it has at least $2r$ marked vertices. Let (T, m) be a tree on n vertices and let v be a vertex of T. We assume that $\mu((T, m)) \geqslant 2r$. We distinguish two cases.

- $\mu((T, m)) = 2r$:

  Let $T_1 = T$, let $m_1 = m$, let $T_2 = (\{v\}, \emptyset)$, and let $m_2\colon V(T_2) \to \{0, 1\}$ be the function equal to 0 on every vertex of T_2. Remark that $(E(T_1), E(T_2)) = (E(T), \emptyset)$ is a partition of $E(T)$, T_2 contains v, and as (T, m) contains (exactly) $2r$ marked vertices, so does (T_1, m). Consequently $((T_1, m_1), (T_2, m_2))$ is an r-good partition of (T, m).

- $\mu((T, m)) > 2r$:

  We distinguish different cases depending on the degree of the root v in T.

  Case 1: $\deg(v) = 1$.

  Let u be the neighbor of v in T, let $T' = T \setminus \{v\}$, and $m' = m_{|V(T')}$. Remark that $\mu((T', m')) \geqslant 2r$ and $|V(T')| = |V(T)| - 1$. By induction hypothesis, (T', m') has an r-good partition $((T_1', m_1'), (T_2', m_2'))$ of root u. We extend it to T by setting $T_1 = T_1'$, $m_1 = m_1'$, $T_2 = (V(T_2') \cup \{v\}, E(T_2') \cup \{v, u\})$, and $m_2 = m_2'$. Notice that T_2 contains v. As the subtree T_1' contains at least r and at most $2r$ marked vertices (induction hypothesis), so does T_1. Also, remark that $(E(T_1), E(T_2))$ is a partition of $E(T)$ and that since $u \in T_2'$, the graph T_2 is connected. Therefore $((T_1, m_1), (T_2, m_2))$ is an r-good partition of T.

  Case 2: $\deg(v) = d > 1$.

  Let $u_1, \ldots, u_d$ be the neighbors of v in T and for every $i \in [\![1, d]\!]$, let C_i be the connected component of $T \setminus \{v\}$ that contains u_i. We also define, for every $i \in [\![1, d]\!]$, the restricted marking function $w_i = m_{|V(C_i)}$.

Subcase (a): there exists $i \in [\![1, d]\!]$ such that $\mu((C_i, w_i)) > 2r$.

Let $T' = (V(C_i) \cup \{v\}, E(C_i) \cup \{u, v\})$ and let $m' = m_{|V(T')}$. Remark that $|V(T')| < |V(T)|$ and $\mu((T', m')) > 2r$. According to the induction hypothesis, (T', m') has an r-good partition $((T'_1, m'_1), (T'_2, m'_2))$ of root v. Similarly as before, we can extend it into an r-good partition $((T_1, m_1), (T_2, m_2))$ of (T, m). This is done by setting:

$$T_1 = T'_1,$$
$$m_1 = m'_1,$$
$$T_2 = (V(T'_2) \cup (V(G) \setminus V(C_i)), E(G) \setminus E(T'_1)), \quad \text{and}$$
$$m_2 : \begin{cases} u \mapsto m'_2(u) & \text{if } u \in V(C_i) \cup \{v\} \\ u \mapsto m(u) & \text{otherwise.} \end{cases}$$

As $((T'_1, m'_1), (T'_2, m'_2))$ is an r-good partition of root v, $v \in V(T'_2)$ and therefore T_2 is connected.

Subcase (b): there exists $i \in [\![1, d]\!]$ such that $r \leqslant \mu((C_i, w_i)) \leqslant 2r$.

Let $T_1 = C_i$ and $T_2 = T[E(T) \setminus E(T_1)]$. In this case, $(E(T_1), E(T_2))$ is a partition of $E(T)$ and T_2 is connected since it contains v, the vertex which is adjacent to the C_j's. Thus, if we set $m_1 = m_{|V(T_1)}$ and $m_2 = m_{|V(T_2)}$, $((T_1, m_1), (T_2, m_2))$ is an r-good partition of (T, m).

Subcase (c): for all $i \in [\![1, d]\!]$, $\mu((C_i, w_i)) < r$.

Let $j = \min \left\{ j \in [\![2, d]\!], \ \sum_{i=1}^{j} \mu((C_i, w_i)) \geqslant r \right\}$. Such value exists since $\mu((T, m)) > 2r$. We then set:

$$T_1 = (\cup_{i \in [\![1, j]\!]} V(C_i) \cup \{v\}, \cup_{i \in [\![1, j]\!]}(E(C_i) \cup \{v, u_i\})),$$
$$m_1 : \begin{cases} v \mapsto 0 \\ u \in V(T_1) \setminus \{v\} \mapsto m(u), \end{cases}$$
$$T_2 = T[E(T) \setminus E(T_1)], \quad \text{and}$$
$$m_2 = m_{|V(T_2)}.$$

By definition of j, $\mu((T_1, m_1)) \geqslant r$ and as for every $i \in [\![1, d]\!]$, $\mu((C_i, w_i)) < r$ we also have $\mu((T_1, m_1)) < 2r$. As before, the pair $((T_1, m_1), (T_2, m_2))$ is an r-good partition of (T, m).

In conclusion, we proved by induction that for every integer r, every tree having at least $2r$ marked vertices has an r-good partition. $\square$

In the sequel we will deal with packings of the graph θ_r, for $r > 1$. The following remark is important.

Remark 1. If G is not biconnected, the number of edge-disjoint models of θ_r in G is equal to the sum of the number of edge-disjoint models of θ_r in every biconnected component of G. This enables us to treat biconnected components separately.

Lemma 2. Let $k > 0, r > 0$ be two integers, and let G be a biconnected graph with $\Delta(G) \geqslant 2kr$. Then $\mathbf{pack}^{\mathbf{e}}_{\theta_r}(G) \geqslant k$.

Proof. As G is biconnected, the removal of a vertex v of maximum degree gives a connected graph. Let T be a minimal tree of $G \setminus \{v\}$ spanning the neighborhood $N_G(v)$ of v. We mark the vertices of T that are elements of $N_G(v)$: this gives the marking function m for T. Let us prove by induction on k that (T, m) has k edge-disjoint marked subtrees $(T_1, m_1), \ldots, (T_k, m_k)$, each containing at least r marked vertices. If we do so, then we are done because $\{\{v\}, T_i\}_{i \in [\![1, k]\!]}$ is a collection of k edge-disjoint θ_r models. In fact, as for every $i \in [\![1, k]\!]$, T_i contains $r' \geqslant r$ vertices adjacent to v in G, contracting the edges of T_i in $G[\{v\} \cup V(T_i)]$ gives the graph $\theta_{r'}$. Let $r > 0$ be an integer.

Base case $k = 1$: Clear.

Induction step $k > 1$: Assume that for every $k' < k$, every tree with at least $2k'r$ vertices marked has k' edge-disjoint subtrees, each with at least r marked vertices. Let (T, m) be a marked tree such that $\mu((T, m)) \geqslant 2kr$. According to Lemma 1, (T, m) has an r-good partition $((T_1, m_1), (T'_1, m'_1))$ such that $r \leqslant \mu((T_1, m_1)) \leqslant 2r$ and $\mu((T'_1, m'_1)) = \mu((T, m)) - \mu((T_1, m_1)) \geqslant 2(k-1)r$. By induction hypothesis, (T'_1, m'_1) has $k - 1$ edge-disjoint marked subtrees $(T_2, m_2), \ldots, (T_k, m_k)$ each containing at least r marked vertices. Remark that as all these trees are subgraphs of T'_1, which is edge-disjoint from T_1 in T, they are edge-disjoint from T_1 as well. Consequently, $(T_1, m_1), (T_2, m_2), \ldots, (T_k, m_k)$ is the family of edge-disjoint subtrees we were looking for. $\square$

3.2. Bounding the treewidth

Lemma 3. There is a function $h_r(k) = O(kr^2 \text{ polylog } kr)$ such that for every positive integers k and r and every graph G, if $\mathbf{tw}(G) \geqslant h_r(k)$, then $\mathbf{pack}^{\mathbf{v}}_{\theta_r}(G) \geqslant k$.

Proof. Let G be a graph and k, r be two positive integers. By Proposition 4, if $\mathbf{tw}(G) \geqslant 2k^2r^2$, then G contains k vertex-disjoint models of θ_r. Therefore, we only have to consider the case where $\mathbf{tw}(G) < 2k^2r^2$.

As $f_{\text{Proposition 5}}(t) = O(\text{polylog } t)$ (cf. Proposition 5 for the definition of $f_{\text{Proposition 5}}$), there are three positive reals $t_0, A \geqslant 1$, and $\alpha \geqslant 1$ such that for every real $t \geqslant t_0$ we have $f_{\text{Proposition 5}}(t) \leqslant A \log^\alpha(t)$. Let $B = \max(0, \max_{i \in [\![1, \lceil t_0 \rceil]\!]} f_{\text{Proposition 5}}(i))$ and observe that for every positive integer i we have $f_{\text{Proposition 5}}(i) \leqslant A \log^\alpha(i) + B$.

Let $h_r(k) = k(2r)^2 \cdot (A \log^\alpha(2k^2r^2) + B)$ for every positive integers k and r. Observe that $h_r(k) = O(kr^2 \text{ polylog } kr)$. We will show that graphs whose treewidth is at least $h_r(k)$ contain k vertex-disjoint models of θ_r. For every positive integers r and k, if $\mathbf{tw}(G) \geqslant h_r(k)$ then we have

$$\mathbf{tw}(G) \geqslant k(2r)^2 \cdot (A \log^\alpha(\mathbf{tw}(G)) + B) \qquad \text{(as we assume } \mathbf{tw}(G) < 2k^2r^2)$$

$$\frac{\mathbf{tw}(G)}{A \log^\alpha(\mathbf{tw}(G)) + B} \geqslant k(2r)^2 \qquad \text{(because } A \log^\alpha(\mathbf{tw}(G)) + B \text{ is positive)}$$

$$\frac{\mathbf{tw}(G)}{f_{\text{Proposition 5}}(\mathbf{tw}(G))} \geqslant k(2r)^2 \qquad \text{(as } \mathbf{tw}(G) \text{ is integer).}$$

Notice that k and $2r$ meet the conditions of Proposition 5. Consequently, there is a partition $G_1, \ldots, G_k$ of G into vertex-disjoint subgraphs such that $\forall i \in [\![1, k]\!]$, $\mathbf{tw}(G_i) \geqslant 2r$. By Proposition 3, each of these subgraphs contains a model of θ_r. Consequently, G contains k vertex-disjoint models of θ_r, as required. $\square$

A very similar proof can be used to show the following lemma, using Proposition 6.

Lemma 4. *There is a function $h_r(k) = O(k^3r \text{ polylog } kr)$ such that, for every positive integers k and r and every graph G, if $\mathbf{tw}(G) \geqslant h_r(k)$, then $\mathbf{pack}_{\theta_r}^{\mathbf{v}}(G) \geqslant k$.*

3.3. From vertices to edges

In this section, we show how an estimation of a vertex-Erdős–Pósa gap can be derived from the bound on the treewidth obtained in Section 3.2. The proof of the two following lemmas are inspired from the proof of [10, Lemma 2].

Lemma 5 (*Adapted from Lemma 2 of [10]*). *Let $k \geqslant 3$, r be two positive integers and G a graph such that $\mathbf{pack}_{\theta_r}^{\mathbf{v}}(G) = k$. Then G has a $(\mathbf{tw}(G) + 1)$-separation triple (V_1, S, V_2) such that $\frac{1}{3}k \leqslant \mathbf{pack}_{\theta_r}^{\mathbf{v}}(G[V_1]) \leqslant \frac{2}{3}k$.*

Proof. Let $(T, \mathcal{V})$ be an optimal nice tree decomposition of G. For all $t \in V(T)$, let H_t be the subset of $V(G)$ equal to $\left(\bigcup_{t' \in \text{desc}_T(t)} V_{t'} \right) \setminus V_t$, that is, informally, the subset of vertices that are in bags *below* V_t but not in V_t. We also define the function $p \colon V(T) \to \mathbb{N}$ as: $\forall t \in V(T)$, $p(t) = \mathbf{pack}_{\theta_r}^{\mathbf{v}}(G[H_t])$, which counts the number of vertex-disjoint models of θ_r in the subgraph of G induced by H_t.

Remark 2. The function p is nondecreasing along every path from a vertex of T to the root of T, because if a vertex $t' \in V(T)$ is a child of a vertex $t \in V(T)$, then $H_{t'} \subseteq H_t$, and thus $\mathbf{pack}_H^{\mathbf{v}}(G[H_{t'}]) \leqslant \mathbf{pack}_H^{\mathbf{v}}(G[H_t])$.

Remark 3. As T is a nice decomposition of G, its vertices can be of four different types. We make remarks about the value taken by p depending on the type of the vertices:

Base node t: $p(t) = 0$, because since t has no descendant, $H_t = \emptyset$;

Introduce node t **with child** t': as the unique element of $V_t \setminus V_{t'}$ cannot appear in the bags of $\text{desc}_T(t')$ (by definition of a tree decomposition), $H_t = H_{t'}$ and then $p(t) = p(t')$;

Forget node t **with child** t': H_t contains one vertex more than $H_{t'}$ therefore $p(t) - p(t') \in \{0, 1\}$;

Join node t **with children** t_1 and t_2: $H_t = H_{t_1} \cup H_{t_2}$, but H_{t_1} and H_{t_2} are disjoint and there is no edge between the vertices of H_{t_1} and of H_{t_2} in $G[H_t]$ (otherwise the set $V_{t_1} = V_{t_2}$ would contain an endpoint of this edge, which also belongs to H_{t_1} or H_{t_2}, and this is contradictory). Thus there is no θ_r-model in $G[H_t]$ that uses (simultaneously) vertices of H_{t_1} and of H_{t_2}, and therefore $p(t) = p(t_1) + p(t_2)$.

Let $t \in V(T)$ be a node such that $p(t) > \frac{2}{3}k$ and such that for every child t' of t, $p(t') \leqslant \frac{2}{3}k$. Let us make some claims about t.

Claim 1. *such a t exists.*

Proof of Claim 1. The value of p on the root r of T is k (because $G[H_r] = G$) and according to the previous remark, the value of p on base nodes is 0. As p is nondecreasing on a path from a base node to the root (see Remark 2), such a vertex t exists. $\diamond$

Claim 2. *t is unique.*

Proof of Claim 2. To show that t is unique, we assume by contradiction that there is another $t' \in V(T)$ with $t' \neq t$ and $p(t') > \frac{2}{3}k$, and such that for every child t'' of t', $p(t'') \leqslant \frac{2}{3}k$. Three cases can occur:

(i) t' is a descendant of t. However, p is nondecreasing along any path from a vertex to the root (Remark 2) and $p(t') \geqslant \frac{2}{3}k$, whereas the value of p for each child of t is at most $\frac{2}{3}k$: this is a contradiction.

(ii) t is a descendant of t'. The same argument applies (symmetric situation).

(iii) t and t' are not in the above situations. Let v be the least common ancestor of t and t'. As p is nondecreasing along any path from a vertex to the root, the child v_t (respectively $v_{t'}$) of v of which t (respectively t') is descendant should be such that $p(v_t) > \frac{2}{3}k$ (respectively $p(v_{t'}) > \frac{2}{3}k$). By definition of v, we have $v_t \neq v_{t'}$. As v is a join node, $p(v) = p(v_t) + p(v_{t'}) > \frac{4}{3}k$, which is impossible. ◇

Claim 3. *t is either a forget node or a join node.*

Proof of Claim 3. By definition of t, the value $p(t)$ is different from the value(s) taken by p over the child(ren) of t. This can only occur in the cases of a join node or a forget node. ◇

We now present a $(\mathbf{tw}(G) + 1)$-separation triple (V_1, S, V_2) of G with the required properties.

Case 1: t is a forget node with t' as child.

Let $S = V_{t'}$, $V_1 = H_{t'}$, and $V_2 = V(G) \setminus (V_1 \cup S)$.

Case 2: t is a join node with t_1, t_2 as children.

As $\frac{2}{3}k < p(t) = p(t_1) + p(t_2)$ (Remark 3), there is $i \in \{1, 2\}$ such that $p(t_i) \geqslant \frac{k}{3}$. Let $S = V_{t_i}$, $V_1 = H_{t_i}$, and $V_2 = V(G) \setminus (V_1 \cup S)$.

In both cases, we have:

(i) $\frac{1}{3}k \leqslant \mathbf{pack}_{\theta_r}^{\mathbf{v}}(G[V_1]) \leqslant \frac{2}{3}k$ by definition of V_1 and t;

(ii) (V_1, S, V_2) is a partition of $V(G)$;

(iii) there is no edge between a vertex in V_1 and a vertex of V_2 (intuitively, S separates V_1 and V_2); and

(iv) $|S| \leqslant \mathbf{tw}(G) + 1$, because S is a bag of an optimal tree decomposition of G.

In the case where t is a forget node, the inequality $\frac{1}{3}k \leqslant \mathbf{pack}_{\theta_r}^{\mathbf{v}}(G[V_1])$ of (i) holds because $p(t') \geqslant p(t) - 1 > \frac{2}{3}k - 1 \geqslant \frac{k}{3}$ (*cf.* Remark 3). To see why (iii) is true, assume by contradiction that there are two vertices $u \in V_1$ and $v \in V_2$ such that $\{u, v\} \in E(G)$. Let $s_0 \in V(T)$ be the child of t such that $S = V_{s_0}$ (*cf.* the two different cases above). By definition of V_1 there is a vertex $s_1 \in V(T)$ of T in $\mathrm{desc}_T(s_0)$ whose bag V_{s_1} contains u. By definition of V_2, the vertex v does not belong to the bag V_{s_0} nor to a bag of a descendant of s_0. Let s_2 be a vertex of T containing u and which is, according to the previous remark, not the bag of a descendant of s_0 nor s_0.

As $(T, \mathcal{V})$ is a tree decomposition of G and $\{u, v\} \in E(G)$, we have the following:

- there is a vertex $s \in V(T)$ whose bag contains both u and v;
- the subgraph of T induced by vertices whose bags contain u (respectively v) is connected.

Consequently there is a path in T from s_1 to s (respectively from s_2 to s) each bag of which contains u (respectively v). As s is on the (only) path of T linking s_1 to s_2, one of u, v belongs to the bag V_s. But this contradicts the fact that (V_1, S, V_2) is a partition of $V(G)$.

We conclude that (V_1, S, V_2) is a $(\mathbf{tw}(G) + 1)$-separation triple of G with the required properties. □

A function $h: \mathbb{R} \to \mathbb{R}$ is said to be *superadditive* if for every x and every y in its domain, $f(x) + f(y) \leqslant f(x + y)$.

Lemma 6 (*Adapted from Lemma 5.4 in [2]*). *Let h_r be a superadditive function such that for every graph G and every positive integers r and k, if $\mathbf{tw}(G) \geqslant h_r(k)$ then $G \geqslant_{\mathrm{m}} (k + 1) \cdot \theta_r$. For every graph G and every positive integer k, if $\mathbf{pack}_{\theta_r}^{\mathbf{v}}(G) = k$ then we have*

$$\mathbf{cover}_{\theta_r}^{\mathbf{v}}(G) \leqslant 3 \cdot h_r(k) \log(k + 1).$$

Proof. We proceed by induction on k.

Base case $k = 0$: Clear.

Induction step $k > 0$: We assume that the lemma holds for every positive integer $k' < k$. Let G be a graph such that $\mathbf{pack}_{\theta_r}^{\mathbf{v}}(G) = k$. First, remark that $\mathbf{tw}(G) < h_r(k)$, otherwise by definition of h_r we would have $\mathbf{pack}_{\theta_r}^{\mathbf{v}}(G) > k$. Thus, by Lemma 5 G contains a $h_r(k)$-separation triple (V_1, S, V_2) such that $k/3 \leqslant \mathbf{pack}_{\theta_r}^{\mathbf{v}}(G[V_1]) \leqslant 2k/3$. This implies that $k_1, k_2 \leqslant \lfloor 2k/3 \rfloor$, where $k_i = \mathbf{pack}_{\theta_r}^{\mathbf{v}}(G[V_i])$ for every $i \in \{1, 2\}$. Also we have $k_1 + k_2 \leqslant k$ as G_1 and G_2 are two vertex-disjoint subgraphs of G.

The triple (V_1, S, V_2) is a partition of $V(G)$, so the following holds:

$$\textbf{cover}^{\textbf{v}}_{\theta_r}(G) \leqslant \textbf{cover}^{\textbf{v}}_{\theta_r}(G[V_1]) + \textbf{cover}^{\textbf{v}}_{\theta_r}(G[V_2]) + |S|$$

$$\leqslant \textbf{cover}^{\textbf{v}}_{\theta_r}(G[V_1]) + \textbf{cover}^{\textbf{v}}_{\theta_r}(G[V_2]) + h_r(k)$$

$$\leqslant 3 \cdot h_r(k_1)\log(k_1 + 1) + 3 \cdot h_r(k_2)\log(k_2 + 1) + h_r(k) \quad \text{(induction hyp.)}.$$

If $k = 1$, then $k_1 = k_1 = 0$ and we have $\textbf{cover}^{\textbf{v}}_{\theta_r}(G) \leqslant h_r(k) \leqslant 3 \cdot h_r(k)\log(k + 1)$. We may now assume $k \geqslant 2$. Observe that in this case, as $k_i \leqslant \lfloor \frac{2}{3}k \rfloor$, we get $k_i + 1 \leqslant \frac{3}{4}(k + 1)$ for every $i \in \{1, 2\}$.

$$\textbf{cover}^{\textbf{v}}_{\theta_r}(G) \leqslant 3 \cdot (h_r(k_1) + h_r(k_2))\log\left(\frac{3(k + 1)}{4}\right) + h_r(k)$$

$$\leqslant 3 \cdot h_r(k)\log\left(\frac{3(k + 1)}{4}\right) + h_r(k) \quad \text{(superadditivity of } h_r)$$

$$\leqslant 3 \cdot h_r(k)\log(k + 1) - 3 \cdot \log(4/3)h_r(k) + h_r(k)$$

$$\leqslant 3 \cdot h_r(k)\log(k + 1).$$

This concludes the proof. $\quad \square$

Corollary 1. *Let f_r be the vertex-Erdős–Pósa gap of θ_r. Then we have*

- $f_r(k) = O(kr^2 \text{ polylog } kr)$;
- $f_r(k) = O(k^3 r \text{ polylog } kr)$.

These estimations follow from Lemmas 3, 4 and 6.

The following lemma shows how to translate a vertex-Erdős–Pósa gap into an edge-Erdős–Pósa gap in the case of θ_r. The main idea of the proof is that if the considered graph has small maximum degree, a small edge-hitting set can be constructed from a small vertex-hitting set. On the other hand, a big maximum degree forces a large packing of θ_r-models.

Lemma 7. *If f_r is the vertex-Erdős–Pósa gap of θ_r, then the edge-Erdős–Pósa gap of θ_r is less than $2kr \cdot f_r(k)$.*

Proof. Let G be a graph, let $r \geqslant 2$ be an integer and let f_r is the vertex-Erdős–Pósa gap of θ_r.

We want to prove that if G contains less than k edge-disjoint models of θ_r, then it has a θ_r-edge-hitting set of size less than $2kr \cdot f_r(k)$.

According to Remark 1, we can assume that G is biconnected. If it is not the case, we consider its biconnected components separately (if it has no biconnected component then the lemma is trivial).

First, remark $\Delta(G) < 2kr$, otherwise by Lemma 2 G would contain at least k edge-disjoint θ_r-models.

Notice that if G does not contain k edge-disjoint θ_r-models, it does not contain k vertex-disjoint θ_r-models either. Consequently, there is a set $X \subseteq V(G)$ meeting every θ_r model of G and such that $|X| \leqslant f_r(k)$. Let us consider the set $Y \subseteq E(G)$ of edges incident to vertices of X, i.e. $Y = \{\{u, v\} \in E(G), \ u \in X\}$. Remark that as $\Delta(G) < 2kr$, we have $|Y| \leqslant 2kr \cdot f_r(k)$. Now, assume that there is a θ_r-model in G not having edges in Y. None of its vertices is in X, which is contradictory. So Y is a θ_r-edge hitting set of the required size. This concludes the proof. $\quad \square$

Corollary 2. *An edge-gap of $O(k^3 r^3)$ for θ_r can be derived from Proposition 2.*

Proof of Theorem 1. It follows from the application of Lemma 7 to the estimations of the vertex-Erdős–Pósa gap of θ_r given in Corollary 1.

4. Further research

The main question, initiated in this paper, is whether for every planar graph J, the class $\mathcal{M}(J)$ satisfies this edge variant of the Erdős–Pósa property. As for the vertex version, it is easy to see that the planarity of J is necessary. For instance, if $J = K_5$, consider as graph G an n-vertex toroidal wall, which is a 3-regular graph embeddable in the torus that contains K_5 as a minor. One can check that G does not contain two edge-disjoint models of K_5, but $\Omega(\sqrt{n})$ edges of G are needed in order to hit all its K_5-models.

Moreover, a second question is: when this property holds, does it hold with a polynomial gap for all graphs? Also, finding lower bounds on this gap for specific graphs is another interesting and complementary question. Let us mention that, as it is the case for the vertex version (see [5,8]), for any non-acyclic planar graph J for which the edge variant of the Erdős–Pósa property holds for $\mathcal{M}(J)$, we have that $f_{\mathcal{M}(J)}(k) = \Omega(k \log k)$. Indeed, let G be an n-vertex cubic graph with treewidth $\Omega(n)$ and girth $\Omega(\log n)$ (such graphs are well-known to exist). Since J is planar, the treewidth of any graph excluding J as a minor is bounded by a constant [18], hence any set of edges of G meeting all models of J has size $\Omega(n)$ (as the removal of an edge may decrease the treewidth by at most two). On the other hand, since J contains a cycle and the girth of G is $\Omega(\log n)$, any

model of J in G contains $\Omega(\log n)$ edges (assuming that J does not have isolated vertices), and therefore G contains $O(n/\log n)$ edge-disjoint models of J (here we have used that the degree of G is bounded), easily implying that $f_{\mathcal{M}(J)}(k) = \Omega(k \log k)$. In particular, it holds that $f_{\mathcal{M}(\theta_r)}(k) = \Omega(k \log k)$ for any $r \geqslant 2$, so a first avenue for further work in this direction is to optimize the gap function $f_{\mathcal{M}(\theta_r)}(k)$ given in Theorem 1.

Finally, when the graphs G (in which the packings or coverings are taken) are restricted to classes of bounded degree, the proof of Lemma 7 can easily be adapted to prove that the bound of the vertex version also holds for the edge version.

References

[1] E. Birmelé, J. Bondy, B. Reed, Brambles, prisms and grids, in: A. Bondy, J. Fonlupt, J.-L. Fouquet, J.-C. Fournier, J.L. Ramírez Alfonsín (Eds.), Graph Theory in Paris, in: Trends in Mathematics, Birkhäuser, Basel, 2007, pp. 37–44.
[2] C. Chekuri, J. Chuzhoy, Large-treewidth graph decompositions and applications, in: Proceedings of the Forty-fifth Annual ACM Symposium on Theory of Computing, STOC'13, ACM, New York, NY, USA, 2013, pp. 291–300.
[3] C. Chekuri, J. Chuzhoy, Polynomial bounds for the grid-minor theorem, in: Proceedings of the 46th Annual ACM Symposium on Theory of Computing, STOC'14, ACM, New York, NY, USA, 2014, pp. 60–69.
[4] R. Diestel, Graph Theory, third ed., in: Graduate Texts in Mathematics, vol. 173, Springer-Verlag, Heidelberg, 2005.
[5] P. Erdős, L. Pósa, On independent circuits contained in a graph, Canad. J. Math. 17 (1965) 347–352.
[6] S. Fiorini, T. Huyhn, G. Joret, personal communication, 2013.
[7] S. Fiorini, G. Joret, I. Sau, Optimal Erdős–Pósa property for pumpkins, Manuscript, 2013.
[8] S. Fiorini, G. Joret, D.R. Wood, Excluded forest minors and the Erdős-Pósa property, Combin. Probab. Comput. 22 (5) (2013) 700–721.
[9] F.V. Fomin, D. Lokshtanov, N. Misra, G. Philip, S. Saurabh, Quadratic upper bounds on the Erdős–Pósa property for a generalization of packing and covering cycles, J. Graph Theory 74 (4) (2013) 417–424.
[10] F.V. Fomin, S. Saurabh, D.M. Thilikos, Strengthening Erdős–Pósa property for minor-closed graph classes, J. Graph Theory 66 (3) (2011) 235–240.
[11] G. Joret, C. Paul, I. Sau, S. Saurabh, S. Thomassé, Hitting and harvesting pumpkins, SIAM J. Discrete Math. 28 (3) (2014) 1363–1390.
[12] K.-I. Kawarabayashi, Y. Kobayashi, Edge-disjoint odd cycles in 4-edge-connected graphs, in: 29th International Symposium on Theoretical Aspects of Computer Science, STACS, 2012, pp. 206–217.
[13] T. Kloks, Treewidth. Computations and Approximations, in: LNCS, vol. 842, Springer, 1994.
[14] A. Leaf, P. Seymour, Treewidth and planar minors, Manuscript, 2012.
[15] J.-F. Raymond, D. Thilikos, Polynomial gap extensions of the erdős–Pósa theorem, in: J. Nešetřil, M. Pellegrini (Eds.), The Seventh European Conference on Combinatorics, Graph Theory and Applications, in: CRM Series, vol. 16, Scuola Normale Superiore, 2013, pp. 13–18.
[16] J.-F. Raymond, D.M. Thilikos, Low Polynomial Exclusion of Planar Graph Patterns, ArXiv e-prints, May 2013.
[17] B.A. Reed, D.R. Wood, Polynomial treewidth forces a large grid-like-minor, European J. Combin. 33 (3) (2012) 374–379.
[18] N. Robertson, P.D. Seymour, Graph Minors. V. Excluding a planar graph, J. Combin. Theory Ser. B 41 (2) (1986) 92–114.