A TIGHT ERDŐS-PÓSA FUNCTION FOR WHEEL MINORS

PIERRE ABOULKER, SAMUEL FIORINI, TONY HUYNH, GWENAËL JORET, JEAN-FLORENT RAYMOND, AND IGNASI SAU

Abstract. Let W_t denote the wheel on $t+1$ vertices. We prove that for every integer $t \geq 3$ there is a constant $c = c(t)$ such that for every integer $k \geq 1$ and every graph G, either G has k vertex-disjoint subgraphs each containing W_t as a minor, or there is a subset X of at most $ck \log k$ vertices such that $G - X$ has no W_t minor. This is best possible, up to the value of c. We conjecture that the result remains true more generally if we replace W_t with any fixed planar graph H.

1. Introduction

Let H be a fixed graph. An *H-model* $\mathcal{M}$ in a graph G is a collection $\{S_x \subseteq G : x \in V(H)\}$ of vertex-disjoint connected subgraphs of G such that S_x and S_y are linked by an edge in G for every edge $xy \in E(H)$. The *vertex set* $V(\mathcal{M})$ of $\mathcal{M}$ is the union of the vertex sets of the subgraphs in the collection. Two H-models $\mathcal{M}$ and $\mathcal{M}'$ are *disjoint* if $V(\mathcal{M}) \cap V(\mathcal{M}') = \varnothing$.

Let $\nu_H(G)$ be the maximum number of pairwise disjoint H-models in G. Let $\tau_H(G)$ be the minimum size of a subset $X \subseteq V(G)$ such that $G - X$ has no H-model. Clearly, $\nu_H(G) \leq \tau_H(G)$. We say that the *Erdős-Pósa property holds for H-models* if there exists a *bounding function* $f : \mathbb{N} \to \mathbb{R}$ such that

$$\tau_H(G) \leq f(\nu_H(G))$$

for every graph G.

Robertson and Seymour [16] proved that the Erdős-Pósa property holds for H-models if and only if H is planar. Their original bounding function was exponential. However, this has been significantly improved by recent breakthrough results of Chekuri and Chuzhoy [3, 4] on the polynomial Grid Theorem.

Theorem 1.1 (Chekuri and Chuzhoy [3])**.** *There exist integers $a, b, c \geq 0$ such that for every planar graph H on h vertices, the Erdős-Pósa property holds for H-models with*

Date: July 5, 2018.

G. Joret is supported by an ARC grant from the Wallonia-Brussels Federation of Belgium. J.-F. Raymond is supported by ERC Consolidator Grant 648527-DISTRUCT and has been supported by the Polish National Science Centre grant PRELUDIUM DEC-2013/11/N/ST6/02706. P. Aboulker, S. Fiorini, and T. Huynh are supported by ERC Consolidator Grant 615640-ForEFront.

bounding function

$$f(k) = ah^b \cdot k \log^c(k+1).$$

If we consider H to be fixed and focus solely on the dependence on k—which is the point of view we take in this paper—Theorem 1.1 gives a $O(k \log^c k)$ bounding function. This is remarkably close to being best possible: If H is planar with at least one cycle, then there is a $\Omega(k \log k)$ lower bound on bounding functions. This follows easily from the original lower bound of Erdős and Pósa for the case where H is a triangle [6]. Alternatively, one can see this by considering n-vertex graphs G with treewidth $\Omega(n)$ and girth $\Omega(\log n)$ (which exist [13]), and notice that $\tau_H(G) = \Omega(n)$ (because removing one vertex decreases treewidth by at most one, and $G - X$ has treewidth $O(1)$ when $G - X$ has no H-minor, by the Grid Theorem) while $\nu_H(G) = O(n/\log n)$ (because each H-model contains a cycle). Thus, a $O(k \log^c k)$ bound is optimal, up to the value of c. While no explicit value for c is given in [3], a quick glance at the proof suggests that it is at least a double-digit integer. In this paper, we put forward the conjecture that a $O(k \log^c k)$ bound holds with $c = 1$.

Conjecture 1.2. *For every planar graph H, the Erdős-Pósa property holds for H-models with a $O(k \log k)$ bounding function.*

If true, Conjecture 1.2 would completely settle the growth rate of the Erdős-Pósa functions for H-models for all planar graphs H (up to the constant factor depending on H). That is, if H is planar with at least one cycle, then the $O(k \log k)$ bound would match the $\Omega(k \log k)$ lower bound mentioned above. And if H is a forest, it is already known that the right order of magnitude is $O(k)$, see [9].

Going back to the $O(k \log^c k)$ bound of Chekuri and Chuzhoy [3], one could initially hope that a value of $c = 1$ could be obtained by optimizing the various steps of their proof. However, any constant c obtained using their general approach necessarily satisfies $c \geq 2$. This is because they obtain Theorem 1.1 as a corollary from the following result.

Theorem 1.3 (Chekuri and Chuzhoy [3])**.** *There exist integers $a', b', c' \geq 0$ such that for all integers $r, k \geq 1$, every graph G of treewidth at least*

$$a' r^{b'} \cdot k \log^{c'}(k+1)$$

has k vertex-disjoint subgraphs $G_1, \dots, G_k$, each of treewidth at least r.

Now, if we fix a planar graph H and if G is such that $\nu_H(G) < k$, then G cannot have k vertex-disjoint subgraphs each of treewidth at least r, where $r = r(H)$ is a constant such that every graph with treewidth at least r contains an H minor. Note that $r(H)$ exists by the Grid Theorem of Robertson and Seymour [16]. Thus, the above theorem implies that G has treewidth $O(k \log^{c'} k)$. The authors of [3] then apply a standard divide-and-conquer approach on an optimal tree decomposition, and obtain a $O(k \log^c k)$ bound on $\tau_H(G)$ (see [3, Lemma 5.4]). This unfortunately results in an extra $\log k$ factor;

$c = c' + 1$. On the other hand, we must have $c' \geq 1$ in Theorem 1.3, as shown again by n-vertex graphs with treewidth $\Omega(n)$ and girth $\Omega(\log n)$. Hence, $c \geq 2$. Therefore, one needs a different approach to prove Conjecture 1.2.

As a side remark, it is natural to conjecture that we could take $c' = 1$ in Theorem 1.3 (at least, if we forget about the precise dependence on r):

Conjecture 1.4. *There is a function $f : \mathbb{N} \to \mathbb{N}$ such that for all integers $r, k \geq 1$, every graph G of treewidth at least*

$$f(r) \cdot k \log(k + 1)$$

has k vertex-disjoint subgraphs $G_1, \dots, G_k$, each of treewidth at least r.

As it turns out, this conjecture is implied by our Conjecture 1.2: It suffices to take H to be the $r \times r$-grid, which has treewidth r. Then either $\nu_H(G) \geq k$, in which case we are done, or $\nu_H(G) < k$, and then there is a subset X of $O(k \log k)$ vertices such that $G - X$ has no H-minor, and hence $G - X$ has treewidth at most $g(r)$ for some function g by the Grid Theorem. Adding X to all bags of an optimal tree decomposition of $G - X$, we deduce that G has treewidth $O(k \log k)$. Thus, this is another motivation to study Conjecture 1.2.

While Conjecture 1.2 remains open in general, it is known to hold for some specific graphs H. For example, the original Erdős-Pósa theorem [6] is simply the assertion that Conjecture 1.2 holds when H is a triangle. This was recently extended to the case where H is an arbitrary cycle [7] (see also [1, 14] for related results). The conjecture also holds when H is a multigraph consisting of two vertices linked by a number of parallel edges [2].

Our main result is that Conjecture 1.2 holds when H is a wheel. A *wheel* is a graph obtained from a cycle by adding a new vertex adjacent to all vertices of the cycle. We denote by W_t the wheel on $t + 1$ vertices.

Theorem 1.5. *For each integer $t \geq 3$, the Erdős-Pósa property holds for W_t-models with a $O(k \log k)$ bounding function.*

We remark that our main theorem implies all the aforementioned special cases. This is because the existence of a $O(k \log k)$ bounding function for H-models is preserved under taking minors of H (see Lemma 2.7). Our result also have the following consequence.

Corollary 1.6. *For every $t \in \mathbb{N}$ there is a function $g \colon \mathbb{N} \to \mathbb{N}$ with $g(k) = O(k \log k)$ such that every $(k \cdot W_t)$-minor free graph has treewidth at most $g(k)$.*

The rest of the paper is organized as follows. In the next section we present some general lemmas about H-models. Since these lemmas are valid for arbitrary planar graphs H, they may be useful in attacking Conjectures 1.2 and 1.4. In Section 3, we specialize to the case of wheels and prove our main theorem. We conclude with some open problems in Section 4.

2. General Tools

In this paper, our graphs are simple (no parallel edges nor loops). Let H, G be two graphs. We let $|G|$ denote $|V(G)|$. We assume the reader is familiar with the notions of graph minors, tree decompositions, and treewidth (see Diestel [5] for an introduction to the area). We let $\mathrm{tw}(G)$ denote the treewidth of G.

An *H-transversal* of G is a set X of vertices of G such that $G - X$ has no H-model. A graph is *minor-minimal* for a given property if it satisfies the property and none of its proper minors does.

We use the following results. The first is an extension of a classic result of Kostochka [11] and Thomason [18], where in addition the size of the K_t-model is logarithmic. (For definiteness, all logarithms are in base 2 in this paper.)

Theorem 2.1 ([12], see also [8, 17])**.** *There is a function $\varphi(t) = O(t\sqrt{\log t})$ such that, if an n-vertex graph has average degree at least $\varphi(t)$, then it contains a K_t-model on $O(\log n)$ vertices.*

The second is a theorem of Fomin, Lokshtanov, Misra and Saurabh [10], whose original purpose was to show that the algorithmic problem of finding a minimum-size H-transversal admits a polynomial-size kernel when H is planar.

Theorem 2.2 ([10])**.** *For every planar graph H, there is a polynomial π such that for every $k \in \mathbb{N}$, every graph G with $\tau_H(G) = k$ and minor-minimal with this property satisfies $|G| \leq \pi(k)$.*

2.1. Minimal counterexamples to the Erdős-Pósa property.

Let H be a graph and let $f : \mathbb{N} \to \mathbb{R}$ be a function. We say that a graph G is a *minimal counterexample* to the Erdős-Pósa property for H-models with bounding function f if the following properties hold:

(i) $\tau_H(G) > f(\nu_H(G))$;
(ii) subject to the above constraint, $\nu_H(G)$ is minimum;
(iii) subject to the above constraints, $|G|$ is minimum;
(iv) subject to the above constraints, $|E(G)|$ is minimum.

Notice that the two last requirements of the above definition imply that a minimal counterexample is a minor-minimal graph satisfying requirements (i) and (ii). The following lemma gives a bound on the size of minimal counterexamples.

Lemma 2.3. *Let H be a planar graph and let $f : \mathbb{N} \to \mathbb{R}$ be a polynomial non-decreasing function. Then there is a polynomial ρ such that, for every minimal counterexample G to the Erdős-Pósa property for H-models with bounding function f, we have $|G| \leq \rho(\nu_H(G))$.*

Proof. Let $k := \nu_H(G)$. Let us first show that $\tau_H(G) = \lfloor f(k) \rfloor + 1$. Let $v \in V(G)$. Observe that $\tau_H(G) \leq \tau_H(G - v) + 1$. By minimality of G, $\tau_H(G - v) \leq f(\nu_H(G - v))$. As $G - v$ is a minor of G, we also have $\nu_H(G - v) \leq \nu_H(G)$. We deduce $\tau_H(G) \leq f(k) + 1$. Since G is a counterexample, we also have $\tau_H(G) > f(k)$. It follows that $\tau_H(G) = \lfloor f(k) \rfloor + 1$.

Now, $\nu_H(G') \leq \nu_H(G)$ holds for every proper minor G' of G, and thus $\tau_H(G') < \tau_H(G)$ (otherwise G would not be a minimal counterexample). Hence G is minor-minimal with the property that $\tau_H(G) = \lfloor f(k) \rfloor + 1$. By Theorem 2.2 we obtain $|G| \leq \pi(\lfloor f(k) \rfloor + 1)$ where π is the polynomial given by that theorem. Therefore, it suffices to take $\rho : t \mapsto \pi(\lfloor f(t) \rfloor + 1)$. ☐

Informally, the following result, originally proved in [9], states that if a graph G has a large H-minor-free induced subgraph with a small 'boundary', then there is a smaller graph G' where the values of ν_H and τ_H are the same.

Theorem 2.4 ([9])**.** *For every planar graph H, there is a computable function $g' : \mathbb{N} \to \mathbb{N}$ such that, for every graph G, if J is an H-minor-free induced subgraph of G such that exactly p vertices of J have a neighbor in $V(G) \setminus V(J)$ and $|J| \geq g'(p)$, then there exists a graph G' such that $\tau_H(G') = \tau_H(G)$, $\nu_H(G') = \nu_H(G)$, and $|G'| < |G|$.*

We can use Theorem 2.4 to upper bound the size of H-minor-free induced subgraphs in minimal counterexamples as follows.

Corollary 2.5. *For every planar graph H, there is a computable and non-decreasing function $g : \mathbb{N} \to \mathbb{N}$ such that, if G is a minimal counterexample to the Erdős-Pósa property for H-models with bounding function f for some function $f : \mathbb{N} \to \mathbb{R}$, then every H-minor free induced subgraph J of G that has exactly p vertices with a neighbor in $V(G) \setminus V(J)$ satisfies $|J| < g(p)$.*

Proof. Let g' be the function in Theorem 2.4. We define the function $g : \mathbb{N} \to \mathbb{N}$ as follows: $g(k) = \max_{i \in \{0,\dots,k\}} g'(i)$. Notice that $g(k) \geq g'(k)$ holds for every $k \in \mathbb{N}$ and that g is non-decreasing. Now, suppose that G is a graph having an H-minor free induced subgraph J with exactly p vertices having a neighbor in $V(G) \setminus V(J)$ and such that $|J| \geq g(p)$. Then, since $g(p) \geq g'(p)$, by Theorem 2.4 there is a graph G' such that $\tau_H(G') = \tau_H(G)$, $\nu_H(G') = \nu_H(G)$, and $|G'| < |G|$. In particular, G cannot be a minimal counterexample to the Erdős-Pósa property for H-models for any bounding function f, a contradiction. ☐

2.2. **Interplays between treewidth and the Erdős-Pósa property.** Given a planar graph H, the standard approach to show that H-models satisfy the Erdős-Pósa property is to first note that $k \cdot H$ (the disjoint union of k copies of H) is also planar. Thus, if $\nu_H(G) < k$ for a graph G, then the treewidth of G is bounded by a function of k and H, by the Grid Theorem [16]. Then one can use a tree decomposition of small width to find a small H-transversal of G. This was first used by Robertson and Seymour [16, Theorem 8.8] in their original proof (see also [19, Theorem 3] and the survey [15, Section 3]). It was subsequently used by several authors to obtain improved bounding functions, most notably by Chekuri and Chuzhoy [3] when deriving their Theorem 1.1 from Theorem 1.3.

As it was already mentioned in the introduction when discussing Conjecture 1.4, the reverse direction holds as well: A bounding function for H-models translates directly to an upper bound on the treewidth of $(k \cdot H)$-minor free graphs, up to an additive term depending only on H:

Lemma 2.6. *Let H be a planar graph, let f be a bounding function for H-models, and let $c = c(H)$ be a constant such that $\mathrm{tw}(G) \leq c$ for every H-minor free graph G. Then, for every $k \geq 1$, every $(k \cdot H)$-minor free graph G has treewidth at most $f(k-1) + c$.*

Proof. Let G be a graph not containing $k \cdot H$ as a minor. Since $\nu_H(G) \leq k-1$ and f is a bounding function for H-models, we deduce $\tau_H(G) \leq f(k-1)$. That is, G has a set X of at most $f(k-1)$ vertices such that $G - X$ is H-minor free. By definition of c, we have $\mathrm{tw}(G - X) \leq c$. Then $\mathrm{tw}(G) \leq c + |X| \leq c + f(k-1)$, as desired. □

Thus combining our main result with Lemma 2.6 gives Corollary 1.6 (stated in the introduction).

We also include the following lemma, which states that if H' is a minor of H, then a bounding function for H'-models can be easily obtained from a bounding function for H-models.

Lemma 2.7. *Let H be a fixed planar graph, let f be a bounding function for H-models, and let $c = c(H)$ be a constant such that $\mathrm{tw}(G) \leq c$ for all H-minor free graphs G. If H' is a minor of H with q connected components, then $k \mapsto f(k) + (qk-1)(c+1)$ is a $O(f(k))$ bounding function for H'-models.*

Proof. Let G be a graph with $\nu_{H'}(G) \leq k$. As H' is a minor of H, we deduce $\nu_H(G) \leq k$. By definition of f, there is a set X of at most $f(k)$ vertices such that $G - X$ is H-minor free. Hence $\mathrm{tw}(G - X) \leq c$. Theorem 8.8 in [16] provides the following upper-bound on τ in graphs of bounded treewidth:

$$\tau_{H'}(G - X) \leq (qk-1)(c+1).$$

Then, $\tau_{H'}(G) \leq \tau_{H'}(G - X) + |X| \leq (qk-1)(c+1) + f(k)$. Finally, since $f(k) \geq k$, we deduce $(qk-1)(c+1) + f(k) = O(f(k))$. □

3. The Proof for Wheels

In this section, we prove our main theorem:

Theorem 1.5. *For each integer $t \geq 3$, the Erdős-Pósa property holds for W_t-models with a $O(k \log k)$ bounding function.*

Proof. To keep track of the dependencies between the constants that we use, we define them here. Recall that t denotes the number of spokes of the wheel that we are considering. Let φ and φ' be constants such that every n-vertex graph of average degree at least φ has a K_{t+1}-model on at most $\varphi' \log n$ vertices (both φ and φ' depend on t, see Theorem 2.1). Let $\alpha, \beta \geq 1$ be constants such that $\rho(n) \leq \alpha n^{\beta}$, for every $n \in \mathbb{N} \setminus \{0\}$, where ρ is the polynomial of Lemma 2.3 for $H = W_t$. Let g denote the function from Corollary 2.5 for $H = W_t$.

We then set

$$\begin{aligned}
&c_1 = g(2t\varphi^2), \quad p = g(2c_1\varphi^2), \quad c_2 = 4p,\\
&\sigma = \max\left\{3\varphi' c_2, 2c_2 + tp, (2t^2p + 1)(c_2 + 2c_1\varphi^2)\right\}, \quad \text{and}\\
&\gamma = \sigma(\beta + \log \alpha).
\end{aligned}$$

Observe that we have $t < c_1 < p < c_2$. Let $f(k) := \gamma \cdot k \log(k + 1)$, for every $k \in \mathbb{N}$. We show that the Erdős-Pósa property holds for W_t-models with bounding function f.

Arguing by contradiction, let G be a minimal counterexample to the Erdős-Pósa property for W_t-models with bounding function f. Let $k := \nu_{W_t}(G)$ Then $k \geq 1$ and $|G|$ is polynomial in k by Lemma 2.3. That is, letting $n := |G|$, we have $n \leq \alpha k^{\beta}$, for the constants α and β defined above.

We first show that G cannot contain a W_t-model of logarithmic size.

Claim 3.1. *G has no W_t-model of size at most $\sigma \log n$.*

Proof. Towards a contradiction, we consider a W_t-model $\mathcal{M}$ of size at most $\sigma \log n$. Notice that $\log n \leq (\beta + \log \alpha) \log(k + 1)$ can be deduced from the aforementioned upper-bound on n. Since $\nu_{W_t}(G - V(\mathcal{M})) \leq k - 1$, by minimality of G,

$$\begin{aligned}
\tau_{W_t}(G) &\leq |V(\mathcal{M})| + \tau_{W_t}(G - V(\mathcal{M}))\\
&\leq \sigma \log n + f(k - 1)\\
&\leq \gamma \log(k + 1) + \gamma(k - 1) \log k\\
&\leq \gamma k \log(k + 1) \leq f(k).
\end{aligned}$$

However, this contradicts the fact that G is a minimal counterexample to the Erdős-Pósa property for W_t-models with bounding function f. $\square$

Note that since $\nu_{W_t}(G) \geq 1$, we have $n \geq t+1 \geq 2$, and thus $\log n \geq 1$ (recall that all logarithms are in base 2). Thus, Claim 3.1 implies in particular that G has no W_t-model of size at most σ.

Let $\mathcal{C}$ be a maximum-size collection of vertex-disjoint cycles in G whose lengths are in the interval $[c_1, c_2]$. Let $\mathcal{P}$ be a maximum-size collection of vertex-disjoint paths of length p in $G - V(\mathcal{C})$, where $V(\mathcal{C}) := \bigcup_{C \in \mathcal{C}} V(C)$. (In this paper, the length of a path is defined as its number of edges.) Finally, let $\mathcal{R}$ be the collection of components of $G - (V(\mathcal{C}) \cup V(\mathcal{P}))$ and let $V(\mathcal{R}) := V(G) - (V(\mathcal{C}) \cup V(\mathcal{P}))$, where $V(\mathcal{P}) := \bigcup_{P \in \mathcal{P}} V(P)$. We point out that the cycles in $\mathcal{C}$ and the paths in $\mathcal{P}$ are subgraphs of G but not necessarily induced subgraphs of G, while the components in $\mathcal{R}$ are induced subgraphs of G. We call the elements of $\mathcal{C} \cup \mathcal{P} \cup \mathcal{R}$ *pieces*.

Observe that, by maximality of $\mathcal{P}$, every path in a piece of $\mathcal{R}$ has length at most $p-1$. This implies that each such piece is W_t-minor free. Indeed, observe that if such a piece R has a W_t-model then R contains a subgraph consisting of a cycle C and a rooted tree T such that T has at most t leaves, $V(T) \cap V(C) = \emptyset$, and the leaves of T collectively have at least t neighbours in C. The cycle C has at most p vertices, and each root-to-leaf path in T has at most p vertices. Thus, this gives a W_t-model with at most $(t+1)p$ vertices. However, this contradicts Claim 3.1 since $(t+1)p \leq \sigma$.

Similarly, each piece in $\mathcal{C}$ (in $\mathcal{P}$, respectively) has at most c_2 (p, respectively) vertices, and these vertices induce a W_t-minor free subgraph of G; otherwise, there would exist a W_t-model of size at most c_2 (resp. p), again a contradiction to Claim 3.1 since $p \leq c_2 \leq \sigma$. These facts will be used often in the rest of the proof.

We say that two distinct pieces K and K' *touch* if some edge of G links some vertex of K to some vertex of K'. Note that, by construction, two distinct pieces in $\mathcal{R}$ cannot touch. A piece is said to be *central* if it is a cycle in $\mathcal{C}$, a path in $\mathcal{P}$, or a piece in $\mathcal{R}$ that touches at least 2φ other pieces. In the next paragraph, we define two auxiliary simple graphs H_s (for small degrees) and H_b (for big degrees) that model how the central pieces are connected through the noncentral pieces. To keep track of the correspondence between the edges of H_s and the noncentral pieces, we put labels on some of these edges.

Initialize both H_s and H_b to the graph whose set of vertices is the set of central pieces and whose set of edges is empty. For each pair of central pieces that touch in G, add an (unlabeled) edge between the corresponding vertices in both H_s and H_b.

Next, while there is some noncentral piece $R \in \mathcal{R}$ that touches two central pieces K and K' that are not yet adjacent in H_b, call $\mathcal{Z}_R$ the set of central pieces that touch R and do the following:

(1) add all (unlabeled) edges to H_{b} between pieces of $\mathcal{Z}_R$ (not already present in H_{b}). This creates a clique on vertex set $\mathcal{Z}_R$ in H_{b}, some of whose edges might have already been there before.
(2) choose a piece $K \in \mathcal{Z}_R$ such that the number of newly added edges of H_{b} incident to K is maximum. Add to H_{s} every edge that links K to another piece of $\mathcal{Z}_R$ (not already present in H_{s}), and label it with R. This creates a star centered at K in H_{s} with all its edges labeled with the noncentral piece R.

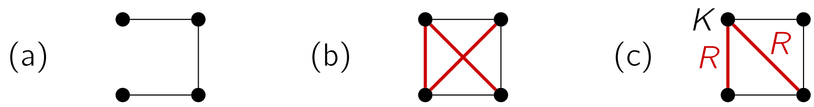

Figure 1. Construction of H_{s} and H_{b}. (a): the vertices of a set $\mathcal{Z}_R$ in H_{s} and H_{b} before step (1). (b): $H_{\mathrm{b}}[\mathcal{Z}_R]$ after step (1). (c): $H_{\mathrm{s}}[\mathcal{Z}_R]$ after step (2).

The edges added during these steps are depicted in thick red lines in the example of Figure 1. By construction, H_{s} is a subgraph of H_{b}. These graphs have the following two crucial properties.

Claim 3.2. *If H_{s} has a W_t-model of size q, then G has a W_t-model of size at most $3c_2q$.*

Proof. Suppose that H_{s} has a W_t-model of size q. Then there exists a subgraph $M_{\mathrm{s}} \subseteq H_{\mathrm{s}}$ with q vertices that can be contracted to W_t. We may assume that the average degree of M_{s} is at most that of W_t, and hence at most 4. That is, $|E(M_{\mathrm{s}})| \leq 2|M_{\mathrm{s}}|$. From the subgraph M_{s}, we construct a subgraph $M \subseteq G$ that can be contracted to M_{s}, and thus also to W_t.

First, for each central piece $K \in V(M_{\mathrm{s}}) \cap (\mathcal{C} \cup \mathcal{P})$, we add all its vertices to M, as well as $|K| - 1$ edges from K in such a way that the subgraph of M induced by $V(K)$ is connected. For each central piece $K \in V(M_{\mathrm{s}}) \cap \mathcal{R}$, we choose some vertex $v_K \in K$ and add it to M. This creates at most $c_2|M_{\mathrm{s}}| = c_2q$ vertices in M.

Second, for each unlabeled edge KK' of M_{s} with $K, K' \in V(M_{\mathrm{s}}) \cap (\mathcal{C} \cup \mathcal{P})$, we choose some edge of G linking K to K' and add it to M. This does not create any new vertex in M.

Third, for each edge KK' of M_{s} that has not been considered so far, we add to M a path linking some vertex of $V(M) \cap V(K)$ to some vertex of $V(M) \cap V(K')$, as follows. If the edge KK' is not labeled, then exactly one of its endpoints is a central piece in $\mathcal{R}$, say K. The path we add to M links v_K to some vertex of K' and is a subgraph of K, except for the last edge and last vertex. Thus, this path has at most $p - 1$ internal vertices. If the edge KK' is labeled with the noncentral piece $R \in \mathcal{R}$, then this edge is part of a star in M_{s} whose edges are all labeled with R. We may assume without loss of generality that K is the center of this star. In this case, the path we add to M links some vertex of K to some vertex of K' and has all its internal vertices in R. Thus, this path has at most p internal vertices.

In total, the addition of these paths to M creates at most $p|E(M_s)| \leq 2c_2|M_s| = 2c_2q$ new vertices in M. The resulting subgraph M has at most $c_2|M_s| + p|E(M_s)| \leq 3c_2q$ vertices. By construction, M can be contracted to M_s, as desired. □

Claim 3.3. *The average degree of H_b is at most φ times the average degree of H_s. The degree of each central piece of $\mathcal{R}$ in H_s is at least 2φ.*

Proof. First, note that edges that appear in H_b but not in H_s must be labeled. Let $R \in \mathcal{R}$ be a noncentral piece, and let r be the number of pieces in $\mathcal{C} \cup \mathcal{P}$ it touches. By definition of noncentral pieces, $r < 2\varphi$. When R is treated in the algorithm used to construct H_b and H_s, if q new edges are added to H_b, then one of the pieces touched by R is incident to at least $2q/r > q/\varphi$ of these new edges and thus at least q/φ new edges are added to H_s. This proves the first part of the claim.

By definition, a piece K of $\mathcal{R}$ is central if it touches at least 2φ other pieces. As two pieces of $\mathcal{R}$ cannot touch, K touches at least 2φ pieces from $\mathcal{C} \cup \mathcal{P}$, that is, at least 2φ other central pieces. Then in the first step of the construction of H_s, all edges have been added from K to these pieces. □

If the average degree of H_s is at least φ, then by definition of φ and φ' at the beginning of the proof, H_s has a K_{t+1}-model of size at most $\varphi' \log |H_s|$, and thus in particular a W_t-model of size at most $\varphi' \log |H_s|$. By Claim 3.2, this gives a W_t-model of size at most $3\varphi' c_2 \log |H_s| \leq 3\varphi' c_2 \log n$ in G, a contradiction to Claim 3.1 since $3\varphi' c_2 \leq \sigma$.

Thus, the average degree of H_s is smaller than φ. Hence, by Claim 3.3, the average degree of H_b is smaller than φ^2. Then strictly more than half of the central pieces have degree less than 2φ in H_s (otherwise at least half of the vertices of H_s have degree at least 2φ, a contradiction to the fact that H_s has average degree less than φ). Similarly, strictly more than half of the central pieces have degree less than $2\varphi^2$ in H_b. Thus there is a central piece whose degree in H_s is less than 2φ, and whose degree in H_b is less than $2\varphi^2$. Choose such a piece K. By Claim 3.3 (second part of the statement), K is either in $\mathcal{C}$ or in $\mathcal{P}$.

In the rest of the proof we use the fact that K has degree less than $2\varphi^2$ in H_b to find a W_t-model of size at most $\sigma \log n$, contradicting Claim 3.1.

If K is the unique central piece in $\mathcal{C} \cup \mathcal{P}$, then $V(K)$ is a W_t-transversal of G since each piece in $\mathcal{R}$ is W_t-minor free. Thus $\tau_{W_t}(G) \leq |K| \leq c_2 \leq f(k)$, contradicting the fact that G is a counterexample.

For each central piece K' adjacent to K in H_b, we consider the collection $\mathcal{R}_{K,K'}$ of all noncentral pieces $R \in \mathcal{R}$ that touch both K and K' ($\mathcal{R}_{K,K'}$ might be empty). Then we consider the subgraph $G_{K'}$ of G induced by $V(K) \cup V(K') \cup V(\mathcal{R}_{K,K'})$.

Let q be an integer equal to t if $K \in \mathcal{C}$, to c_1 if $K \in \mathcal{P}$.

Our next goal is to show that for every central piece K' adjacent to K in H_b, there exists a set of strictly less than q vertices that separates K from K' in $G_{K'}$. Thus fix a piece K' adjacent to K in H_b. By Menger's theorem, it suffices to show that the maximum number of vertex-disjoint K–K' paths in $G_{K'}$ is strictly less than q. Assume for contradiction that $G_{K'}$ contains q vertex-disjoint K–K' paths.

By taking the paths to be as short as possible, we may assume that only their endpoints are in K and K', all their internal vertices are in pieces in $\mathcal{R}_{K,K'}$, and each such path intersects at most one piece in $\mathcal{R}_{K,K'}$ and thus has length at most $p + 1$.

Assume first that $K \in \mathcal{C}$, and so $q = t$. In this case $G_{K'}$ contains a small W_t-model as follows. Let T be a smallest tree in K' containing all the endpoints of our paths in K'. The center vertex of the wheel is then modeled by the union of T and the t K–K' paths (minus their endpoints in K). If $K' \in \mathcal{C} \cup \mathcal{P}$, then obviously $|V(T)| \leq c_2$, and the model thus has at most $2c_2 + tp$ vertices. If $K' \in \mathcal{R}$, then $|V(T)| \leq tp$ since each path in K' has length at most $p-1$; moreover, $\mathcal{R}_{K,K'}$ is empty in this case, implying that the model has at most $c_2 + tp$ vertices. Therefore, in both cases the resulting model has at most $2c_2 + tp$ vertices, which contradicts Claim 3.1 since $2c_2 + tp \leq \sigma$.

Assume now that $K \in \mathcal{P}$. Since $t < c_1$, by the previous case we may assume that $K' \in \mathcal{P} \cup \mathcal{R}$. Since there are $q = c_1$ vertex-disjoint K–K' paths in $G_{K'}$, two of these paths intersect K on two vertices that are at distance at least $c_1 - 1$ on the path K, which allows us to construct a cycle in $G_{K'}$ of length at least c_1 and at most $4p$: The cycle might use all the vertices of K and at most p vertices of K', which is at most $2p$ vertices, and might intersect at most two pieces of $\mathcal{R}_{K,K'}$, using at most p vertices in each of them. This is a contradiction to the maximality of $\mathcal{C}$: The length of this cycle is in the interval $[c_1, c_2]$ and yet the cycle is vertex disjoint from all cycles in $\mathcal{C}$.

Therefore, for each K' adjacent to K in H_b, there exists a set $X(K')$ with less than q vertices meeting all the K–K' paths in $G_{K'}$.

Let $X := \bigcup_{K'} X(K')$ where the union is taken over all central pieces K' adjacent to K in H_b. Note that $|X| \leq 2q\varphi^2$ since there are at most $2\varphi^2$ such pieces K' and for every K' we have $|X(K')| \leq q$.

We also note that X separates K from all other central pieces in G. To see this, let K'' be a central piece distinct from K and let Q be a K''–K path in G. Let K' be the last central piece that Q meets before reaching K. Then Q contains a K'–K path that is contained in $G_{K'}$, which must contain a vertex from X.

Let J be the union of the components of $G - X$ that intersect K. Observe that $V(K)$ is not completely included in X: If $K \in \mathcal{C}$, then $|K| \geq c_1 > 2t\varphi^2 \geq |X|$, and if $K \in \mathcal{P}$, then $|K| = p > 2c_1\varphi^2 \geq |X|$. Thus J is not empty. Note also that X separates J from the rest of the graph.

Suppose that the subgraph G' of G induced by $X \cup V(J)$ is W_t-minor free. Thus, by Corollary 2.5, $|G'| < g(|X|)$. We deduce

$$\begin{aligned} |X| + |J| &< g(|X|) && \text{since } |G'| = |X| + |J| \\ |K| &< g(|X|) && \text{since } |J| \geq |K| - |X| \\ |K| &< g(2q\varphi^2) && \text{since } g \text{ is non-decreasing.} \end{aligned}$$

Hence, if $K \in \mathcal{C}$, then $c_1 \leq |K| < g(2t\varphi^2)$, and if $K \in \mathcal{P}$, then $p \leq |K| < g(2c_1\varphi^2)$. Since $c_1 = g(2t\varphi^2)$ and $p = g(2c_1\varphi^2)$, we get a contradiction in both cases.

Thus, we may assume that G' contains a W_t-model. Let M be a subgraph of G' containing W_t as a minor with $|V(M)| + |E(M)|$ minimum. (We remark that here we take M to be a subgraph instead of just a model as before because we will need to consider the edges of that subgraph in the proof.) To finish the proof, it is now enough to prove that M has at most $\sigma \log n$ vertices, since by Claim 3.1 this will give us the desired contradiction.

Let $R(J) := J[V(\mathcal{R})]$. Thus $R(J)$ consists of a number of disjoint pieces or subgraphs of pieces of $\mathcal{R}$. Note that M might use all vertices of $V(K) \cup X$ (which is fine); what we need to prove is that it cannot use too many vertices of $R(J)$.

First, suppose that M is fully contained in some piece of $\mathcal{R}$. Since the vertices of M can be covered with $2t$ paths, and each path in the piece has length less than p, it follows that $|M| \leq 2tp \leq \sigma$ and we are done.

Thus we may assume that M is not contained in some piece of $\mathcal{R}$, and thus in particular M is not contained in $R(J)$ (since M is connected). By the above remark, we also know that each component of $M[V(R(J))]$ contains at most $2tp$ vertices. Since M has maximum degree at most t (by minimality of $|V(M)| + |E(M)|$), there are at most $t|V(K) \cup X|$ edges of M with one endpoint in $V(K) \cup X$ and the other in $R(J)$. Hence M intersects $R(J)$ on at most $2t^2p|V(K) \cup X|$ vertices. Therefore, M has at most $2t^2p|V(K) \cup X| + |V(K) \cup X|$ vertices. Since $|V(K) \cup X| \leq |K| + |X| \leq c_2 + 2c_1\varphi^2$, we deduce that $|M| \leq (2t^2p + 1)(c_2 + 2c_1\varphi^2) \leq \sigma$, as desired. □

+

4. Conclusion

One obvious extension of our result for wheels would be to prove it for all planar graphs. Note that the first steps of our proof work for any such H: Starting with G a minimal counterexample for some bounding function and some value k, we have that G has $n \leq k^{O(1)}$ vertices. Thus, in order to get a contradiction, it is enough to show that there is a $O(\log n)$-size H-model in G. Unfortunately, the rest of our proof is specific to wheels and does not generalize.

Let us mention another possible extension of our result. Strengthening the $O(k \log k)$ bound from [7], Mousset, Noever, Škorić, and Weissenberger [14] recently showed that there is a constant $c > 0$ such that for every $\ell \geq 3$, models of the ℓ-cycle C_ℓ have the Erdős-Pósa property with bounding function $ck \log k + \ell k$. In particular, the constant c in front of the $k \log k$ term is independent of ℓ. We expect that a similar property holds for wheels:

Conjecture 4.1. *There is a constant $c > 0$ and a function $g : \mathbb{N} \to \mathbb{N}$ such that for all integers $t \geq 3$, W_t-models have the Erdős-Pósa property with bounding function*

$$ck \log k + g(t)k.$$

Acknowledgements. We thank Stefan van Zwam for valuable discussions at an early stage of this project. We also thank Hong Liu for spotting a slight inaccuracy in the proof of Theorem 1.5 in an earlier version of this paper. Finally, we thank the referees for their careful reading and many helpful comments.

References

[1] E. Birmelé, J. A. Bondy, and B. A. Reed. The Erdős–Pósa property for long circuits. *Combinatorica*, 27(2):135–145, 2007.

[2] D. Chatzidimitriou, J.-F. Raymond, I. Sau, and D. M. Thilikos. An $O(\log \mathrm{OPT})$-approximation for covering and packing minor models of θ_r. *Algorithmica*, to appear in 2017. arXiv:1510.03945.

[3] C. Chekuri and J. Chuzhoy. Large-treewidth graph decompositions and applications. In *Proceedings of the 45th annual ACM Symposium on Theory of Computing*, pages 291–300. ACM, 2013. arXiv:1304.1577.

[4] C. Chekuri and J. Chuzhoy. Polynomial bounds for the grid-minor theorem. *Journal of the ACM*, 63(5):40:1–40:65, 2016. arXiv:1305.6577.

[5] R. Diestel. *Graph theory*, volume 173 of *Graduate Texts in Mathematics*. Springer, Heidelberg, fourth edition, 2010.

[6] P. Erdős and L. Pósa. On independent circuits contained in a graph. *Canadian Journal of Mathematics*, 17:347–352, 1965.

[7] S. Fiorini and A. Herinckx. A tighter Erdős–Pósa function for long cycles. *Journal of Graph Theory*, 77(2):111–116, 2014. arXiv:1205.0940.

[8] S. Fiorini, G. Joret, D. O. Theis, and D. R. Wood. Small minors in dense graphs. *European Journal of Combinatorics*, 33(6):1226–1245, 2012. arXiv:1005.0895.

[9] S. Fiorini, G. Joret, and D. R. Wood. Excluded forest minors and the Erdős–Pósa property. *Combinatorics, Probability and Computing*, 22(5):700–721, 2013. arXiv:1204.5192.

[10] F. V. Fomin, D. Lokshtanov, N. Misra, and S. Saurabh. Planar $\mathcal{F}$-Deletion: Approximation, Kernelization and Optimal FPT Algorithms. In *Proceedings of the 53rd Annual IEEE Symposium on Foundations of Computer Science*, pages 470–479. IEEE Computer Society, 2012. arXiv:1204.4230.

[11] A. V. Kostochka. Lower bound of the Hadwiger number of graphs by their average degree. *Combinatorica*, 4(4):307–316, 1984.

[12] R. Montgomery. Logarithmically small minors and topological minors. *Journal of the London Mathematical Society*, 91(1):71–88, 2015. arXiv:1309.7886.

[13] M. Morgenstern. Existence and explicit constructions of $q+1$ regular Ramanujan graphs for every prime power q. *J. Comb. Theory, Ser. B*, 62(1):44–62, 1994.

[14] F. Mousset, A. Noever, N. Škorić, and F. Weissenberger. A tight Erdős–Pósa function for long cycles. *Journal of Combinatorial Theory, Series B*, 125:21–32, 2017. arXiv:1603.07588.

[15] J.-F. Raymond and D. M. Thilikos. Recent techniques and results on the Erdős–Pósa property. *Discrete Applied Mathematics*, 231:25–43, 2017. arXiv:1603.04615.

[16] N. Robertson and P. D. Seymour. Graph minors. V. Excluding a planar graph. *Journal of Combinatorial Theory, Series B*, 41(1):92–114, 1986.

[17] A. Shapira and B. Sudakov. Small complete minors above the extremal edge density. *Combinatorica*, 35(1):75–94, 2015. arXiv:1208.3568.

[18] A. Thomason. An extremal function for contractions of graphs. *Mathematical Proceedings of the Cambridge Philosophical Society*, 95(2):261–265, 1984.

[19] C. Thomassen. On the presence of disjoint subgraphs of a specified type. *Journal of Graph Theory*, 12(1):101–111, 1988.

(P. Aboulker, S. Fiorini, T. Huynh)
Département de Mathématique
Université Libre de Bruxelles
Brussels, Belgium

Email address: {pierreaboulker@gmail.com, sfiorini@ulb.ac.be, tony.bourbaki@gmail.com}

(G. Joret)
Département d'Informatique
Université Libre de Bruxelles
Brussels, Belgium

Email address: gjoret@ulb.ac.be

(J.-F. Raymond)
Logic and Semantics Research Group
Technische Universität Berlin
Berlin, Germany

Email address: raymond@tu-berlin.de

(I. Sau)
CNRS, LIRMM, Université de Montpellier
Montpellier, France

Email address: ignasi.sau@lirmm.fr