

Excluded Forest Minors and the Erdős–Pósa Property

SAMUEL FIORINI¹, GWENAËL JORET² and DAVID R. WOOD^{3†}

¹Département de Mathématique, Université Libre de Bruxelles, Brussels, Belgium
(e-mail: sfiorini@ulb.ac.be)

²Département d'Informatique, Université Libre de Bruxelles, Brussels, Belgium
(e-mail: gjoret@ulb.ac.be)

³School of Mathematical Sciences, Monash University, Melbourne, Australia
(e-mail: david.wood@monash.edu)

Received 31 May 2012; revised 23 May 2013; first published online 8 July 2013

A classical result of Robertson and Seymour states that the set of graphs containing a fixed planar graph H as a minor has the so-called Erdős–Pósa property; namely, there exists a function f depending only on H such that, for every graph G and every positive integer k , the graph G has k vertex-disjoint subgraphs each containing H as a minor, or there exists a subset X of vertices of G with $|X| \leq f(k)$ such that $G - X$ has no H -minor (see Robertson and Seymour, *J. Combin. Theory Ser. B* **41** (1986) 92–114). While the best function f currently known is exponential in k , a $O(k \log k)$ bound is known in the special case where H is a forest. This is a consequence of a theorem of Bienstock, Robertson, Seymour and Thomas on the pathwidth of graphs with an excluded forest-minor. In this paper we show that the function f can be taken to be linear when H is a forest. This is best possible in the sense that no linear bound is possible if H has a cycle.

2010 *Mathematics subject classification*: Primary 05C83
Secondary 05C35

1. Introduction

Let \mathcal{F} be a finite set of graphs, which we will typically think of as a set of excluded (or forbidden) minors. Given a graph G , an \mathcal{F} -packing (or simply *packing*) in G is a collection of vertex-disjoint subgraphs of G each containing a member of \mathcal{F} as a minor. The maximum size of such a collection is denoted by $\nu_{\mathcal{F}}(G)$. A dual notion is that of an \mathcal{F} -transversal (or simply *transversal*) of G , which is defined as a subset X of vertices of G such that $G - X$ contains no member of \mathcal{F} as a minor. The minimum size of such a set X is denoted by $\tau_{\mathcal{F}}(G)$.

† Research of David Wood is supported by the Australian Research Council.

Thus if \mathcal{G} is a proper minor-closed class of graphs and \mathcal{F} is the corresponding set of minimal forbidden minors (which is finite by the graph minor theorem [27]), then $\tau_{\mathcal{F}}(G)$ is the minimum number of vertices to remove from G to obtain a graph in the class \mathcal{G} , while $\nu_{\mathcal{F}}(G)$ is the maximum number of vertex-disjoint forbidden minors in G .

These two graph invariants generalize some classical invariants in graph theory and combinatorial optimization. If $\mathcal{F} = \{K_2\}$ then it is easily seen that $\nu_{\mathcal{F}}(G)$ is the maximum size of a matching in G , while $\tau_{\mathcal{F}}(G)$ is the minimum size of a vertex cover in G . Similarly, if $\mathcal{F} = \{K_3\}$, then $\nu_{\mathcal{F}}(G)$ is the maximum size of a cycle packing in G , and $\tau_{\mathcal{F}}(G)$ is the minimum size of a cycle transversal (also known as a feedback vertex set) of G .

Clearly $\tau_{\mathcal{F}}(G) \geq \nu_{\mathcal{F}}(G)$ for every \mathcal{F} and every graph G . If at least one graph in \mathcal{F} is planar then the two parameters are ‘tied’ to each other in the following sense. There exists a function f depending only on \mathcal{F} such that $\tau_{\mathcal{F}}(G) \leq f(\nu_{\mathcal{F}}(G))$ for every graph G . This was shown by Robertson and Seymour [26]¹ and is a consequence of their well-known excluded grid theorem [26]. It is often referred to as the *Erdős–Pósa property* of \mathcal{F} (or more accurately, of the set of graphs contractible to some graph in \mathcal{F}), because Erdős and Pósa [11] proved the existence of such a function f when $\mathcal{F} = \{K_3\}$.

Robertson and Seymour’s result is best possible in the sense that no such function f exists if no graph in \mathcal{F} is planar (see [26]). However, the function f that follows from their proof is huge (exponential). This is because, as mentioned above, the proof relies on their excluded grid theorem, namely for every $r \in \mathbb{N}$ there exists a minimum integer $g(r)$ such that every graph with no $r \times r$ -grid minor has treewidth at most $g(r)$, and the current best upper bound on $g(r)$ is exponential: $g(r) \in 2^{O(r^2 \log r)}$ [20, 23]. (As for lower bounds, only $g(r) \in \Omega(r^2 \log r)$ is known [28].)

The situation changes drastically if instead of simply assuming that some graph in \mathcal{F} is planar, we further require that \mathcal{F} contains a forest F . Bienstock, Robertson, Seymour and Thomas [2] proved that every graph with pathwidth at least $|F| - 1$ contains F as a minor. Using this result, one can derive without much difficulty that $\tau_{\mathcal{F}}(G) \in O_{\mathcal{F}}(k \log k)$, where $k = \nu_{\mathcal{F}}(G)$; see Proposition 2.6 in Section 2. As far as we are aware, this is the best bound that is currently known.

In this paper we prove that a linear bound holds when \mathcal{F} contains a forest.

Theorem 1.1. *Let \mathcal{F} be a finite set of graphs containing at least one forest. Then there exists a computable constant $c = c(\mathcal{F})$ such that $\tau_{\mathcal{F}}(G) \leq c \cdot \nu_{\mathcal{F}}(G)$ for every graph G .*

If for some $t \in \mathbb{N}$ we let \mathcal{F} be the (finite) set of minimal forbidden minors for the class of graphs with pathwidth at most t , then \mathcal{F} contains at least one tree. This follows from the fact that trees have unbounded pathwidth (for instance, the complete binary tree of height h has pathwidth h). Using that \mathcal{F} is computable [1, 22], we obtain the following corollary from Theorem 1.1.

¹ It should be noted that this result is proved in [26] in the special case where \mathcal{F} consists of a single planar graph H . However, the general case follows by a straightforward modification of the proof of (8.8) in [26].

Corollary 1.2. *For every $t \in \mathbb{N}$ there exists a computable constant $c = c(t)$ such that, for every graph G and every $k \in \mathbb{N}$, either G contains k vertex-disjoint subgraphs each with pathwidth at least $t + 1$, or G has a vertex subset X of size at most $c \cdot k$ such that $G - X$ has pathwidth at most t .*

Interestingly, Corollary 1.2 becomes false if we replace pathwidth by treewidth, even for $t = 1$. A graph has treewidth at least 2 if and only if it contains a cycle, and there are graphs with no k vertex-disjoint cycles such that every cycle transversal has size $\Omega(k \log k)$ (see [11]). More generally, for every fixed set \mathcal{F} containing a planar graph but no forest, there are graphs G with $v_{\mathcal{F}}(G) = k$ and $\tau_{\mathcal{F}}(G) \in \Omega_{\mathcal{F}}(k \log k)$; this follows from the existence of n -vertex graphs G with treewidth $\Omega(n)$ and girth $\Omega(\log n)$.² In this sense Theorem 1.1 is best possible.

Our proof of Theorem 1.1 can very briefly be described as follows. First suppose for simplicity that $\mathcal{F} = \{T\}$ with T a tree; the general case can be reduced to this one without much difficulty. We show that every graph G either has a model of T of constant size, or a reduction operation can be applied to G , producing a graph G' with $v_{\mathcal{F}}(G) = v_{\mathcal{F}}(G')$ and $\tau_{\mathcal{F}}(G) = \tau_{\mathcal{F}}(G')$ which is smaller than G . In the first case we remove all vertices of the model and apply induction, while in the second case we are done by applying induction on G' . The reduction operation is described in Section 3. Section 4 is devoted to the proof that every *reduced* graph contains a constant-size T -model. Section 5 describes a number of algorithmic applications of (the proof of) Theorem 1.1.

We conclude the Introduction by mentioning two recent related results. First, Diestel, Kawarabayashi and Wollan [9] studied the $\mathcal{F} = \{K_t\}$ case and showed that, while $\tau_{\mathcal{F}}(G)$ cannot be bounded from above by a function of $v_{\mathcal{F}}(G)$ when $t \geq 5$, such a function exists for graphs with high enough vertex-connectivity compared to $v_{\mathcal{F}}(G)$ and t . There exists a function f such that, for every $t, k \geq 0$, every $(k(t - 3) + 14t + 14)$ -connected graph G either has an \mathcal{F} -packing of size k or an \mathcal{F} -transversal of size at most $f(k, t)$. Moreover, for fixed t the connectivity requirement is best possible up to an additive constant.

Second, Fomin, Saurabh and Thilikos [15] considered restricting the graphs G under consideration to belong to some fixed proper minor-closed family of graphs \mathcal{G} , and showed that if $\mathcal{F} = \{H\}$ with H a connected planar graph, then there exists a constant $c = c(\mathcal{F}, \mathcal{G})$ such that $\tau_{\mathcal{F}}(G) \leq c \cdot v_{\mathcal{F}}(G)$ for every $G \in \mathcal{G}$. Note that, as mentioned above, there is no such linear bound if we do not impose $G \in \mathcal{G}$.

² Indeed, by [26] graphs containing no member of \mathcal{F} as a minor have treewidth at most c for some constant $c = c(\mathcal{F})$ (since \mathcal{F} contains a planar graph). Thus $G - X$ has treewidth at most c for every \mathcal{F} -transversal X of G , and hence $\tau_{\mathcal{F}}(G) = \Omega(n)$. On the other hand, every subgraph of G containing a member of \mathcal{F} as a minor has a cycle, and thus has $\Omega(\log n)$ vertices, implying $v_{\mathcal{F}}(G) = O(n/\log n)$.

We remark that to obtain n -vertex graphs G with treewidth $\Omega(n)$ and girth $\Omega(\log n)$ it suffices to consider d -regular n -vertex expanders of girth $\Omega(\log n)$ for some fixed d , such as the Ramanujan graphs constructed by Lubotzky, Phillips and Sarnak [24], for instance. The fact that the treewidth is $\Omega(n)$ is a direct consequence of their positive vertex-expansion (or see [17, Proposition 1]).

2. Definitions and preliminaries

All graphs in this paper are finite, simple, and undirected, unless otherwise stated. Let $V(G)$ and $E(G)$ denote the vertex and edge sets of a graph G . Let $[i, j] := \{i, i + 1, \dots, j\}$.

A *tree decomposition* of a graph G is a pair $(T, \{B_x : x \in V(T)\})$ where T is a tree, and $\{B_x : x \in V(T)\}$ is a family of subsets of $V(G)$ (called *bags*) such that

- $\bigcup_{x \in V(T)} B_x = V(G)$,
- for every edge $uv \in E(G)$, there exists $x \in V(T)$ with $u, v \in B_x$, and
- for every vertex $u \in V(G)$, the set $\{x \in V(T) : u \in B_x\}$ induces a subtree of T .

The *width* of a tree decomposition $(T, \{B_x : x \in V(T)\})$ is $\max\{|B_x| - 1 : x \in V(T)\}$. The *treewidth* of G is the minimum width among all tree decompositions of G . See [8] for an introduction to the theory of treewidth. A *path decomposition* of G is a tree decomposition of G where the underlying tree T is a path. We will denote such a decomposition simply by the sequence B_1, \dots, B_p of its bags (in order). The *pathwidth* of G is the minimum width of a path decomposition of G .

A graph H is a *minor* of a graph G if H can be obtained from a subgraph of G by contracting edges. (Note that, since we only consider simple graphs, loops and parallel edges created during an edge contraction are deleted.) An *H -model* in G is a collection $M = \{S_x : x \in V(H)\}$ of vertex-disjoint connected subgraphs of G (called *branch sets*) such that, for every edge $xy \in E(H)$, some edge in G joins a vertex in S_x to a vertex in S_y . The *vertex set* of M is $V(M) := \cup\{V(S_x) : x \in V(H)\}$. Clearly, H is a minor of G if and only if G contains an H -model.

The following result of Bienstock, Robertson, Seymour and Thomas [2] will be routinely used in our proofs (see [7] for a short proof).

Theorem 2.1 ([2]). *Every graph with pathwidth at least $t - 1$ contains every forest on t vertices as a minor.*

Note that the bound in Theorem 2.1 is best possible.

We proceed with a few simple lemmas that will be used in our proof of Theorem 1.1. The following lemma is a special case of (8.6) in [26].

Lemma 2.2 ([26]). *Let P be a path and let $\mathcal{P}_1, \dots, \mathcal{P}_m$ be families of subpaths of P . Let $x_1, \dots, x_m \geq 0$ be integers, and let $k := x_1 + \dots + x_m$. Suppose that for each $i \in [1, m]$ there are k members of \mathcal{P}_i that are pairwise vertex-disjoint. Then there exist x_i members $P_1^i, \dots, P_{x_i}^i$ in \mathcal{P}_i for each $i \in [1, m]$ such that $P_1^1, \dots, P_{x_1}^1, P_1^2, \dots, P_{x_2}^2, \dots, P_1^m, \dots, P_{x_m}^m$ are all pairwise vertex-disjoint.*

The next lemma can be derived from the proof of (8.8) in [26]. We provide a proof for completeness.

Lemma 2.3. *Let \mathcal{F} be a set of $q \geq 1$ graphs, each with at most r components, and let $t \geq 1$. Then $\tau_{\mathcal{F}}(G) \leq 2qrt \cdot v_{\mathcal{F}}(G)$ for every graph G with pathwidth strictly less than t .*

Proof. The claim is trivially true if $v_{\mathcal{F}}(G) = 0$, so assume $v_{\mathcal{F}}(G) \geq 1$. Let $\mathcal{F} = \{H_1, \dots, H_q\}$. For each $i \in [1, q]$, let $H_{i,1}, \dots, H_{i,c_i}$ denote the components of H_i , where $c_i \leq r$. Let $s := v_{\mathcal{F}}(G) + 1$. Let B_1, \dots, B_p denote a path decomposition of G of width at most $t - 1$. Let P denote the path on p vertices with vertex set $\{B_1, \dots, B_p\}$ (in order).

For each $i \in [1, q]$ and $j \in [1, c_i]$, let $\mathcal{M}_{i,j}$ be the set of $H_{i,j}$ -models in G . Observe that for each $M \in \mathcal{M}_{i,j}$ we have that $V(M)$ induces a connected subgraph of G , since $H_{i,j}$ is connected; thus vertices in $V(M)$ appear in consecutive bags of the path decomposition, and hence M defines a corresponding subpath P_M of P . Notice that a sufficient condition for two models $M \in \mathcal{M}_{i,j}$ and $M' \in \mathcal{M}_{i',j'}$ being vertex-disjoint is that P_M and $P_{M'}$ are vertex-disjoint; this will be used below.

First suppose that there exists an index $i \in [1, q]$ such that, for each $j \in [1, c_i]$, there are sc_i pairwise vertex-disjoint paths in $\{P_M : M \in \mathcal{M}_{i,j}\}$. Then by Lemma 2.2 (with $m = c_i$, $k = sc_i$ and $x_1 = \dots = x_m = s$), one can find s $H_{i,j}$ -models $M_{i,j}^1, \dots, M_{i,j}^s$ in $\mathcal{M}_{i,j}$ for each $j \in [1, c_i]$ such that $M_{i,1}^1, \dots, M_{i,1}^s, M_{i,2}^1, \dots, M_{i,2}^s, \dots, M_{i,c_i}^1, \dots, M_{i,c_i}^s$ are all pairwise vertex-disjoint. In particular, G has s vertex-disjoint H_i -models, which implies $v_{\mathcal{F}}(G) \geq s$, a contradiction. Hence there is no such index i .

For each $i \in [1, q]$, consider an index $j \in [1, c_i]$ such that $\{P_M : M \in \mathcal{M}_{i,j}\}$ has no sc_i pairwise vertex-disjoint paths. Then, as is well known, one can find a subset of at most (in fact, strictly less than) sc_i vertices of P that meet all the paths in $\{P_M : M \in \mathcal{M}_{i,j}\}$; let X_i denote the union of the bags in the path decomposition corresponding to such a subset (thus $|X_i| \leq sc_i t$). Observe that $G - X_i$ has no $H_{i,j}$ -minor, and hence no H_i -minor. Therefore, $X_1 \cup \dots \cup X_q$ is an \mathcal{F} -transversal of G of size at most $sc_1 t + \dots + sc_q t \leq sqrt = qrt \cdot v_{\mathcal{F}}(G) + qrt \leq 2qrt \cdot v_{\mathcal{F}}(G)$. □

If \mathcal{F} is a set consisting of a single graph H , we simply write $v_H(G)$ and $\tau_H(G)$ for the invariants $v_{\mathcal{F}}(G)$ and $\tau_{\mathcal{F}}(G)$, respectively.

Suppose that \mathcal{F} is a finite set of graphs containing a forest F , and let T be a tree on $|F|$ vertices containing F . The next lemma allows us to reduce the task of finding a small \mathcal{F} -transversal of G to that of finding a small $\{T\}$ -transversal of G . This will be used in the proof of Theorem 1.1.

Lemma 2.4. *Let \mathcal{F} be a set of $q \geq 1$ graphs, each with at most r components, containing a forest F on t vertices. Let T be a tree on t vertices with $F \subseteq T$. Then*

$$v_T(G) \leq v_{\mathcal{F}}(G) \leq \tau_{\mathcal{F}}(G) \leq \tau_T(G) + 2qrt \cdot v_{\mathcal{F}}(G)$$

for every graph G .

Proof. The inequalities $v_T(G) \leq v_{\mathcal{F}}(G) \leq \tau_{\mathcal{F}}(G)$ are obvious; let us show that $\tau_{\mathcal{F}}(G) \leq \tau_T(G) + 2qrt \cdot v_{\mathcal{F}}(G)$. Let X be a $\{T\}$ -transversal of G with $|X| = \tau_T(G)$. Then $G - X$ has no T -minor, and hence has pathwidth at most $t - 2$ by Theorem 2.1. By Lemma 2.3, $G - X$ has an \mathcal{F} -transversal Y with $|Y| \leq 2qrt \cdot v_{\mathcal{F}}(G - X) \leq 2qrt \cdot v_{\mathcal{F}}(G)$. It follows that $X \cup Y$ is an \mathcal{F} -transversal of G of size at most $\tau_T(G) + 2qrt \cdot v_{\mathcal{F}}(G)$. □

As mentioned in the Introduction, it is not difficult to derive from Theorem 2.1 that, when \mathcal{F} is a finite set of graphs containing a forest, $\tau_{\mathcal{F}}(G) \in O_{\mathcal{F}}(k \log k)$ where $k = v_{\mathcal{F}}(G)$. We conclude this section with a proof of this statement, for completeness.

First we consider the case where \mathcal{F} consists of a single tree.

Lemma 2.5. *Let T be a tree on t vertices. Then $\tau_T(G) \leq 3(t + 1)k \log_2((t + 1)k) - t$ for every graph G with $v_T(G) = k$.*

Proof. The proof is by induction on k . If the pathwidth of G is at least $t(k + 1) - 1$, then G contains every forest on at most $t(k + 1)$ vertices as a minor, and thus contains $k + 1$ disjoint copies of T as a minor. This implies $v_T(G) \geq k + 1$, a contradiction. Hence, G has pathwidth at most $t(k + 1) - 2$.

Let B_1, \dots, B_p be a path decomposition of G of width at most $t(k + 1) - 2$. By modifying the decomposition if necessary, we may assume that $|B_i \Delta B_{i+1}| = 1$ for each $i \in [1, p - 1]$.

For each $i \in [1, p]$, let $L_i := G[(B_1 \cup \dots \cup B_{i-1}) - B_i]$ and $\ell_i := v_T(L_i)$, and similarly let $R_i := G[(B_{i+1} \cup \dots \cup B_p) - B_i]$ and $r_i := v_T(R_i)$. Observe that $\ell_i + r_i \leq k$, since L_i and R_i are vertex-disjoint subgraphs of G . Also, since $|B_i \Delta B_{i+1}| = 1$, we have $\ell_{i+1} \in \{\ell_i, \ell_i + 1\}$ for each $i \in [1, p - 1]$. Thus there exists an index $j \in [1, p]$ such that $\ell_j \leq \lceil \frac{k}{2} \rceil$ and $r_j \leq \lfloor \frac{k}{2} \rfloor$.

For the base case of the induction, namely $k = 1$, we have $\ell_j \leq 1$ and $r_j = 0$. If $\ell_j = 0$ then $Z := B_j$ is an $\{T\}$ -transversal of G of size at most $t(k + 1) - 1$. Otherwise, $\ell_j = 1$ and $j > 1$ (since $\ell_1 = 0$). We may further assume that j is chosen so that $\ell_{j-1} = 0$. Then $Z := B_{j-1} \cup B_j$ is an $\{T\}$ -transversal of G of size at most $t(k + 1)$ (since $|B_{j-1} \Delta B_j| = 1$). Thus in each case the transversal Z has size at most $t(k + 1) = 2t \leq 3(t + 1) \log_2(t + 1) - t$.

For the inductive step, assume $k \geq 2$. By induction there are $\{T\}$ -transversals X and Y of L_j and R_j , respectively, such that

$$|X| \leq 3(t + 1) \left\lceil \frac{k}{2} \right\rceil \log_2 \left((t + 1) \left\lceil \frac{k}{2} \right\rceil \right) - t \quad \text{and} \quad |Y| \leq 3(t + 1) \left\lfloor \frac{k}{2} \right\rfloor \log_2 \left((t + 1) \left\lfloor \frac{k}{2} \right\rfloor \right) - t.$$

Then $X \cup Y \cup B_j$ is an $\{T\}$ -transversal of G of size at most

$$\begin{aligned} & 3(t + 1) \left\lceil \frac{k}{2} \right\rceil \log_2 \left((t + 1) \left\lceil \frac{k}{2} \right\rceil \right) - t + 3(t + 1) \left\lfloor \frac{k}{2} \right\rfloor \log_2 \left((t + 1) \left\lfloor \frac{k}{2} \right\rfloor \right) - t + t(k + 1) - 1 \\ & \leq 3(t + 1)k \log_2(t + 1) + 3(t + 1)k \log_2 \left(\frac{k+1}{2} \right) + tk - t - 1 \\ & = 3(t + 1)k \log_2(t + 1) + 3(t + 1)k \log_2(k + 1) - 2tk - 3k - t - 1 \\ & \leq 3(t + 1)k \log_2(t + 1) + 3(t + 1)k \log_2 \left(\frac{3k}{2} \right) - 2tk - 3k - t - 1 \\ & \leq 3(t + 1)k \log_2((t + 1)k) - t, \end{aligned}$$

as desired. □

Combining Lemmas 2.4 and 2.5, we obtain the aforementioned $O_{\mathcal{F}}(k \log k)$ bound.

Proposition 2.6. *Let \mathcal{F} be a set of $q \geq 1$ graphs, each with at most r components, containing a forest on t vertices. Then $\tau_{\mathcal{F}}(G) \leq 3(t + 1)k \log_2((t + 1)k) + 2qrk - t$ for every graph G with $v_{\mathcal{F}}(G) = k$.*

3. A reduction operation

First we need to introduce some definitions. A *rooted graph* is a pair (G, R) where G is a graph and $R = (v_1, \dots, v_k)$ is a (possibly empty) ordered subset of vertices of G called *roots*. The notions of minors and models generalize in a natural way to rooted graphs. A rooted graph (H, R') with $R' = (w_1, \dots, w_\ell)$ is a *minor* of (G, R) if $k = \ell$ and there exists a collection $\{V_u : u \in V(H)\}$ of disjoint vertex subsets of G (the *branch sets*), each inducing a connected subgraph in G , such that $v_i \in V_{w_i}$ for each $i \in [1, k]$, and there exists an edge between V_u and V_v in G for every $uv \in E(H)$. The collection $M := \{V_u : u \in V(H)\}$ is called a *model* of (H, R') in (G, R) . We write $V(M)$ to denote the set $\cup_{u \in V(H)} V_u$ of vertices of the model M .

Isomorphism between rooted graphs is defined in the expected way (the i th root of the first graph is required to be mapped to the i th root of the second). We write $R' \subseteq R$ for an ordered subset of an ordered set R (thus, the ordering of R' is consistent with the ordering of R). Also, if X is an (unordered) set then $R - X$ is ordered as in R .

For our reduction operation we will need to keep track of rooted minors of bounded sizes of a given rooted graph (G, R) . For an integer $q \geq 1$, the q -*folio* of (G, R) is the $2^{|R|}$ -tuple which, for each set $R' \subseteq R$, records the set of all minors (H, R') of (G, R) with $|H| \leq q$ (keeping one member per isomorphism class). For integers $p \geq 0$ and $q \geq 1$, the p -*deletion q -folio* of (G, R) is the $(p + 1)$ -tuple which, for each $i \in [0, p]$, records for each $X \subseteq R$ with $|X| \leq i$ the set of all q -folios of $(G - (X \cup Y), R - X)$ taken over all subsets $Y \subseteq V(G) - R$ with $|X| + |Y| = i$. (Informally speaking, we record the different q -folios obtained by deleting i vertices in or outside the set of roots, for each $i \in [0, p]$.)

The following lemma on p -deletion q -folios is an adaptation of Lemma 2.2 in [18], its proof relies on standard monadic second-order logic techniques.

Lemma 3.1. *Let $p, r, z \geq 0$ and $q \geq 1$. Then there exists a computable function $g(p, q, r, z)$ such that, for every rooted graph (G, R) such that G has treewidth at most z and $|R| = r$, there exists a rooted graph (G', R) with $|G'| \leq g(p, q, r, z)$ and such that (G, R) and (G', R) have the same p -deletion q -folio.*

Proof. Fix an ordered set R with $|R| = r$. Let $\mathcal{G}(R, q)$ denote the set of (non-isomorphic) rooted graphs (G, R) with $|G| \leq q$. Clearly $|\mathcal{G}(R, q)|$ is bounded from above by a function of r and q .

Let $\mathcal{S}(R, q)$ denote the set of vectors s having one entry $s(R')$ per subset $R' \subseteq R$ and such that $s(R') \subseteq \mathcal{G}(R', q)$. Again, $|\mathcal{S}(R, q)|$ is bounded from above by a function of r and q . By definition, for each rooted graph (G, R) there is exactly one vector in $\mathcal{S}(R, q)$ that encodes the q -folio of (G, R) . (Note that, on the other hand, some vectors $s \in \mathcal{S}(R, q)$ might not correspond to the q -folio of any rooted graph (G, R) .)

Let $\mathcal{D}(R, q, p)$ denote the set of vectors $d = (D_0, \dots, D_p)$ such that, for each $i \in [0, p]$, D_i has one entry $D_i(X)$ for each $X \subseteq R$ with $|X| \leq i$, and $D_i(X) \subseteq \mathcal{S}(R - X, q)$ for each such set X . The size of $\mathcal{D}(R, q, p)$ is bounded from above by a function of r, q , and p . By the definition of $\mathcal{D}(R, q, p)$, the p -deletion q -folio of a rooted graph (G, R) is encoded by exactly one vector in $\mathcal{D}(R, q, p)$.

It is well known that the property that a fixed graph H is a minor of a graph G can be expressed by a monadic second-order logic (MSO) sentence φ_H (see Grohe [16] for an introduction to monadic second-order logic). Similarly, for each $s \in \mathcal{S}(R, q)$ one can write an MSO-sentence φ_s such that a rooted graph (G, R) satisfies φ_s if and only if its q -folio equals s . More generally, for each $d \in \mathcal{D}(R, q, p)$ one can express by an MSO-sentence φ_d the property that d is the p -deletion q -folio of (G, R) .

By Seese’s theorem [29], there is an algorithm which given an MSO-sentence φ and an integer t decides whether there exists a graph of treewidth at most t satisfying φ . It is also known that there is a computable function $h(\varphi, t)$ such that if an MSO-sentence φ is realizable by (or more accurately, admits a model consisting of) a graph of treewidth at most t , then there is such a graph with at most $h(\varphi, t)$ vertices. This can be derived using (among others) a theorem of Thatcher and Wright [30], as explained in the proof of Lemma 2.2 in [18].

Now clearly the set $\mathcal{D}(R, q, p)$ is computable. It follows that, for each $d \in \mathcal{D}(R, q, p)$, we can use Seese’s algorithm to compute the subset $\mathcal{D}_z \subseteq \mathcal{D}(R, q, p)$ of vectors $d \in \mathcal{D}(R, q, p)$ such that there exists a rooted graph (G, R) of treewidth at most z that satisfies φ_d . Moreover, for each $d \in \mathcal{D}_z$, we can compute a rooted graph (G', R) whose p -deletion q -folio equals d and such that $|G'| \leq h(\varphi_d, z)$. The lemma follows by setting $g(p, q, r, z) := \max\{h(\varphi_d, z) : d \in \mathcal{D}_z\}$. □

A *separation* of a graph G is a pair (G_1, G_2) of two induced subgraphs of G such that $G = G_1 \cup G_2$, its *order* is $|V(G_1) \cap V(G_2)|$. Let us point out that $G_1 \subseteq G_2$ or $G_2 \subseteq G_1$ is allowed in this definition, that is, we do not require $V(G_1) - V(G_2)$ and $V(G_2) - V(G_1)$ to be non-empty.

Lemma 3.2. *For every tree T there exists a computable function b such that, for every graph G having a separation (G_1, G_2) of order t with $|G_1| \geq b(t)$ such that G_1 has no T -minor, there exists a graph G' satisfying $v_T(G') = v_T(G)$, $\tau_T(G') = \tau_T(G)$, and $|G'| < |G|$.*

Proof. Since G_1 has no T -minor, its pathwidth is at most $|T| - 2$ by Theorem 2.1, and hence so is its treewidth. We will prove the lemma with $b(t) := g(t, q, t, |T| - 2) + 1$, where $q := t(|T| + 1)$ and g is the function from Lemma 3.1. Let $X := V(G_1) \cap V(G_2)$ and let R be an arbitrary ordering of X . Using Lemma 3.1, let (G'_1, R) be a rooted graph with the same t -deletion q -folio as (G_1, R) , and with $|G'_1| \leq g(t, q, t, |T| - 2) < b(t) \leq |G_1|$. Assume without loss of generality that $V(G'_1) \cap V(G_2) = X$. Let G' be the graph obtained from $G'_1 \cup G_2$ by removing every edge linking two vertices in X that is included in G_2 but not in G'_1 . (Thus $G[X]$ and $G'[X]$ could possibly be distinct graphs.) Then $|G'| < |G|$, and it remains to show that $v_T(G') = v_T(G)$ and $\tau_T(G') = \tau_T(G)$.

First we prove that $v_T(G') = v_T(G)$ by showing that for $(K, L) \in \{(G, G'), (G', G)\}$ we have $v_T(L) \geq v_T(K)$. If $K = G$, let $K_1 := G_1$ and $L_1 := G'_1$, otherwise let $K_1 := G'_1$ and $L_1 := G_1$. Consider $k := v_T(K)$ vertex-disjoint T -models M_1, \dots, M_k in K , where $M_i = \{V_u^i : u \in V(T)\}$ for each $i \in [1, k]$. We may assume that these models are ordered so that $V(M_1), \dots, V(M_r)$ each intersects $V(K_1)$, while none of $V(M_{r+1}), \dots, V(M_k)$ intersects

$V(K_1)$, for some index $\ell \in [1, k]$. Observe that, since T is connected and K_1 has no T -minor, each of $V(M_1), \dots, V(M_\ell)$ must intersect X . It follows that $\ell \leq t$.

For each $i \in [1, \ell]$, let $X_i := V(M_i) \cap X$ (thus X_i and X_j are disjoint for $i \neq j$), let A_i be the set of vertices $u \in V(T)$ such that V_u^i intersects X_i , and, for every $u \in A_i$, let $s_{u,i} := |V_u^i \cap X_i|$ and arbitrarily partition $V_u^i \cap V(K_1)$ into $s_{u,i}$ parts $W_{u,1}^i, \dots, W_{u,s_{u,i}}^i$, so that each part induces a connected subgraph of K_1 and contains exactly one vertex from $V_u^i \cap X_i$. Also let B_i be the set of vertices $u \in V(T)$ such that V_u^i intersects $V(K_1)$ but not X_i (thus A_i and B_i are disjoint). For $u \in B_i$, let $s_{u,i} := 1$ and let $W_{u,1}^i := V_u^i$.

Let H^* be the graph with one vertex per triple (i, u, j) where $i \in [1, \ell], u \in A_i \cup B_i$, and $j \in [1, s_{u,i}]$, and where two distinct vertices (i, u, j) and (i', u', j') are adjacent if and only if there is an edge in K_1 connecting a vertex of $W_{u,j}^i$ to a vertex of $W_{u',j'}^{i'}$. Observe that V_u^i corresponds to one vertex of H^* if $V_u^i \subseteq V(K_1) - X$, and to s_u vertices of H^* otherwise. Let R^* be the ordering of $X_1 \cup \dots \cup X_\ell$ induced by R . By the definition of H^* , the rooted graph (H^*, R^*) is a minor of (K_1, R^*) . Since

$$|H^*| = \sum_{i=1}^{\ell} \sum_{u \in A_i} s_{u,i} + \sum_{i=1}^{\ell} |B_i| = \sum_{i=1}^{\ell} |X_i| + \sum_{i=1}^{\ell} |B_i| \leq |X| + \ell|T| \leq t + t|T| = q,$$

(H^*, R^*) is in the q -folio of (K_1, R) . (More precisely, the R^* -entry of the q -folio of (K_1, R) contains (H^*, R^*) .) Hence (H^*, R^*) is also in the q -folio of (L_1, R) , since the q -folios of (K_1, R) and (L_1, R) are the same (up to isomorphism, as always). Now consider a model of (H^*, R^*) in (L_1, R^*) , and let $Y_{u,j}^i$ denote the branch set corresponding to the vertex (i, u, j) of H^* . Define, for each $i \in [1, k]$ and $u \in V(T)$, the set Z_u^i as follows:

$$Z_u^i := \begin{cases} (V_u^i - V(K_1)) \cup Y_{u,1}^i \cup \dots \cup Y_{u,s_{u,i}}^i & \text{if } i \leq \ell \text{ and } u \in A_i \cup B_i, \\ V_u^i & \text{otherwise.} \end{cases}$$

For each $i \in [1, k]$, we have that $L[Z_u^i]$ is connected for each $u \in V(T)$, which follows from the fact that $K[V_u^i]$ is connected. Also, there exists an edge between Z_u^i and Z_v^i in L for each $uv \in E(T)$. Let $M'_i := \{Z_u^i : u \in V(T)\}$. Since Z_u^i is disjoint from Z_v^j for distinct $i, j \in [1, k]$ and $u, v \in V(T)$, we deduce that M'_1, \dots, M'_k are k vertex-disjoint models of T in L . Hence $v_T(K) \leq v_T(L)$, as desired.

Now we show that $\tau_T(G') = \tau_T(G)$ by showing that for $(K, L) \in \{(G, G'), (G', G)\}$ we have $\tau_T(L) \geq \tau_T(K)$. As before, if $K = G$, let $K_1 := G_1$ and $L_1 := G'_1$; otherwise let $K_1 := G'_1$ and $L_1 := G_1$.

Let J be a minimum-size transversal of L , and let $J_1 := J \cap V(L_1)$. We have $|J_1| \leq |X| = t$, because $(J - J_1) \cup X$ is also a transversal of L , and thus must have size at least that of J . (Here we use the fact that L_1 has no T -minor, and that T is connected.) Since the t -deletion q -folios of (K_1, R) and (L_1, R) are the same, there exists a subset I_1 of vertices of K_1 with $|I_1| = |J_1|$ such that the q -folios of $(L_1 - J_1, R - J_1)$ and $(K_1 - I_1, R - I_1)$ are the same (by definition).

We claim that the set $I := (J - J_1) \cup I_1$, which has the same size as J , is a transversal of K . Arguing by contradiction, assume not, and let $M := \{V_u : u \in V(T)\}$ denote a T -model in $K - I$. Let A be the subset of vertices $u \in V(T)$ such that V_u intersects X , and let B be the set of those such that $V_u \subseteq V(K_1) - X$ (thus A and B are disjoint). For

every $u \in A$, let $s_u := |V_u \cap X|$, and partition $V_u \cap V(K_1)$ into s_u parts $W_{u,1}, \dots, W_{u,s_u}$, so that each part induces a connected subgraph of K_1 and contains exactly one vertex from $V_u \cap X$. Further, let $s_u := 1$ and $W_{u,1} := V_u$ for every $u \in B$. We define H^* as the minor of $K_1 - I_1$ modelled by the $W_{u,j}$ s, exactly as before: H^* has one vertex for every couple (u, j) where $u \in A \cup B$ and $j \in \{1, \dots, s_u\}$, and two distinct vertices (u, j) and (u', j') of H^* are adjacent if and only if there is an edge between $W_{u,j}$ and $W_{u',j'}$ in $K_1 - I_1$.

Let R^* be the ordered subset of R induced by the vertices in $V(M) \cap X$. Thus (H^*, R^*) is a minor of $(K_1 - I_1, R^*)$. Since $|H^*| = \sum_{u \in A} s_u + |B| \leq |X| + |T| \leq q$, the rooted graph (H^*, R^*) is in the q -folio of $(K_1 - I_1, R - I_1)$, and hence also in the q -folio of $(L_1 - J_1, R - J_1)$. Consider a model of (H^*, R^*) in $(L_1 - J_1, R^*)$; let $W'_{u,j}$ denote the branch set corresponding to the vertex (u, j) of H^* in that model. For every $u \in V(T)$, let

$$V'_u := \begin{cases} (V_u - V(K_1)) \cup W'_{u,1} \cup \dots \cup W'_{u,s_u} & \text{if } u \in A_i \cup B_i, \\ V_u & \text{otherwise.} \end{cases}$$

It can be checked that V'_u is disjoint from V'_v for every distinct $u, v \in V(T)$, that $L[V'_u]$ is connected for every $u \in V(T)$, and that there is an edge between V'_u and V'_v in L for every $uv \in E(T)$. We deduce that $M' := \{V'_u : u \in V(T)\}$ is a model of T in $L - J$, a contradiction. Therefore, the set I is a transversal of K , and $\tau_T(K) \leq \tau_T(L)$, as claimed. \square

Even though we will not use this fact, we would like to point out that Lemma 3.2 remains true more generally if the tree T is replaced by a connected planar graph H .

Lemma 3.3. *For every connected planar graph H there exists a computable function b such that, for every graph G having a separation (G_1, G_2) of order t with $|G_1| \geq b(t)$ such that G_1 has no H -minor, there exists a graph G' satisfying $v_H(G') = v_H(G)$, $\tau_H(G') = \tau_H(G)$, and $|G'| < |G|$.*

This can be shown using exactly the same proof as for Lemma 3.2, the only difference being that G_1 now has treewidth at most c_H for some constant c_H (by [26]), and thus the function $b(t)$ should now be defined as $b(t) := g(t, q, t, c_H) + 1$.

4. Proof of Main Theorem

Let us start by recalling the overview of the proof of Theorem 1.1 given in the Introduction for the special case where \mathcal{F} consists of a single tree T . The proof is by induction on $|G|$; we show that one can always either apply the reduction operation described in Lemma 3.2 (in which case we are done by induction), or find a T -model in G of size at most c . In the latter case, we delete all vertices of the model and apply induction on the remaining graph. The key step is thus proving that T -models of constant size can be found in G when the reduction operation cannot be applied. This is achieved by Lemma 4.6, which is the main lemma of this section.

Our basic strategy for finding a T -model of constant size when the reduction operation cannot be applied to G is to consider a ball B of large but constant radius around an arbitrary vertex of G . If all vertices in B have bounded degree then B has bounded size,

and it turns out that it is not difficult to show that B must contain T as a minor (in which case we found a T -model of constant size) because otherwise the reduction operation could have been applied to G . However, difficulties arise when B contains vertices of high degrees. This is why in Lemma 4.6 we do not directly try and find a T -model of constant size, but instead look for a subgraph of G of constant size having some specified pathwidth t . This is enough, since by Theorem 2.1 every graph with pathwidth at least $|T| - 1$ contains T as a minor, and allows us to set up an inductive argument that handles vertices of high degree in G .

First we need to introduce a few definitions and results. Lemmas 4.1, 4.2, 4.3, 4.4, and 4.5 below are short lemmas that will be used in the proof of Lemma 4.6.

Given a graph G and $t \in \mathbb{N}$, a pw - t -separation of G is a separation (G_1, G_2) of G such that G_1 has a path decomposition B_1, \dots, B_p of width at most t with $V(G_1) \cap V(G_2) = B_1 \cup B_p$. Notice that the order of such a separation is at most $2t + 2$.

Lemma 4.1. *Let G be a graph of pathwidth at most t with $k \geq 0$ marked vertices. Then there is a pw - t -separation (G_1, G_2) of G with*

$$|G_1| \geq \frac{|G| - k(t + 1)}{k + 1}$$

such that no vertex in $V(G_1) - V(G_2)$ is marked.

Proof. Consider a path decomposition B_1, \dots, B_p of G with width at most t . If $k = 0$, then $G_1 := G$ and $G_2 := G[B_1 \cup B_p]$ defines the desired pw - t -separation (G_1, G_2) of G . Now assume $k \geq 1$.

For each marked vertex v , choose an arbitrary bag that contains v and mark it. Let m_1, \dots, m_ℓ denote the indices of the marked bags. Observe that $1 \leq \ell \leq k$ (note that $\ell < k$ if some bag is marked more than once). Also let $m_0 := 1$ and $m_{\ell+1} := p$.

For $1 \leq i < j \leq p$, let $Y_{i,j} := (B_{i+1} \cup \dots \cup B_{j-1}) - (B_i \cup B_j)$. Observe that $Y_{i,j}$ is disjoint from $Y_{i',j'}$ when $j \leq i'$ (this is a consequence of the axioms of path decompositions). In particular, the $\ell + 1$ sets $Y_{1,m_1}, Y_{m_1,m_2}, \dots, Y_{m_{\ell-1},m_\ell}, Y_{m_\ell,p}$ are pairwise disjoint. Hence, among them there is a set $Y_{a,b}$ satisfying

$$|Y_{a,b}| \geq \frac{\sum_{j=0}^{\ell} |Y_{m_j,m_{j+1}}|}{\ell + 1} \geq \frac{|G| - \sum_{j=1}^{\ell} |B_{m_j}|}{\ell + 1} \geq \frac{|G| - \ell(t + 1)}{\ell + 1} \geq \frac{|G| - k(t + 1)}{k + 1}.$$

On the other hand, none of $Y_{1,m_1}, Y_{m_1,m_2}, \dots, Y_{m_{\ell-1},m_\ell}, Y_{m_\ell,p}$ contains a marked vertex (as follows again from the axioms of path decompositions). Thus $G_1 := G[Y_{a,b} \cup B_a \cup B_b]$ and $G_2 := G[V(G) - Y_{a,b}]$ define a separation (G_1, G_2) of G with $V(G_1) \cap V(G_2) = B_a \cup B_b$ such that

$$|G_1| \geq \frac{|G| - k(t + 1)}{k + 1}$$

and $V(G_1) - V(G_2)$ ($= Y_{a,b}$) has no marked vertex. Since B_a, B_{a+1}, \dots, B_b is a path decomposition of G_1 with width at most t , we deduce that (G_1, G_2) is the desired pw - t -separation of G . □

Lemma 4.2. *Let (G_1, G_2) be a pw- t -separation of a graph G with $k \geq 0$ marked vertices, and let ℓ be an integer satisfying $1 \leq \ell \leq \lceil |G_1|/(k + 1) \rceil$. Then there exists a pw- t -separation (G'_1, G'_2) of G with $G'_1 \subseteq G_1$, $G_2 \subseteq G'_2$, and $|G'_1| = \ell$, such that no vertex in $V(G'_1) - V(G'_2)$ is marked.*

Proof.

Case 1: $\ell \leq t + 1$. Let Y be an arbitrary subset of ℓ vertices of G_1 , let $G'_1 := G[Y]$ and let $G'_2 := G$. Then $V(G'_1) - V(G'_2)$ is empty and G'_1 has a trivial path decomposition of width at most t consisting of a single bag $B_1 := Y$. Hence (G'_1, G'_2) is a pw- t -separation of G as desired.

Case 2: $\ell > t + 1$. Let B_1, \dots, B_p be a path decomposition of G_1 of width at most t such that $B_1 \cup B_p = V(G_1) \cap V(G_2)$. We may assume that $|B_i \triangle B_{i+1}| = 1$ for each $i \in [1, p - 1]$, by modifying the path decomposition if necessary (observe that this can be done without changing the first and last bags).

Mark bags B_1, B_p , and for each marked vertex v mark an arbitrarily chosen bag that contains v (thus the same bag can be chosen several times). Let q be the total number of marked bags, and let m_1, \dots, m_q denote the indices of these bags, in order. (Thus $m_1 = 1$ and $m_q = p$.) Since $|B_1| \leq t + 1$ and $\ell > t + 1$, it follows that $p \geq 2$, and hence $q \geq 2$. Let $Y_{a,b} := B_a \cup B_{a+1} \cup \dots \cup B_b$ for $1 \leq a < b \leq p$. Since $Y_{m_1, m_2} \cup \dots \cup Y_{m_{q-1}, m_q} = V(G_1)$ and $q - 1 \leq k + 1$, there exists $j \in [1, q - 1]$ such that

$$|Y_{m_j, m_{j+1}}| \geq \left\lceil \frac{|G_1|}{q - 1} \right\rceil \geq \left\lceil \frac{|G_1|}{k + 1} \right\rceil \geq \ell.$$

Now, since $|B_i \triangle B_{i+1}| = 1$ for each $i \in [1, p - 1]$ and since $\ell > |B_{m_j}|$, there exists an index $b \in [m_j + 1, m_{j+1}]$ such that $|Y_{m_j, b}| = \ell$. It follows from the axioms of a path decomposition that no vertex in $Y_{m_j, b} - (B_{m_j} \cup B_b)$ is marked. Therefore, setting $G'_1 := G[Y_{m_j, b}]$ and $G'_2 := G - (Y_{m_j, b} - (B_{m_j} \cup B_b))$, and using that B_{m_j}, \dots, B_b is a path decomposition of G'_1 of width at most t , we deduce that (G'_1, G'_2) is a pw- t -separation of G with the required properties. □

With a slight abuse of notation, we simply write (G, v) to denote the rooted graph (G, R) where R consists only of the vertex v . The *height* of a rooted tree (T, v) is the maximum length of a path from v to a leaf, where the length of a path is the number of its edges. (Thus the height of (T, v) is zero if $|T| = 1$.) For $u \in V(T)$, the notions of *parent*, *ancestors*, *children*, *descendants* of u in (T, v) are defined as expected. When the root v is clear from the context, we simply denote by T_u the subtree of (T, v) induced by u and all its descendants. For $h \geq 1$, the *complete binary tree of height h* , denoted \mathcal{B}_h , is the unique rooted tree of height h where the root has degree 2, every other non-leaf vertex has degree 3, and every path from the root to a leaf has length exactly h . It is known, and not difficult to prove, that the pathwidth of \mathcal{B}_h equals h .

Lemma 4.3. *Let (T, v) be a rooted tree of height h and maximum degree Δ . If (T, v) does not contain \mathcal{B}_{k+1} as a (rooted) minor, then $|T| \leq (h + 1)^{k+1}(\Delta + 1)^{k+1}$.*

Proof. The proof is by induction on k . If $k = 0$, then the tree T is a path, and thus satisfies $|T| = h + 1 \leq (h + 1)(\Delta + 1)$. For the inductive step, assume $k \geq 1$. We may assume that \mathcal{B}_k is a minor of (T, v) , because otherwise we are done by induction. Let $u \in V(T)$ be a vertex such that (T_u, u) contains \mathcal{B}_k as a minor and u is at maximum distance from v in T . Let q denote that distance. Let u_1, \dots, u_ℓ denote the children of u . Observe that $\ell \geq 2$ since (T_u, u) has a \mathcal{B}_k -minor, $k \geq 1$, and q is maximum. Let v_1, \dots, v_{q+1} be the vertices on the path from v to u in T , in order. Thus $v_1 = v$ and $v_{q+1} = u$. For each $i \in [1, q]$, let $w_{i,1}, \dots, w_{i,a_i}$ denote the children of v_i in T that are distinct from v_{i+1} (note that possibly $a_i = 0$).

We claim that, for each $i \in [1, q]$ and $j \in [1, a_i]$, the rooted tree $(T_{w_{i,j}}, w_{i,j})$ has no \mathcal{B}_k minor. Indeed, otherwise a model of \mathcal{B}_k in $(T_{w_{i,j}}, w_{i,j})$ could be combined with a model of \mathcal{B}_k in $(T_{v_{i+1}}, v_{i+1})$ (which exists, by the definition of u) and the v_1-v_i path in T to give a model of \mathcal{B}_{k+1} in (T, v) , a contradiction. By induction, we thus have $|T_{w_{i,j}}| \leq h^k(\Delta + 1)^k$, since the height of $(T_{w_{i,j}}, w_{i,j})$ is at most $h - 1$.

By the definition of u , the rooted tree (T_{u_i}, u_i) has no \mathcal{B}_k minor either for each $i \in [1, \ell]$. (Recall that u_1, \dots, u_ℓ are the children of u .) Hence $|T_{u_i}| \leq h^k(\Delta + 1)^k$ by induction. It follows that

$$\begin{aligned} |T| &= q + 1 + \sum_{i=1}^q \sum_{j=1}^{a_i} |T_{w_{i,j}}| + \sum_{i=1}^{\ell} |T_{u_i}| \\ &\leq h + 1 + \left[\sum_{i=1}^q \sum_{j=1}^{a_i} h^k(\Delta + 1)^k \right] + \ell h^k(\Delta + 1)^k \\ &\leq h^k(\Delta + 1)^k + h(\Delta + 1)h^k(\Delta + 1)^k + \Delta h^k(\Delta + 1)^k \\ &= h^{k+1}(\Delta + 1)^{k+1} + h^k(\Delta + 1)^{k+1} \\ &\leq (h + 1)^{k+1}(\Delta + 1)^{k+1}, \end{aligned}$$

as desired. (In the second inequality, we used that $h + 1 \leq h^k(\Delta + 1)^k$, which follows from the fact that $h, k \geq 1$ and $\Delta \geq 2$.) □

Lemma 4.4. *If (T, v) is a rooted tree on t vertices and (G, w) is a connected rooted graph with pathwidth at least $2t - 2$, then (G, w) contains (T, v) as a minor.*

Proof. Let (T', v') be a copy of (T, v) . Let T^* be the (unrooted) tree obtained from the disjoint union of T and T' by identifying v and v' . Thus T^* has $2t - 1$ vertices. Since G has pathwidth at least $2t - 2$, it contains a T^* -minor (by Theorem 2.1), and hence has a model of T^* . Since G is connected, we may assume that every vertex of G is in a branch set of this model. Let B_v denote the branch set of vertex $v = v'$. Exchanging T and T' if necessary, we may assume that the root w of (G, w) is in a branch set corresponding to a vertex from T . Thus we can find a path P in G with one endpoint being w and the other endpoint in B_v such that P avoids all branch sets of vertices in $V(T') - \{v'\}$. Now, replacing the branch set B_v by $B_v \cup V(P)$ and taking all branch sets corresponding to vertices of T' distinct from v' , we obtain a model of (T', v') in (G, w) . Therefore, (G, w) contains (T, v) as a minor. □

The following lemma is well known; a proof is included for completeness (see Corollary 3.1 in [10] for a similar result).

Lemma 4.5. *Let G_1, G_2, G_3 be three connected graphs, each with pathwidth at least k , and let v_i be an arbitrary vertex of G_i , for $i = 1, 2, 3$. Let G be the graph obtained from the disjoint union of G_1, G_2, G_3 by adding a new vertex v adjacent to v_1, v_2, v_3 . Then G has pathwidth at least $k + 1$.*

Proof. Arguing by contradiction, assume G has a path decomposition B_1, \dots, B_p of width at most k . For every $w \in V(G)$, let $[\ell_w, r_w]$ denote the interval of indices j such that $w \in B_j$. For each $i \in [1, 3]$, let $I_i := \cup\{[\ell_w, r_w] : w \in V(G_i)\}$. Since G_i is connected, I_i is again an interval; we denote by ℓ_i and r_i its left and right endpoints, respectively. Re-indexing G_1, G_2, G_3 if necessary, we may assume that $\ell_1 \leq \ell_2$ and $r_2 \leq r_3$. Since $vv_i \in E(G)$, it follows that $I_i \cap [\ell_v, r_v] \neq \emptyset$ for each $i \in [1, 3]$. This implies that $I_1 \cup I_3 \cup [\ell_v, r_v]$ is again an interval, and this interval contains I_2 . Hence, every bag B_j ($j \in [1, p]$) that includes a vertex of $V(G_2)$ contains also at least one vertex from $V(G) - V(G_2)$. Therefore, $B_1 \cap V(G_2), \dots, B_p \cap V(G_2)$ is a path decomposition of G_2 of width at most $k - 1$, a contradiction. □

The next lemma is our main tool. Informally, it shows that every connected graph G with pathwidth at least t and no pw- $(t - 1)$ -separation (G_1, G_2) of G with $|G_1|$ ‘big’ has a connected subgraph G' of constant size with pathwidth at least t . It turns out that a slightly stronger statement is easier to prove by induction, namely that for every vertex w of G one can find such a subgraph G' containing it. We note that the vertex w will be an arbitrary vertex when this lemma is used in the proof of Theorem 1.1.

Lemma 4.6. *There exists a computable function $f : \mathbb{N} \times \mathbb{N} \rightarrow \mathbb{N}$ such that for every $t, r \in \mathbb{N}$, every connected graph G of pathwidth at least t , and every vertex $w \in V(G)$, at least one of the following holds.*

- (a) *There exists a connected subgraph G' of G with $w \in V(G')$ such that G' has pathwidth at least t and $|G'| \leq f(t, r)$.*
- (b) *There exists a pw- $(t - 1)$ -separation (G_1, G_2) of G with $|G_1| \geq r$.*

Proof. The proof is by induction on t . For $t = 0$, the vertex w itself provides a connected subgraph G' of G of pathwidth 0, and thus the claim holds with $f(0, r) := 1$. Now assume $t \geq 1$. If $r \leq 1$, letting G_1 be the graph induced by an arbitrary vertex of G and letting $G_2 := G$, we obtain a pw- $(t - 1)$ -separation (G_1, G_2) of G with $|G_1| \geq r$. Thus we may also assume $r \geq 2$.

Let

$$\begin{aligned}
 r_1 &:= (r + t - 1) + (r + t)2t + r, \\
 r_2 &:= (r + t - 1)(1 + f(t - 1, r_1)) + (r + t)2t + r, \\
 r_3 &:= (r + t - 1)(1 + f(t - 1, r_1) + f(t - 1, r_2)) + (r + t)2t + r.
 \end{aligned}$$

Let Δ, ε , and d be constants defined as follows:

$$\Delta := (r + t) \left(\sum_{i=1}^3 f(t - 1, r_i) + 2t + 1 \right) + r - 1,$$

$$\varepsilon := \frac{1}{2r + t},$$

$$d := \max \left\{ \frac{2t \ln(\Delta + 1)}{\ln(1 + \varepsilon)}, \left(\frac{2t}{\ln(1 + \varepsilon)} \right)^2 \right\}.$$

We will prove the claim with

$$f(t, r) := \lceil \max\{\Delta^{d+1}, \Delta + d + 1\} \rceil.$$

(We remark that the ceiling is there only to ensure that $f(t, r)$ is an integer.) We may assume that $|G| > f(t, r)$, since otherwise we are done with $G' = G$.

Case 1: Every vertex at distance at most d from w has degree at most Δ . For $i = 0, 1, \dots, d$, let H_i be the subgraph of G induced by all vertices at distance at most i from w , and let J_i be an arbitrary breadth-first search tree of H_i from w . Since $|H_d| \leq \Delta^{d+1} \leq f(t, r) < |G|$, we deduce that $V(H_i) - V(H_{i-1})$ is not empty for each $i \in [1, d]$. The graph H_d is connected, includes the vertex w , and has at most $f(t, r)$ vertices. Thus we are done if the pathwidth of H_d is at least t . So let us assume that H_d has pathwidth at most $t - 1$ (and thus, in particular, J_d has pathwidth at most $t - 1$).

The tree \mathcal{B}_t , which has pathwidth t , cannot be a minor of (J_d, w) . Hence we have

$$|H_d| = |J_d| \leq (d + 1)^t (\Delta + 1)^t \tag{4.1}$$

by Lemma 4.3.

By the definition of d ,

$$d \geq \frac{d}{2} + \frac{\sqrt{d} \ln(d + 1)}{2} \geq \frac{t \ln(\Delta + 1)}{\ln(1 + \varepsilon)} + \frac{t \ln(d + 1)}{\ln(1 + \varepsilon)},$$

which implies $(1 + \varepsilon)^d \geq (d + 1)^t (\Delta + 1)^t$. Thus, if $|H_i| > (1 + \varepsilon)|H_{i-1}|$ for each $i \in [1, d]$, then

$$|H_d| > (1 + \varepsilon)^d \geq (d + 1)^t (\Delta + 1)^t,$$

contradicting (4.1). Hence there exists $j \in [1, d]$ such that $|H_j| \leq (1 + \varepsilon)|H_{j-1}|$. Mark all $k := |H_j| - |H_{j-1}|$ vertices of H_j that are in $V(H_j) - V(H_{j-1})$. We have $k \geq 1$, since $V(H_j) - V(H_{j-1})$ is not empty, and also $k \leq \varepsilon|H_{j-1}| \leq \varepsilon|H_j|$. Since the pathwidth of H_j is at most that of H_d , by Lemma 4.1 there is a pw- $(t - 1)$ -separation (L_1, L_2) of H_j such that no vertex in $V(L_1) - V(L_2)$ is marked and such that

$$|L_1| \geq \frac{|H_j| - kt}{k + 1} \geq \frac{|H_j| - \varepsilon t |H_j|}{2k} \geq \frac{1 - \varepsilon t}{2\varepsilon} = r.$$

Now, the fact that no vertex in $V(L_1) - V(L_2)$ is marked implies that $V(L_1) - V(L_2) \subseteq V(H_{j-1})$, and hence no vertex in $V(L_1) - V(L_2)$ is adjacent in G to a vertex in $V(G) - V(L_1)$. Therefore, $G_1 := L_1$ and $G_2 := G - (V(L_1) - V(L_2))$ defines a pw- $(t - 1)$ -separation (G_1, G_2) of G with $|G_1| \geq r$. This concludes the proof of Case 1.

Case 2: Some vertex at distance at most d from w has degree more than Δ . Let x be such a vertex and let x_1, \dots, x_p denote its neighbours, where $p \geq \Delta + 1$.

First we prove a few easy claims.

Claim 4.7. *If there exists a connected subgraph H of G of pathwidth at least t such that $|H| \leq f(t, r) - d$ and $x \in V(H)$, then the lemma holds.*

Proof. Let P be a shortest w - x path in G and let $G' := G[V(H) \cup V(P)]$. Then G' is connected and has pathwidth at least t ; moreover, $w \in V(G')$ and $|G'| \leq |H| + d \leq f(t, r)$. □

Claim 4.8. *If there exists $X \subseteq V(G)$ with*

$$|X| \leq \frac{f(t, r) - d - r}{r + t} - 2t$$

such that $G[X]$ is connected, $x \in X$, and $H := G - X$ has a $\text{pw-}(t - 1)$ -separation (H_1, H_2) with $|H_1| \geq (r + t - 1)|X| + (r + t)2t + r$, then the lemma holds.

Proof. Using Lemma 4.2 with $k = 0$ and $\ell = (r + t - 1)|X| + (r + t)2t + r$, we may assume that the $\text{pw-}(t - 1)$ -separation (H_1, H_2) of H has been chosen so that $|H_1| = (r + t - 1)|X| + (r + t)2t + r$.

Let $J := G[V(H_1) \cup X]$. If J has pathwidth at least t , then since H_1 has pathwidth at most $t - 1$ there is a unique component J' of J of pathwidth at least t , namely the one containing X (recall that $G[X]$ is connected). Since $x \in V(J')$ and $|J'| \leq |J| = |H_1| + |X| = (r + t)|X| + (r + t)2t + r \leq f(t, r) - d$, we are done by Claim 4.7. Thus we may assume that J has pathwidth at most $t - 1$.

Let $Y := X \cup (V(H_1) \cap V(H_2))$ and mark all vertices of J that are in Y . Since $|J| = (r + t)|X| + (r + t)2t + r \geq (r + t)|Y| + r$, by Lemma 4.1 there is a $\text{pw-}(t - 1)$ -separation (J_1, J_2) of J with $|J_1| \geq r$ such that $Y \subseteq V(J_2)$. Observe that no vertex in $V(J_1) - Y$ is adjacent in G to a vertex in $V(G) - V(J_1)$. Hence, $G_1 := J_1$ and $G_2 := G - (V(J_1) - Y)$ defines a $\text{pw-}(t - 1)$ -separation (G_1, G_2) of G with $|G_1| \geq r$. □

Claim 4.9. *If there exists $X \subseteq V(G)$ with $|X| \leq 1 + \sum_{i=1}^3 f(t - 1, r_i)$ such that $G[X]$ is connected, $x \in X$, and every component of $G - X$ that includes some vertex in $\{x_1, \dots, x_p\}$ has pathwidth at most $t - 1$, then the lemma holds.*

Proof. Let q be the number of neighbours of x that are not in X ; let us assume without loss of generality that these neighbours are x_1, \dots, x_q . Let H be the union of all components of $G - X$ that include at least one of these vertices. Thus H has pathwidth at most $t - 1$ by the assumption of the claim. Since $|H| \geq q \geq p - |X| \geq \Delta + 1 - |X| \geq (r + t - 1)|X| + (r + t)2t + r$, by Lemma 4.1 there is a $\text{pw-}(t - 1)$ -separation (H_1, H_2) of H with $|H_1| \geq (r + t - 1)|X| + (r + t)2t + r$. The claim follows then from Claim 4.8, since

$$|X| \leq \sum_{i=1}^3 f(t - 1, r_i) + 1 \leq \frac{f(t, r) - d - r}{r + t} - 2t. \quad \square$$

Let $X_0 := \{x\}$. Apply the following argument first with $j = 1$, then $j = 2$, then $j = 3$. Since

$$|X_{j-1}| \leq 1 + \sum_{i=1}^{j-1} f(t-1, r_i),$$

the graph $G[X_{j-1}]$ is connected, and $x \in X_{j-1}$, by Claim 4.9 we may assume that there is a component C of $G - X_{j-1}$ with pathwidth at least t that includes a neighbour x_{n_j} of x (otherwise we are done). By induction (on t), either C has a connected subgraph H_j of pathwidth at least $t - 1$ with $|H_j| \leq f(t - 1, r_j)$ and $x_{n_j} \in V(H_j)$, or there is a pw- $(t - 2)$ -separation (C_1, C_2) of C with $|C_1| \geq r_j$. In the second case, (C_1, C'_2) is a pw- $(t - 2)$ -separation of $G - X_{j-1}$, where C'_2 is the union of C_2 with all components of $G - X_{j-1}$ distinct from C . Since

$$|X_{j-1}| \leq 1 + \sum_{i=1}^{j-1} f(t-1, r_i) \leq \frac{f(t, r) - d - r}{r + t} - 2t$$

and by the definition of r_j we have $|C_1| \geq r_j \geq (r + t - 1)|X_{j-1}| + (r + t)2t + r$, we are done by Claim 4.8. In the first case, let $X_j := X_{j-1} \cup V(H_j)$. Observe that $|X_j| \leq 1 + \sum_{i=1}^j f(t - 1, r_i)$ and that $G[X_j]$ is connected (because of the vertex x).

Now that the sets X_0, X_1, X_2, X_3 are defined, let P be a shortest path in G from w to X_3 . Thus $|P| \leq d + 1$, since $x \in X_3$ and x is at distance at most d from w in G . Since H_i is connected and has pathwidth at least $t - 1$ for each $i \in [1, 3]$, and since H_1, H_2, H_3 are pairwise vertex-disjoint, the graph $G[X_3]$ has pathwidth at least t by Lemma 4.5. It follows that $G' := G[X_3 \cup V(P)]$ is a connected subgraph of G that includes w , with pathwidth at least t , and satisfying

$$|G'| \leq |X_3| + |P| - 1 \leq 1 + \sum_{i=1}^3 f(t-1, r_i) + d \leq f(t, r).$$

This concludes the proof. □

Now we turn to the proof of Theorem 1.1.

Proof of Theorem 1.1. The proof is in two steps. First we prove the theorem in the special case where \mathcal{F} consists of a single tree, which we then use to handle the general case.

First consider the case where $\mathcal{F} = \{T\}$, where T is a tree on t vertices. Let $r := \max\{b(i) : 0 \leq i \leq 2t - 2\}$, where b is the function from Lemma 3.2. Let G be an arbitrary graph. Let $c := \max\{1, 2(t - 1), r, f(t - 1, r)\}$, where f is the function from Lemma 4.6. We will show that $\tau_{\mathcal{F}}(G) \leq c \cdot v_{\mathcal{F}}(G)$ by induction on $|G|$. If $|G| = 1$ then the claim obviously holds since $c \geq 1$. Now assume that $|G| > 1$. Since T is connected, we have that $v_{\mathcal{F}}(G) = v_{\mathcal{F}}(G_1) + \dots + v_{\mathcal{F}}(G_\ell)$ and $\tau_{\mathcal{F}}(G) = \tau_{\mathcal{F}}(G_1) + \dots + \tau_{\mathcal{F}}(G_\ell)$, where G_1, \dots, G_ℓ are the components of G . We may thus assume that G is connected, since otherwise we are done by applying induction on its components.

If the pathwidth of G is at most $t - 2$, then we directly obtain that $\tau_{\mathcal{F}}(G) \leq 2(t - 1) \cdot v_{\mathcal{F}}(G) \leq c \cdot v_{\mathcal{F}}(G)$ by Lemma 2.3. Hence assume that G has pathwidth at least $t - 1$.

Apply Lemma 4.6 on G (with w an arbitrary vertex of G) to obtain one of the two possible outcomes. If the outcome is a pw- $(t - 2)$ -separation (G_1, G_2) of G with $|G_1| \geq r$, then by Lemma 4.2 (with no marked vertices) we may assume that $|G_1| = r \leq c$. If G_1 contains T as a minor then $v_{\mathcal{F}}(G - V(G_1)) \leq v_{\mathcal{F}}(G) - 1$. By induction, there is an \mathcal{F} -transversal Y of $G - V(G_1)$ of size at most $c \cdot v_{\mathcal{F}}(G - V(G_1)) \leq c \cdot v_{\mathcal{F}}(G) - c$, and hence $V(G_1) \cup Y$ is an \mathcal{F} -transversal of G of size at most $c \cdot v_{\mathcal{F}}(G)$. If G_1 has no T -minor, then since $q := |V(G_1) \cap V(G_2)| \leq 2t - 2$ and $r \geq b(q)$, by Lemma 3.2 there is a graph G' with $v_{\mathcal{F}}(G') = \tau_{\mathcal{F}}(G)$, $\tau_{\mathcal{F}}(G') = \tau_{\mathcal{F}}(G)$, and $|G'| < |G|$. Since by induction $\tau_{\mathcal{F}}(G') \leq c \cdot v_{\mathcal{F}}(G')$, we are done in this case.

Suppose now that the outcome of Lemma 4.6 is a subgraph G' of G of pathwidth at least $t - 1$ with $|G'| \leq f(t - 1, r) \leq c$. Then, since G' contains T as a minor (by Theorem 2.1), $v_{\mathcal{F}}(G - V(G')) \leq v_{\mathcal{F}}(G) - 1$. By induction, there is an \mathcal{F} -transversal Y of $G - V(G')$ of size at most $c \cdot v_{\mathcal{F}}(G - V(G')) \leq c \cdot v_{\mathcal{F}}(G) - c$, and hence $V(G') \cup Y$ is an \mathcal{F} -transversal of G of size at most $c \cdot v_{\mathcal{F}}(G)$. This concludes the proof of the case where \mathcal{F} consists of a single tree.

Now let \mathcal{F} be an arbitrary finite set of graphs containing a forest F . Let $q := |\mathcal{F}|$ and let r denote the maximum number of components of a graph in \mathcal{F} . Let $t := |F|$. Let T be an arbitrary tree obtained from F by adding edges. By the proof above, there is a constant c' such that $\tau_T(G) \leq c' \cdot v_T(G)$ for every graph G . Using Lemma 2.4, we obtain

$$\tau_{\mathcal{F}}(G) \leq \tau_T(G) + 2qrt \cdot v_{\mathcal{F}}(G) \leq c' \cdot v_T(G) + 2qrt \cdot v_{\mathcal{F}}(G) \leq (c' + 2qrt) \cdot v_{\mathcal{F}}(G)$$

for every graph G . Therefore, the theorem holds with $c := c' + 2qrt$. □

We remark that, while Corollary 1.2 was deduced from Theorem 1.1 and the computability of the obstruction set for graphs of pathwidth at most t (see [1, 22]), it can alternatively be derived directly from Lemma 4.6.

5. Algorithmic implications

While the focus of this article is not algorithms, we would nevertheless like to point out a few algorithmic implications of our results.

First, the proof of Theorem 1.1 can be turned into a polynomial-time algorithm that, for fixed \mathcal{F} , computes in polynomial time an \mathcal{F} -packing and an \mathcal{F} -transversal of a given input graph G differing in size by at most a factor c . This is explained in part by the fact that the graph (G', R) in Lemma 3.1 can be computed in polynomial time using standard monadic second-order logic techniques (as done in [18], for instance), and the same is true for G' in Lemma 3.2. Moreover, it is not difficult to extend Lemma 3.2 to obtain that, given an \mathcal{F} -packing in G' , one can compute an \mathcal{F} -packing in G of the same size in polynomial time, and similarly given an \mathcal{F} -transversal of G' , one can compute an \mathcal{F} -transversal in G of no larger size in polynomial time. (This is needed when applying a reduction operation in

the proof of Theorem 1.1, since after having obtained an \mathcal{F} -packing and an \mathcal{F} -transversal of a reduced graph G' , we have to ‘lift back’ these to the input graph G .)

Furthermore, the outcomes of Lemmas 4.1 and 4.2 can easily be computed in polynomial time (for fixed t). The same is true for Lemma 4.6 (for fixed t and r), as its main computational steps are (i) breadth-first searches, (ii) calls to Lemmas 4.1 and 4.2, and (iii) tests of whether a graph has pathwidth at most i for some $i \leq t$ (which can be done in linear time when t is fixed: see [3]), and each is performed at most linearly many times (in fact, only a constant number of times for (i) and (ii)).

Since the proof of Theorem 1.1 makes at most linearly many calls to Lemmas 3.2, 4.6, and 2.3, it remains to show that Lemma 2.3 can be realized in polynomial time. While the proof itself does not directly yield such an algorithm, we note that the problem of finding a maximum \mathcal{F} -packing and that of finding a minimum \mathcal{F} -transversal can both be defined in monadic second-order logic. Hence, by Courcelle’s theorem [4], both problems can be solved in linear time on graphs of bounded pathwidth, and therefore an \mathcal{F} -packing and an \mathcal{F} -transversal such as promised by Lemma 2.3 can be found in linear time.

Second, a closer inspection shows that the running time of the algorithm sketched above is not only polynomial for fixed \mathcal{F} , but is moreover of the form $O(g(\mathcal{F}) \cdot n^\alpha)$ for n -vertex graphs, where g is a function depending only on \mathcal{F} , and α is a constant independent of \mathcal{F} . Building on this observation we now sketch a modification of the proof of Theorem 1.1 to obtain a *single-exponential* fixed-parameter tractable (FPT) algorithm for the problem of computing a minimum-size \mathcal{P}_t -transversal of a graph when parametrized by the size of the optimum, where \mathcal{P}_t is the finite set of minimal excluded minors for the class of graphs with pathwidth strictly less than t . Here single-exponential FPT means that the running time of the algorithm is $O(d^k \cdot n^\alpha)$ on instances G with $\tau_{\mathcal{P}_t}(G) = k$, where d is a constant depending only on t , and as before α is an absolute constant. Finding such an algorithm was posed as an open problem by Philip *et al.* [25], who gave one for the $t = 2$ case.

Given a pair (G, k) , our (recursive) FPT algorithm either finds a \mathcal{P}_t -transversal (abbreviated transversal) of G of size at most k , or correctly answers that there is no such transversal. The algorithm can be briefly described as follows (leaving some details to the reader).

- (I) If G is not connected, let G_1, \dots, G_p denote its components, and for each $i \in [1, p]$ and $\ell \in [0, k]$ recurse on (G_i, ℓ) . Given the results of these recursive calls, decide whether G has a transversal of size at most k .³
- (II) If G is connected but has pathwidth strictly less than t , return an empty transversal.
- (III) If G is connected with pathwidth at least t , apply Lemma 4.6 on G (with r defined as in the proof of Theorem 1.1).
 - (III.a) If the outcome of Lemma 4.6 is a pw- $(t - 1)$ -separation (G_1, G_2) of G , apply Lemma 3.2 to obtain a smaller graph G' and recurse on (G', k) ; if G' has no transversal of size at most k then so does G , otherwise ‘lift’ the transversal of G' found back to G and return it.

³ Here we use the fact that all graphs in \mathcal{P}_t are connected, which implies that X is a transversal of G if and only if $X \cap V(G_i)$ is a transversal of G_i for each $i \in [1, p]$.

(III.b) If the outcome of Lemma 4.6 is a connected subgraph G' of G with pathwidth at least t and at most $f(t,r)$ vertices, *branch* on every non-empty subset of $V(G')$ of size at most k , namely, for every such subset Y , recurse on $(G - Y, k - |Y|)$.

Observe that every transversal of G contains at least one vertex of G' . If a transversal of $G - Y$ of size at most $k - |Y|$ is found for some subset Y , return the union of that transversal with Y . If no transversal is found for any of the subsets Y , then G has no transversal of size at most k .

It should be clear that the algorithm finds a transversal of size at most k if there is one. The running time is $O(d^k \cdot n^\alpha)$ with $d := 2^{f(t,r)}$ and α some absolute constant, since we perform at most d recursive calls in a branching step, and in each one the parameter is decreased by at least one.

We conclude this section by mentioning that the optimization problem of finding a minimum \mathcal{F} -transversal in a graph has been considered for various families \mathcal{F} in several recent works: see [5, 6, 12, 14, 19, 21, 25].

Very recently, and independently of this work, Fomin *et al.* [13] obtained a single-exponential FPT algorithm for the minimum \mathcal{F} -transversal problem for every finite set \mathcal{F} of *connected* graphs containing at least one planar graph. Since all graphs in \mathcal{P}_t are connected, this includes in particular a single-exponential FPT algorithm for finding a minimum-size \mathcal{P}_t -transversal. The authors of [13] also presented a *randomized* (Monte Carlo) constant-factor approximation algorithm for finding a minimum-size \mathcal{F} -transversal when \mathcal{F} is a finite set of graphs containing at least one planar graph (here the graphs in \mathcal{F} are not assumed to be connected). We note that, while our constant-factor approximation algorithm described at the beginning of this section is restricted to the case where \mathcal{F} contains a forest, it is fully deterministic, and provides an \mathcal{F} -packing of size within a constant factor of optimal as well.

6. An open problem

For a finite set \mathcal{F} of graphs containing a forest, let

$$\rho(\mathcal{F}) := \sup \frac{\tau_{\mathcal{F}}(G)}{v_{\mathcal{F}}(G)}$$

over all graphs G with $v_{\mathcal{F}}(G) > 0$. By Theorem 1.1 this quantity is well defined. When \mathcal{F} consists of a single forest F , we simply write $\rho(F)$ for $\rho(\mathcal{F})$. For $t \geq 1$, let $\phi(t) := \max\{\rho(F) : F \text{ forest, } |F| = t\}$.

A natural next step would be to investigate the order of magnitude of $\phi(t)$. Our proof gives an upper bound which is exponential in t , and almost certainly far from the truth. In particular, it would be interesting to decide whether $\phi(t)$ is polynomial in t . As for lower bounds, we have that $\rho(F) \geq |F|$ for every forest F , as can be seen by taking G to be a large complete graph, and hence $\phi(t) \geq t$. But we do not know of any superlinear lower bound.

Acknowledgements

We would like to acknowledge the participation of Bundit Laekhanukit and Guylain Naves in the early stages of this research project. We thank them for several stimulating discussions. We also thank Dimitrios Thilikos for useful discussions regarding folios. Finally, we thank the anonymous referee for her/his thorough reading and very helpful comments.

References

- [1] Adler, I., Grohe, M. and Kreutzer, S. (2008) Computing excluded minors. In *Proc. Nineteenth Annual ACM–SIAM Symposium on Discrete Algorithms*, ACM, pp. 641–650.
- [2] Bienstock, D., Robertson, N., Seymour, P. and Thomas, R. (1991) Quickly excluding a forest. *J. Combin. Theory Ser. B* **52** 274–283.
- [3] Bodlaender, H. L. and Kloks, T. (1996) Efficient and constructive algorithms for the pathwidth and treewidth of graphs. *J. Algorithms* **21** 358–402.
- [4] Courcelle, B. (1990) Graph rewriting: An algebraic and logic approach. In *Handbook of Theoretical Computer Science, Vol. B*, Elsevier, pp. 193–242.
- [5] Cygan, M., Lokshtanov, D., Pilipczuk, M., Pilipczuk, M. and Saurabh, S. (2011) On the hardness of losing width. In *Proc. 6th International Symposium on Parameterized and Exact Computation, Saarbrücken, Germany* (D. Marx and P. Rossmanith, eds), Vol. 7112 of *Lecture Notes in Computer Science*, Springer, pp. 159–168.
- [6] Cygan, M., Pilipczuk, M., Pilipczuk, M. and Wojtaszczyk, J. (2010) An improved FPT algorithm and quadratic kernel for pathwidth one vertex deletion. In *Parameterized and Exact Computation* (V. Raman and S. Saurabh, eds), Vol. 6478 of *Lecture Notes in Computer Science*, Springer, pp. 95–106.
- [7] Diestel, R. (1995) Graph minors I: A short proof of the path-width theorem. *Combin. Probab. Comput.* **4** 27–30.
- [8] Diestel, R. (2010) *Graph Theory*, Vol. 173 of *Graduate Texts in Mathematics*, fourth edition, Springer.
- [9] Diestel, R., Kawarabayashi, K. and Wollan, P. (2012) The Erdős–Pósa property for clique minors in highly connected graphs. *J. Combin. Theory Ser. B* **102** 454–469.
- [10] Ellis, J. A., Sudborough, I. H. and Turner, J. S. (1994) The vertex separation and search number of a graph. *Inf. Comput.* **113** 50–79.
- [11] Erdős, P. and Pósa, L. (1965) On independent circuits contained in a graph. *Canad. J. Math.* **17** 347–352.
- [12] Fiorini, S., Joret, G. and Pietropaoli, U. (2010) Hitting diamonds and growing cacti. In *Integer Programming and Combinatorial Optimization* (F. Eisenbrand and F. Shepherd, eds), Vol. 6080 of *Lecture Notes in Computer Science*, Springer, pp. 191–204.
- [13] Fomin, F. V., Lokshtanov, D., Misra, N. and Saurabh, S. (2012) Foundations of Computer Science (FOCS). In *Proc. IEEE 53rd Annual Symposium on, Planar F-Deletion: Approximation, Kernelization and Optimal FPT Algorithms*, pp. 470–479.
- [14] Fomin, F. V., Lokshtanov, D., Misra, N., Philip, G. and Saurabh, S. (2011) Hitting forbidden minors: Approximation and kernelization. In *Proc. 28th International Symposium on Theoretical Aspects of Computer Science (STACS 2011)* (T. Schwentick and C. Dürr, eds), Vol. 9, Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, pp. 189–200.
- [15] Fomin, F. V., Saurabh, S. and Thilikos, D. M. (2011) Strengthening Erdős–Pósa property for minor-closed graph classes. *J. Graph Theory* **66** 235–240.
- [16] Grohe, M. (2008) Logic, graphs, and algorithms. In *Logic and Automata: History and Perspectives* (J. Flum, E. Grädel and T. Wilke, eds), Amsterdam University Press, pp. 357–422.

- [17] Grohe, M. and Marx, D. (2009) On tree width, bramble size, and expansion. *J. Combin. Theory Ser. B* **99** 218–228.
- [18] Grohe, M., Kawarabayashi, K., Marx, D. and Wollan, P. (2011) Finding topological subgraphs is fixed-parameter tractable. In *Proc. 43rd Annual ACM Symposium on Theory of Computing: STOC '11*, ACM, pp. 479–488.
- [19] Joret, G., Paul, C., Sau, L., Saurabh, S. and Thomassé, S. (2011) Hitting and harvesting pumpkins. In *Algorithms: ESA 2011* (C. Demetrescu and M. Halldórsson, eds), Vol. 6942 of *Lecture Notes in Computer Science*, Springer, pp. 394–407.
- [20] Kawarabayashi, K. and Kobayashi, Y. (2012) Linear min-max relation between the treewidth of H -minor-free graphs and its largest grid. In *Proc. 29th International Symposium on Theoretical Aspects of Computer Science: STACS 2012* (C. Dürr and T. Wilke, eds), Vol. 14 of *Leibniz International Proceedings in Informatics*, Schloss Dagstuhl–Leibniz-Zentrum für Informatik, pp. 278–289.
- [21] Kim, E. J., Paul, C. and Philip, G. A single-exponential FPT algorithm for the K_4 -minor cover problem. In *Proc. 13th Scandinavian Symposium and Workshops on Algorithm Theory: SWAT 2012*, to appear. [arXiv:1204.1417](https://arxiv.org/abs/1204.1417)
- [22] Lagergren, J. (1998) Upper bounds on the size of obstructions and intertwiners. *J. Combin. Theory Ser. B* **73** 7–40.
- [23] Leaf, A. and Seymour, P. Treewidth and planar minors. Preprint. <http://web.math.princeton.edu/~pds/papers/treewidth/paper.pdf>
- [24] Lubotzky, A., Phillips, R. and Sarnak, P. (1988) Ramanujan graphs. *Combinatorica* **8** 261–277.
- [25] Philip, G., Raman, V. and Villanger, Y. (2010) A quartic kernel for pathwidth-one vertex deletion. In *Graph Theoretic Concepts in Computer Science* (D. Thilikos, ed.), Vol. 6410 of *Lecture Notes in Computer Science*, Springer, pp. 196–207.
- [26] Robertson, N. and Seymour, P. D. (1986) Graph minors V: Excluding a planar graph. *J. Combin. Theory Ser. B* **41** 92–114.
- [27] Robertson, N. and Seymour, P. D. (2004) Graph minors XX: Wagner’s conjecture. *J. Combin. Theory Ser. B* **92** 325–357.
- [28] Robertson, N., Seymour, P. and Thomas, R. (1994) Quickly excluding a planar graph. *J. Combin. Theory Ser. B* **62** 323–348.
- [29] Seese, D. (1991) The structure of the models of decidable monadic theories of graphs. *Ann. Pure Appl. Logic* **53** 169–195.
- [30] Thatcher, J. W. and Wright, J. B. (1968) Generalized finite automata theory with an application to a decision problem of second-order logic. *Math. Systems Theory* **2** 57–81.