



Delineating Half-Integrality of the Erdős-Pósa Property for Minors: The Case of Surfaces

Christophe Paul 

LIRMM, Univ Montpellier, CNRS, Montpellier, France

Evangelos Protopapas 

LIRMM, Univ Montpellier, CNRS, Montpellier, France

Dimitrios M. Thilikos 

LIRMM, Univ Montpellier, CNRS, Montpellier, France

Sebastian Wiederrecht 

Discrete Mathematics Group, Institute for Basic Science, Daejeon, South Korea

Abstract

In 1986 Robertson and Seymour proved a generalization of the seminal result of Erdős and Pósa on the duality of packing and covering cycles: A graph has the *Erdős-Pósa property* for minors if and only if it is planar. In particular, for every non-planar graph H they gave examples showing that the Erdős-Pósa property does not hold for H . Recently, Liu confirmed a conjecture of Thomas and showed that every graph has the *half-integral* Erdős-Pósa property for minors. Liu's proof is non-constructive and to this date, with the exception of a small number of examples, no constructive proof is known.

In this paper, we initiate the delineation of the half-integrality of the Erdős-Pósa property for minors. We conjecture that for every graph H , there exists a *unique* (up to a suitable equivalence relation on graph parameters) graph parameter EP_H such that H has the Erdős-Pósa property in a minor-closed graph class \mathcal{G} if and only if $\sup\{\text{EP}_H(G) \mid G \in \mathcal{G}\}$ is finite. We prove this conjecture for the class \mathcal{H} of Kuratowski-connected shallow-vortex minors by showing that, for every non-planar $H \in \mathcal{H}$, the parameter $\text{EP}_H(G)$ is *precisely* the maximum order of a Robertson-Seymour counterexample to the Erdős-Pósa property of H which can be found as a minor in G . Our results are constructive and imply, for the first time, parameterized algorithms that find either a packing, or a cover, or one of the Robertson-Seymour counterexamples, certifying the existence of a half-integral packing for the graphs in \mathcal{H} .

2012 ACM Subject Classification Theory of computation \rightarrow Design and analysis of algorithms; Mathematics of computing \rightarrow Graph theory

Keywords and phrases Erdős-Pósa property, Erdős-Pósa pair, Graph parameters, Graph minors, Universal obstruction, Surface containment

Digital Object Identifier 10.4230/LIPIcs.ICALP.2024.114

Category Track A: Algorithms, Complexity and Games

Funding *Christophe Paul*: Supported by the French-German Collaboration ANR/DFG Project UTMA (ANR-20-CE92-0027).

Evangelos Protopapas: Supported by the French-German Collaboration ANR/DFG Project UTMA (ANR-20-CE92-0027).

Dimitrios M. Thilikos: Supported by the French-German Collaboration ANR/DFG Project UTMA (ANR-20-CE92-0027) and the MEAE and the MESR via the Franco-Norwegian project PHC Aurora project n. 51260WL (2024).

Sebastian Wiederrecht: Supported by the Institute for Basic Science (IBS-R029-C1).



© Christophe Paul, Evangelos Protopapas, Dimitrios M. Thilikos, and Sebastian Wiederrecht;
licensed under Creative Commons License CC-BY 4.0

51st International Colloquium on Automata, Languages, and Programming (ICALP 2024).

Editors: Karl Bringmann, Martin Grohe, Gabriele Puppis, and Ola Svensson;

Article No. 114; pp. 114:1–114:19



Leibniz International Proceedings in Informatics

LIPIcs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



1 Introduction

In 1965 Erdős and Pósa published a paper [9] proving the following min-max duality theorem.

For every positive integer k and every graph G , either G contains k pairwise vertex-disjoint cycles, or there exists a set $S \subseteq V(G)$ with $|S| = \mathcal{O}(k \cdot \log(k))$ such that $G - S$ has no cycles.

This result has since become central in both graph theory and algorithm design [36, 4, 21, 45, 26]. A collection of pairwise vertex-disjoint cycles is called a (cycle) *packing*, while a set S as above is commonly referred to as a (cycle) *cover* or *transversal*. In a more general context, one may consider any family \mathcal{M} of graphs and define $\text{pack}_{\mathcal{M}}(G)$ to be the largest size of a *packing* of members of \mathcal{M} in G , while $\text{cover}_{\mathcal{M}}(G)$ is the minimum size of a set $S \subseteq V(G)$ such that $G - S$ contains¹ no member of \mathcal{M} . Clearly $\text{pack}_{\mathcal{M}}(G) \leq \text{cover}_{\mathcal{M}}(G)$. We say that \mathcal{M} has the *Erdős-Pósa property* (EP-property) in a graph class \mathcal{G} if there exists a function f such that, for every $G \in \mathcal{G}$, it holds that $\text{cover}_{\mathcal{M}}(G) \leq f(\text{pack}_{\mathcal{M}}(G))$.

If we now fix some graph H and select \mathcal{M}_H to be the class of all graphs containing H as a minor, we enter the realm of the *Graph Minors Series* of Robertson and Seymour. In Graph Minors V. [36], as an implication of their min-max duality between the treewidth of a graph and its largest grid-minor, they prove that

For every graph H , \mathcal{M}_H has the EP-property in the class of all graphs if and only if H is planar. (1)

The tools and ideas of Erdős-Pósa-type dualities have since found many applications and interpretations [34, 27, 18, 3, 12, 33]. Moreover, the study of Erdős-Pósa dualities has led to important advances in structural graph theory. As an example, the proof for the directed version of Erdős and Pósa's result [35], known as *Younger's Conjecture* has paved the way for proving the *Directed Grid Theorem* [24].

Half-integral Erdős-Pósa. We call a collection \mathcal{C} of subgraphs of G a *half-integral packing* of \mathcal{M} in G if every graph in \mathcal{C} belongs to \mathcal{M} and no vertex of G belongs to more than two of them. We define $\frac{1}{2}\text{-pack}_{\mathcal{M}}(G)$ to be the maximum size of such a half-integral packing. Accordingly, \mathcal{M} has the $\frac{1}{2}$ EP-property in a graph class \mathcal{G} if there exists a function f such that, for every $G \in \mathcal{G}$, it holds that $\text{cover}_{\mathcal{M}}(G) \leq f(\frac{1}{2}\text{-pack}_{\mathcal{M}}(G))$.

Attempting to generalize Robertson and Seymour's seminal result on planar graphs, Robin Thomas conjectured the following relaxation of the EP-property (see [22, 26]).

For every graph H , \mathcal{M}_H has the $\frac{1}{2}$ EP-property in the class of all graphs. (2)

The above conjecture was recently proven by Liu [26]. As before, it is apparent from the definition that $\frac{1}{2}\text{-pack}_{\mathcal{M}}(G) \leq 2 \cdot \text{cover}_{\mathcal{M}}(G)$. Hence, Liu's theorem reveals a min-max duality between half-integral packing and covering in all graphs. Moreover, it is a consequence of the Graph Minors Theorem [37] that for every graph H and every graph parameter $\mathfrak{p} \in \{\text{pack}_{\mathcal{M}_H}, \frac{1}{2}\text{-pack}_{\mathcal{M}_H}, \text{cover}_{\mathcal{M}_H}\}$ one can decide in time $f_{H,\mathfrak{p}}(k)|V(G)|^3$ if $\mathfrak{p}(G) \geq k$ (or $\mathfrak{p}(G) \leq k$ in the case where $\mathfrak{p} = \text{cover}_{\mathcal{M}_H}$) [10] for some function $f_{H,\mathfrak{p}}$.

In light of the above results, it appears that the story of the Erdős-Pósa property in the regime of graph minors, from both a structural and an algorithmic perspective, is quite complete. However, we should stress the following two points.

First: The algorithm from [10] is inherently *non-constructive*. Indeed, while for $\text{pack}_{\mathcal{M}_H}$ and $\text{cover}_{\mathcal{M}_H}$ constructive algorithms are known [40, 39, 23], with the exception of some small special cases [22], *no* such results exist for $\frac{1}{2}\text{-pack}_{\mathcal{M}_H}$, not even approximation algorithms.

¹ At this point we consider containment to be defined through the subgraph relation.

Second: Let \mathcal{C} be a graph class and let p be a graph parameter. We say that p is *bounded* in \mathcal{G} if there exists $c \in \mathbb{N}$ such that, for every $G \in \mathcal{G}$, it holds that $p(G) \leq c$. The proof of the “if” direction of (1) was based on the fact that, for every H , \mathcal{M}_H has the EP-property in every graph class of bounded treewidth. This leads to the following *intermediate* question: *For which graph parameters p it holds that \mathcal{M}_H has the EP-property in every class where p is bounded?* To be specific, if we fix some graph H , is it possible to find a graph parameter EP_H such that \mathcal{M}_H has the EP-property in some minor-closed² graph class \mathcal{G} if and only if EP_H is bounded in \mathcal{G} ? Indeed, we conjecture that for every graph H , such a graph parameter exists and *precisely* delineates the half-integrality of the Erdős-Pósa property of \mathcal{M}_H .

► **Conjecture 1.** *For every graph H , there exists a minor-monotone graph parameter EP_H such that \mathcal{M}_H has the Erdős-Pósa property in a minor-closed graph class \mathcal{G} if and only if EP_H is bounded in \mathcal{G} .*

Notice that for any planar graph P , we can simply set EP_P to be the constant zero-function and thus, Conjecture 1 trivially holds for all planar graphs, because of (1). However, for non-planar graphs, the existence of such a parameter does not follow from any known results. Even if EP_H would exist for some particular non-planar graph H , it would be desirable to have some constructive, and ideally canonical, characterization of EP_H . That is, we aim at a description of EP_H that allows for algorithmic applications.

There are reasons to believe that EP_H exists and moreover has some canonical representation. It has recently been shown in [31], that this assertion is tied to the conjecture that graphs are ω^2 -well quasi ordered by minors, which is a wide open question in order theory (see the classic result of Thomas in [44] for the most advanced result on this conjecture).

The contribution of this paper is *resolving* Conjecture 1 for an infinite family of non-planar graphs. Moreover, our results are constructive and provide a canonical representation of EP_H yielding parameterized approximation algorithms³ for $1/2\text{-pack}_H$ for any H in our family.

1.1 The threshold of half-integrality

In Graph Minors V [36], towards proving the “only if” direction of (1), Robertson and Seymour gave counterexamples of graphs where non-planar graphs cannot have the EP-property. Let us investigate such an example for the graph K_5 . One may embed K_5 in both the projective plane and the torus, but it is impossible to have two disjoint drawings of K_5 in either of them.

Consider the two graphs in the middle of Figure 1 and notice that the number of cycles and paths can be scaled. We call the infinite sequences defined by such “scalable graphs” *parametric graphs*⁴. These parametric graphs are the *handle grid* \mathcal{H} and the *cross-cap grid* \mathcal{C} and represent the torus and the projective plane respectively. None of them contains two disjoint copies of graphs from \mathcal{M}_{K_5} , both have a half-integral packing of $\Omega(k)$ members of \mathcal{M}_{K_5} , and any minimum-size cover of all \mathcal{M}_{K_5} has $\Omega(k)$ vertices.


The seminal theorem of Reed [34] on the $1/2\text{EP}$ -property of odd cycles exhibits exactly this kind of behaviour. Reed showed that odd cycles have the EP-property in every *odd-minor*⁵-closed graph class excluding an *Escher-wall*, while the Escher-wall itself is a counterexample

² A graph class is *minor-closed* if it contains all minors of its graphs.

³ This means that our algorithms run in time $f(k) \cdot |V(G)|^{\mathcal{O}(1)}$ for some computable function f where k is the size of the half-integral packing we are looking for.

⁴ We postpone the formal definition of parametric graphs to a later point. See Section 2.

⁵ *Odd-minors* are a variant of the minor relation that preserves the parity of cycles. For example, bipartite graphs are exactly the K_3 -odd-minor-free graphs. We refer the interested reader to [14] for a formal definition.




■ **Figure 1** The parametric graphs representing the annulus grid \mathcal{A}_k , the handle grid \mathcal{H}_k , the cross-cap grid \mathcal{C}_k , and the shallow-vortex grid \mathcal{V}_k .

to the EP-property of odd cycles. Here, the k -Escher-wall is obtained by taking exactly the bipartite graphs from the parametric graph \mathcal{C} , representing the projective plane, and subdividing each of the “crossing” edges once. The result is a non-bipartite graph where every odd cycle must use an odd number of these subdivided edges. In the realm of odd-minors, this establishes a positive instance of Conjecture 1: pick EP_{K_3} as the maximum k for which G contains the k -Escher-wall as an odd minor.

It is tempting to suspect that Reed’s strategy can apply for the Erdős-Pósa property for minors. That is, for K_5 , the two parametric graphs \mathcal{H} and \mathcal{C} are essentially the only counterexamples for the EP-property of \mathcal{M}_{K_5} and excluding both of them as minors always yields a class in which \mathcal{M}_{K_5} exhibits the EP-property. Notice that this would imply that the \subseteq -minimal minor-closed classes where the EP-property fails for \mathcal{M}_{K_5} are precisely two: the class of graphs embeddable in the projective plane and the class of graphs embeddable in the torus. Clearly, both these two classes have bounded Euler genus. Our next step is to observe that this is not true in general.

Kuratowski-connectivity. We say that a graph G is *Kuratowski-connected* if for every separation (A, B) of G of order at most 3, if there is a component C of $G[A \setminus B]$ and a component D of $G[B \setminus A]$, such that every vertex in $A \cap B$ has a neighbour in $V(C)$ and a neighbour in $V(D)$, then one of $G[A]$, $G[B]$ can be drawn in a disc Δ with $A \cap B$ drawn in the boundary of Δ . We denote by \mathcal{K} the set of all Kuratowski-connected graphs. This definition was introduced by Robertson, Seymour, and Thomas as a tool for their characterization of *linklessly embeddable graphs* via a finite set of minimal obstructions [38] (see also [46, 30]).



■ **Figure 2** The two first parametric graphs serve as counterexamples for the Erdős-Pósa property of the graph J . The third parametric graph is a counterexample for the Erdős-Pósa property of K_8 . All three parametric graphs have unbounded Euler-genus. For the first two this is witnessed by a large packing of $K_{3,3}$ while the last one can be observed to contain $K_{3,r}$ as a minor.

Consider the graph J obtained by identifying two adjacent vertices of $K_{3,3}$ with two vertices of K_5 and observe that J is not Kuratowski-connected. Similar to K_5 , there cannot be two disjoint drawings of $K_{3,3}$ on the torus. So, if we take the parametric graph representing the torus (\mathcal{H}_k) or the projective plane (\mathcal{C}_k) from Figure 1 and paste “many” copies of $K_{3,3}$ around the “outer cycle”, we obtain a parametric graph without two disjoint J -minors but where no small vertex-set can hit all J -minors (see the two first graphs in Figure 2).

Shallow-vortex minors. There is a second property, that poses a similar issue. In [41] Thilikos and Wiederrecht introduced the parametric graph \mathcal{V} of *shallow-vortex grids* where \mathcal{V}_k is obtained from the annulus grid \mathcal{A}_k by adding k consecutive crossings in its internal cycle (see the fourth graph in Figure 1 for an illustration of \mathcal{V}_9). The class \mathcal{V} of *shallow-vortex minors* was defined in [41] as the class containing all minors of \mathcal{V}_k , for all $k \in \mathbb{N}$. Notice that K_8 is a Kuratowski-connected graph. It was shown by Curticapean and Xia [6] that K_8 is not a shallow-vortex minor. However, this is the case for $K_{3,r}$, for every $r \in \mathbb{N}$, which implies that the parametric graph \mathcal{V}_k has unbounded Euler-genus. If we now paste the k extra crossings of \mathcal{V}_k to the “outer cycle” of \mathcal{C}_k , we obtain a parametric graph that is a counterexample for the EP-property of K_8 but which is of unbounded Euler-genus (see the last graph in Figure 2). These observations indicate that, if we want to understand the graphs for which the counterexamples of Robertson and Seymour precisely define the boundary to the $1/2$ EP-property, we have to consider the graphs in $\mathcal{K} \cap \mathcal{V}$.

Our contribution. The main combinatorial result (stated in Theorem 2 in its full generality) is that Conjecture 1 holds, for every graph H that is Kuratowski-connected and a shallow vortex minor. Moreover, for every such non-planar H , $\text{EP}_H(G)$ is equivalent to the exclusion of the parametric graphs representing some particular set of surfaces where H embeds. Therefore, for the non-planar graphs $H \in \mathcal{K} \cap \mathcal{V}$, the boundary between the Erdős-Pósa property and its half-integral relaxation is drawn *precisely* by a set of surfaces, depending on H . Notice, that the class $\mathcal{K} \cap \mathcal{V}$ encompasses, apart from planar graphs, several important graphs such as $K_5, K_{3,3}, K_{4,4}, K_6, K_7$, and the entire *Petersen family*. These last observations imply that our results extend, both algorithmically and combinatorially, to the half-integral packing of *links* and *knots*.

2 Notation and definitions

Let us introduce some notation in order to present our results in full generality. A *minor antichain* is a family \mathcal{A} of graphs such that no graph $G_1 \in \mathcal{A}$ is a minor of another graph $G_2 \in \mathcal{A} \setminus \{G_1\}$. Since we focus on the minor relation, we refer to minor antichains simply as *antichains*. Let us denote by \mathbb{K} the collection of all antichains \mathcal{A} where every member of \mathcal{A} is Kuratowski-connected. Moreover, let us denote by \mathbb{V} the collection of all antichains containing at *least* one shallow-vortex minor. Finally, let \mathbb{P} be the collection of all antichains containing at least one planar graph and set $\mathbb{H} := \mathbb{K} \cap \mathbb{V}$ and $\mathbb{H}^- := \mathbb{H} \setminus \mathbb{P}$.

The Erdős-Pósa property for antichains. Let H and G be graphs. A subgraph $H' \subseteq G$ is an H -*host* in G if H is a minor of H' . An H -*packing* in G is a collection of pairwise vertex-disjoint H -hosts in G . An H -*cover* is a set $S \subseteq V(G)$ such that $G - S$ is H -minor-free. A *half-integral H -packing* is a collection of H -hosts in G such that no vertex of G belongs to more than two of them.

Given an antichain \mathcal{Z} , we say that a subgraph $H' \subseteq G$ is a \mathcal{Z} -host in G if it is an H -host for some $H \in \mathcal{Z}$. A \mathcal{Z} -packing is an H -packing of some $H \in \mathcal{Z}$ and a \mathcal{Z} -cover is an H -cover for all $H \in \mathcal{Z}$, finally a *half-integral \mathcal{Z} -packing* is a half-integral H -packing for some $H \in \mathcal{Z}$. We define the two graph parameters $\text{cover}_{\mathcal{Z}}$ and $\text{pack}_{\mathcal{Z}}$ as follows.

$$\text{cover}_{\mathcal{Z}}(G) := \min\{k \mid G \text{ has an } \mathcal{Z}\text{-cover of size } k\} \text{ and}$$

$$\text{pack}_{\mathcal{Z}}(G) := \max\{k \mid G \text{ has an } \mathcal{Z}\text{-packing of size } k\}.$$

We say that \mathcal{Z} has the *Erdős-Pósa property* in a graph class \mathcal{G} if there exists some function $f: \mathbb{N} \rightarrow \mathbb{N}$ such that $\text{cover}_{\mathcal{Z}}(G) \leq f(\text{pack}_{\mathcal{Z}}(G))$, for all $G \in \mathcal{G}$.

Equivalence of graph parameters. We use \mathcal{G}_{all} for the class of all graphs. Given two graph parameters $p, q: \mathcal{G}_{\text{all}} \rightarrow \mathbb{N}$, we say that p and q are *equivalent*, and write $p \sim q$, if there exists a function $f: \mathbb{N} \rightarrow \mathbb{N}$ such that, for every graph G , $p(G) \leq f(q(G))$ and $q(G) \leq f(p(G))$. We refer to the function f as the *gap* of this equivalence.

Our result is the identification of a graph parameter EP such that \mathcal{Z} has the Erdős-Pósa property in a minor-closed graph class \mathcal{G} with single-exponential gap if and only if EP is bounded in \mathcal{G} , for every $\mathcal{Z} \in \mathbb{H}$.

Surfaces and embeddability. We consider a containment relation \preceq between surfaces where we write $\Sigma \preceq \Sigma'$ if the surface Σ' can be obtained by adding handles or cross-caps to the surface Σ . The *empty surface* will be denoted by Σ^{\emptyset} and the surface obtained by adding h handles and c cross-caps to the sphere $\Sigma^{(0,0)}$ is denoted by $\Sigma^{(h,c)}$. Its *Euler-genus* is defined to be $2h + c$. Notice that, by Dyck's Theorem [8], we may assume that $c \leq 2$ for all surfaces. Let \mathbb{S} be a set of surfaces. We say that \mathbb{S} is *closed*, if $\Sigma \in \mathbb{S}$ and $\Sigma' \preceq \Sigma$ imply that $\Sigma' \in \mathbb{S}$ and that it is *proper*, if it does not contain all surfaces. If \mathbb{S} is closed and proper we define the “surface obstruction set” $\text{sobs}(\mathbb{S})$ as the set of all \preceq -minimal surfaces which do not belong to \mathbb{S} . It is easy to observe that $\text{sobs}(\mathbb{S})$ always consists of one or two surfaces [43]. Notice that $\text{sobs}(\emptyset) = \{\Sigma^{\emptyset}\}$, $\text{sobs}(\{\Sigma^{\emptyset}\}) = \{\Sigma^{(0,0)}\}$, $\text{sobs}(\{\Sigma^{\emptyset}, \Sigma^{(0,0)}\}) = \{\Sigma^{(1,0)}, \Sigma^{(0,1)}\}$, and, for a more complicated example, $\text{sobs}(\{\Sigma^{\emptyset}, \Sigma^{(0,0)}, \Sigma^{(0,1)}, \Sigma^{(0,2)}\}) = \{\Sigma^{(1,0)}\}$.

We say that a graph G is *embeddable* in a surface Σ (or Σ -embeddable) if it has a drawing in Σ without crossings. The *Euler genus* of a graph G , denoted by $\text{eg}(G)$, is the smallest Euler genus of a surface where G is embeddable.

Parametric graphs and Dyck-grids. A *parametric graph* is a sequence $\mathcal{G} = \langle \mathcal{G}_i \rangle_{i \in \mathbb{N}}$ of graphs indexed by non-negative integers. We say that \mathcal{G} is *minor-monotone* if for every $i \in \mathbb{N}$ we have that \mathcal{G}_i is a minor of \mathcal{G}_{i+1} . All parametric graphs considered in this paper are minor-monotone. We write $\mathcal{G}^{(1)} \lesssim \mathcal{G}^{(2)}$ for two minor-monotone parametric graphs $\mathcal{G}^{(1)}$ and $\mathcal{G}^{(2)}$ if there exists a function $f: \mathbb{N} \rightarrow \mathbb{N}$ such that for every $i \in \mathbb{N}$ it holds that $\mathcal{G}_i^{(1)}$ is a minor of $\mathcal{G}_{f(i)}^{(2)}$. A *minor-monotone parametric family* is a finite collection of $\mathfrak{G} = \{\mathcal{G}^{(j)} \mid j \in [r]\}$ of minor-monotone parametric graphs such that for distinct $i, j \in [r]$ it holds that $\mathcal{G}^{(i)} \not\lesssim \mathcal{G}^{(j)}$ and $\mathcal{G}^{(j)} \not\lesssim \mathcal{G}^{(i)}$. We define the minor-monotone parameter

$$p_{\mathfrak{G}}(G) := \max\{k \mid \text{there exists } i \in [r] \text{ such that } G \text{ contains } \mathcal{G}_k^{(i)} \text{ as a minor}\}. \quad (3)$$

The three parametric graphs $\mathcal{A} = \langle \mathcal{A}_k \rangle_{k \in \mathbb{N}}$, $\mathcal{H} = \langle \mathcal{H}_k \rangle_{k \in \mathbb{N}}$, and $\mathcal{C} = \langle \mathcal{C}_k \rangle_{k \in \mathbb{N}}$ are defined as follows: The *annulus grid* \mathcal{A}_k is the $(4k, k)$ -cylindrical grid⁶ depicted in the far left of Figure 1. The *handle grid* \mathcal{H}_k (resp. *cross-cap grid* \mathcal{C}_k) is obtained by adding in \mathcal{A}_k edges as indicated in the middle left (resp. middle right) part of Figure 1. We refer to the added edges as *transactions* of the handle grid \mathcal{H}_k or of the cross-cap grid \mathcal{C}_k .

Let now $h \in \mathbb{N}$ and $c \in [0, 2]$. We define the parametric graph $\mathcal{D}^{(h,c)} = \langle \mathcal{D}_k^{(h,c)} \rangle_{k \in \mathbb{N}}$ by taking one copy of \mathcal{A}_k , h copies of \mathcal{H}_k , and $c \in [0, 2]$ copies of \mathcal{C}_k , then “cut” them along the dotted red line, as in Figure 1, and join them together in the cyclic order $\mathcal{A}_k, \mathcal{H}_k, \dots, \mathcal{H}_k, \mathcal{C}_k, \dots, \mathcal{C}_k$, as indicated in Figure 3.




Figure 3 The Dyck-grid $\mathcal{D}_8^{1,2}$. The simple and the exceptional cycles are drawn in orange.

We call the graph $\mathcal{D}_k^{(h,c)}$ the *Dyck-grid of order k with h handles and c cross-caps*. Given some surface $\Sigma = \Sigma^{(h,c)}$, we say that the graph D is the $(\Sigma; d)$ -Dyck-grid if $D = \mathcal{D}_d^{(h,c)}$ and we use \mathcal{D}^Σ to denote the parametric graph $\langle \mathcal{D}_i^\Sigma \rangle_{i \in \mathbb{N}}$, where \mathcal{D}_i^Σ is the $(\Sigma; i)$ -Dyck-grid.

Let us return to our antichain $\mathcal{Z} \in \mathbb{H}^-$. We denote by $\mathbb{S}_{\mathcal{Z}}$ the set of surfaces where none of the graphs in \mathcal{Z} can be embedded. Notice that $\mathbb{S}_{\mathcal{Z}}$ is closed and proper and, for every $\Sigma \in \text{sobs}(\mathbb{S}_{\mathcal{Z}})$, there exists some $H \in \mathcal{Z}$ such that H embeds in Σ .

3 Our results

We associate with \mathcal{Z} the parametric family $\mathcal{D}_{\mathcal{Z}} := \{\mathcal{D}^\Sigma \mid \Sigma \in \text{sobs}(\mathbb{S}_{\mathcal{Z}})\}$. Let $\text{EP}_{\mathcal{Z}} := \text{p}_{\mathcal{D}_{\mathcal{Z}}}$. Our combinatorial result determines precisely when a member in \mathbb{H}^- has the Erdős-Pósa in some minor-closed graph class.

► **Theorem 2.** *For every $\mathcal{Z} \in \mathbb{H}^-$, for every minor-closed graph class \mathcal{G} , \mathcal{Z} has the Erdős-Pósa property in \mathcal{G} if and only if $\text{EP}_{\mathcal{Z}}$ is bounded in \mathcal{G} .*

Let $h_{\mathcal{Z}} := \max\{|V(H)| \mid H \in \mathcal{Z}\}$ and $\gamma_{\mathcal{Z}} := \max\{\text{eg}(H) \mid H \in \mathcal{Z}\}$. The engine that drives the proof of Theorem 2 and which represents our first main algorithmic result is the following.

► **Theorem 3.** *There exists a function $f_3 : \mathbb{N}^4 \rightarrow \mathbb{N}$ such that, for every antichain $\mathcal{Z} \in \mathbb{H}^-$, there exists an algorithm such that, given $k, t \in \mathbb{N}$ and a graph G , outputs one of the following:*

- a \mathcal{D}_t^Σ -host in G , for some $\Sigma \in \text{sobs}(\mathbb{S}_{\mathcal{Z}})$, or
- an \mathcal{Z} -packing of size at least k in G , or
- an \mathcal{Z} -cover of size at most $f_3(\gamma_{\mathcal{Z}}, h_{\mathcal{Z}}, t, k)$ in G .

Moreover, the algorithm runs in time $2^{2^{\mathcal{O}_{\gamma_{\mathcal{Z}}}(\text{poly}(t)) + \mathcal{O}_{h_{\mathcal{Z}}}(\text{poly}(k))}} \cdot |V(G)|^3 \cdot (\log(|V(G)|))^2$ and $f_3(\gamma_{\mathcal{Z}}, h_{\mathcal{Z}}, t, k) = 2^{\mathcal{O}_{\gamma_{\mathcal{Z}}}(\text{poly}(t)) + \mathcal{O}_{h_{\mathcal{Z}}}(\text{poly}(k))}$.

By a recent result of Gavaille and Hilaire [13], it holds that there exists some constant c such that for every $\mathcal{Z} \in \mathbb{H}^-$ and $\Sigma \in \text{sobs}(\mathbb{S}_{\mathcal{Z}})$, there exists some $H \in \mathcal{Z}$ such that H is a minor of $\mathcal{D}_{c\gamma_{\mathcal{Z}}^4 h_{\mathcal{Z}}^2}^\Sigma$. Moreover, as observed in [43], every Dyck-grid of big enough order

⁶ An $(n \times m)$ -cylindrical grid is a Cartesian product of a cycle on n vertices and a path on m vertices.

contains a large half-integral packing of itself of smaller order. Combining these two results with Theorem 3, yields the following (constructive) parameterized approximation algorithm for $1/2\text{-pack}_{\mathcal{Z}}$.

► **Theorem 4.** *There exists a function $f_4 : \mathbb{N}^2 \rightarrow \mathbb{N}$ such that, for every antichain $\mathcal{Z} \in \mathbb{H}^-$, there exists an algorithm such that, given $k \in \mathbb{N}$ and a graph G , outputs one of the following:*

1. *a half-integral \mathcal{Z} -packing of size at least k in G , or*
2. *an \mathcal{Z} -cover of size at most $f_4(h_{\mathcal{Z}}, k)$ in G .*

Moreover, the algorithm runs in time⁷ $2^{2^{\text{poly}_{h_{\mathcal{Z}}}(k)}} \cdot |V(G)|^3 \cdot (\log(|V(G)|))^2$ and $f_4(h_{\mathcal{Z}}, k) = 2^{\text{poly}_{h_{\mathcal{Z}}}(k)}$.

We wish to stress that, given the combinatorial bounds of Theorem 3, we may directly apply the minor-checking algorithm of [23] for the two first outcomes of Theorem 3 and the algorithm of [29] for its third outcome. Both these algorithms are quadratic on $|V(G)|$ and this implies alternative quadratic algorithms to those in Theorem 3 and Theorem 4. However, this would come with the cost of enormous parametric dependencies on k .

3.1 Some implications of our results

Half-integral Erdős-Pósa for linked pairs and knots. As mentioned above, $\mathbb{H} = \mathbb{K} \cap \mathbb{V}$ contains several antichains of particular interest. A first example is the *Petersen family*, which is exactly the (minor) obstruction set⁸ for the so-called *linklessly embeddable* graphs (in short, *link-less* graphs). Indeed, the origin of the definition of Kuratowski-connectivity comes from the paper of Robertson, Seymour, and Thomas [38], where this obstruction set was found. All obstructions for link-less graphs as well as those for knot-less graphs are Kuratowski-connected. Moreover, as the shallow-vortex minor K_6 (resp. K_7) is a member of the obstruction set of link-less (resp. knot-less) graphs, we also have that both these obstruction sets belong to \mathbb{H}^- . This insight allows us to apply Theorem 4 to topological objects such as *links* and *knots*.

Let G be a graph and let $\mathcal{C} = \{C_1, \dots, C_k\}$ be a collection of subgraphs of G . The *intersection graph* of \mathcal{C} is the graph $I(\mathcal{C}) = (\mathcal{C}, E_{\mathcal{C}})$ where $CC' \in E_{\mathcal{C}}$ if and only if $C \cap C'$ is not the empty graph. We say that \mathcal{C} is a *collection of double cycles* (resp. *cycles*) if each C_i is union of two disjoint cycles (resp. a cycle).

Given a collection \mathcal{C} of double cycles (resp. cycles) of G , and some \mathbb{R}^3 -embedding of G , we say that \mathcal{C} is a $1/2$ -*packing* of links (resp. knots) if for every $i \in [k]$, the two components of C_i are linked (resp. the cycle C_i is knotted) in this particular embedding (see [1] for more on links and knots). The *half-integral linked pair* (resp. *knot*) *packing number* of a graph G , denoted by $1/2\text{-lppack}(G)$ (resp. $1/2\text{-knpack}(G)$), is the maximum k such that, for every \mathbb{R}^3 -embedding of G , there exists a $1/2$ -packing of links (resp. knots) in G of size k . Both $1/2\text{-lppack}(G)$ and $1/2\text{-knpack}(G)$ are minor-monotone parameters, therefore we know (non-constructively) that there exists an algorithm for checking whether $1/2\text{-lppack}(G) \geq k$ ($1/2\text{-knpack}(G) \geq k$) in time $f(k) \cdot |V(G)|^2$. Up to now, no constructive (on k) algorithm is known for these problems. Our results imply the following.

⁷ Given two functions $\chi, \psi : \mathbb{N} \rightarrow \mathbb{N}$, we write $\chi(n) = \mathcal{O}_x(\psi(n))$ in order to denote that there exists a computable function $f : \mathbb{N} \rightarrow \mathbb{N}$ such that $\chi(n) = \mathcal{O}(f(n) \cdot \psi(n))$. We also use $\chi(n) = \text{poly}_x(\psi(n))$ instead of $\chi(n) = \mathcal{O}_x((\psi(n))^c)$, for some $c \in \mathbb{N}$.

⁸ The obstruction set of some minor-closed class \mathcal{G} is the set $\text{obs}(\mathcal{G})$ of the minor-minimal graphs that are not in \mathcal{G} .

► **Theorem 5.** *There exists a function $f : \mathbb{N} \rightarrow \mathbb{N}$ and algorithms that, given a graph G and a $k \in \mathbb{N}$, outputs either that $1/2\text{-lppack}(G) \geq k$ (resp. $1/2\text{-knpack}(G) \geq k$) or a vertex set A of at most $f(k)$ vertices such that $G - A$ has a link-less (knot-less) \mathbb{R}^3 -embedding. Moreover, both algorithms run in time $2^{2^{\text{poly}(k)}} \cdot |V(G)|^3 \cdot (\log(|V(G)|))^2$ and $f(k) = 2^{\text{poly}(k)}$.*

Theorem 5 implies that both the parameter $1/2\text{-lppack}$ as well as the parameter $1/2\text{-knpack}$ admit FPT-approximation algorithms with exponential approximation gap. Moreover, in case the output is that $1/2\text{-lppack}(G) \geq k$ (resp. $1/2\text{-knpack}(G) \geq k$), the algorithms output a $1/2$ -packing of k graphs certifying that, every \mathbb{R}^3 -embedding of G contains a $1/2$ -packing \mathcal{C} of k links (resp. knots) such that $I(\mathcal{C})$ is either edgeless or a clique. We stress that the above algorithms become constructive (on k) as we know the obstructions of link-less/knot-less graphs or at least an upper bound to their size. For the later class not such bound is known.

Other implications of our results, related to canonical approximate characterizations of the parameters we study, are discussed in the conclusion section (Section 4).

3.2 Outline of the proof

We begin the description of the main ideas of our proof with the definition of a tree decomposition.

Tree decompositions. Let G be a graph. A *tree decomposition* of a graph G is a pair (T, β) where T is a tree and $\beta : V(T) \rightarrow 2^{V(G)}$ is a function, whose images are called the *bags* of \mathcal{T} , such that $\bigcup_{t \in V(T)} \beta(t) = V(G)$, for every $e = xy \in E(G)$, there exists $t \in V(T)$ with $\{x, y\} \subseteq \beta(t)$, and for every $v \in V(G)$, the set $\{t \in V(T) \mid v \in \beta(t)\}$ induces a subtree of T . We refer to the vertices of T as the *nodes* of the tree decomposition \mathcal{T} . The *width* of \mathcal{T} is the value $\max_{t \in V(T)} |\beta(t)| - 1$. The *treewidth* of G , denoted by $\text{tw}(G)$, is the minimum width over all tree decompositions of G .

The classic approach. In order to facilitate the presentation of our proof, let us briefly explain the two main ideas of the proof that planar graphs enjoy the Erdős-Pósa property in the set of all graphs. The key ingredient is that every planar graph is a minor of a graph of sufficiently large treewidth. The proof follows in two steps.

Step 1. Assuming that $\text{pack}_H(G) \leq k$, based on the grid theorem by Robertson and Seymour, we may assume that the treewidth of G is bounded by some function of k .

Step 2. With the tree decomposition (T, β) of G at hand, we build an H -cover A of G by adding to it (if any exists) an adhesion $D_{xy} = \beta(x) \cap \beta(y)$ such that both $G_x := G[\beta(V(T_x)) \setminus D_{xy}]$ and $G_y := G[\beta(V(T_y)) \setminus D_{xy}]$ contain H as a minor (here T_x and T_y are the two components of $T - xy$) and then recursing on the corresponding tree decompositions of G_x and G_y . If $\text{pack}_H(G) \leq k$, eventually this procedure returns an H -cover of size at most $k \cdot (\text{tw}(G) + 1)$.

Throughout the present outline we describe arguments that can be paralleled to the two steps above. Moreover, in each step we explain the challenges that are met and the way we deal with them in our proof.

For simplicity, instead of an antichain \mathcal{Z} , we consider a non-planar graph H that is Kuratowski-connected and a shallow-vortex minor. We denote by \mathbb{S}_H the set of all surfaces where H cannot be embedded and by $\mathbb{S}'_H := \text{sobs}(\mathbb{S}_H)$ the corresponding surface obstruction set. We stress that the graphs in $\mathfrak{D}_H = \{\mathcal{D}^\Sigma \mid \Sigma \in \mathbb{S}'_H\}$ can be seen as “generators of

half-integrality". Indeed, it is possible to prove that, for every $t \in \mathbb{N}$, $\text{pack}_H(\mathcal{D}_t^\Sigma) \leq 1$, and $\text{cover}_H(\mathcal{D}_t^\Sigma) = \Theta(1/2 \cdot \text{pack}_H(\mathcal{D}_t^\Sigma)) = \Omega(t)$. This already proves the easy direction of Theorem 2.

Let $\mathcal{T} = (T, \beta)$ be a tree decomposition of a graph G . For each $t \in V(T)$, we define the *adhesions* of t as the sets in $\{\beta(t) \cap \beta(d) \mid d \text{ adjacent with } t\}$ and the maximum size of them is called the *adhesion* of t . The *adhesion* of \mathcal{T} is the maximum adhesion of a node of \mathcal{T} . The *torso* of \mathcal{T} on a node t is the graph, denoted by G_t , obtained by adding edges between every pair of vertices of $\beta(t)$ which belongs to a common adhesion of t .

We now consider a graph G where $\text{pack}_H(G) \leq k$ and we assume that G excludes as a minor the Dyck grid \mathcal{D}_t^Σ , for every $\Sigma \in \mathbb{S}'_H$. Under these circumstances, our aim is to find an H -cover whose size is bounded by some function of t and k .

Graphs excluding Dyck grids. As a first step, we need a deeper understanding of how the graphs excluding \mathcal{D}_t^Σ look like. In general, the structure of graphs excluding a given graph as a minor is given by the Graph Minors Structure theorem (in short GMST). However, the formal definition of GMST involves complicated concepts which we prefer not to introduce in this brief outline. Instead we give a more compact statement, proved in [43].

Given a graph H and a set $A \subseteq V(G)$, we say that H is an A -minor of G if there is a collection $\mathcal{S} = \{S_v \mid v \in V(H)\}$ of pairwise vertex-disjoint connected⁹ subsets of $V(G)$, each containing at least one vertex of A and such that, for every edge $xy \in E(H)$, the set $S_x \cup S_y$ is connected in G . Given an annotated graph (G, A) where G is a graph and $A \subseteq V(G)$, we define $\text{tw}(G, A)$ as the maximum treewidth of an A -minor of G . A streamlined way to restate the GMST is the following.

► **Proposition 6** ([43]). *There exists a function $f : \mathbb{N} \rightarrow \mathbb{N}$ such that every graph G excluding a graph on k vertices as a minor, has a tree decomposition (T, β) where, for every $t \in V(T)$, the torso G_t contains some set A_t where $\text{tw}(G_t, A_t) \leq f(k)$ and such that $G_t - A_t$ can be embedded in a surface of Euler genus at most $f(k)$.*

To deal with the exclusion of Dyck grids (corresponding to surfaces), we need a more refined version of Proposition 6 that works for every (closed and proper) set of surfaces \mathbb{S} . In this direction, Thilikos and Wiederrecht defined in [43] an extension of treewidth, namely \mathbb{S} -tw, where for a graph G ,

$\mathbb{S}\text{-tw}(G)$ is the minimum k for which G has a tree decomposition (T, β) where, for every $t \in V(T)$, the torso G_t contains some set A_t where $\text{tw}(G_t, A_t) \leq k$ and $G_t - A_t$ is (4) embeddable in a surface in \mathbb{S} .

The main result of [43] is that in order to exclude the graphs in $\mathfrak{D}_H = \{\mathcal{D}_t^\Sigma \mid \Sigma \in \text{cobs}(\mathbb{S})\}$, we have to fix the surface of Proposition 6 to be one of the surfaces in \mathbb{S} .

► **Proposition 7.** *For every closed and proper set of surfaces \mathbb{S} , there exists some function $f : \mathbb{N} \rightarrow \mathbb{N}$ such that, for every graph G , if G excludes all graphs in $\{\mathcal{D}_t^\Sigma \mid \Sigma \in \text{sobs}(\mathbb{S})\}$ as minors, then $\mathbb{S}\text{-tw}(G) \leq f(t)$.*

Notice that the above proposition already gives us the grid theorem when applied for the set \mathbb{S}_\emptyset containing the empty surface Σ^\emptyset . It is easy to verify that $\text{tw} + 1 = \mathbb{S}_\emptyset\text{-tw}$. As $\text{sobs}(\mathbb{S}_\emptyset) = \{\Sigma^{(0,0)}\}$, Proposition 7 implies that graphs excluding $\mathcal{D}_t^{\Sigma^{(0,0)}} = \mathcal{A}_t$ have bounded treewidth (see Figure 1 for an example of an annulus grid).

⁹ A set $X \subseteq V(G)$ is *connected* in G if the induced subgraph $G[X]$ is a connected graph.


From small treewidth modulators to small size modulators. Proposition 7 gives valuable information on the structure of the graphs that exclude the “half-integrality generators” in $\mathfrak{D}_H = \{\mathcal{D}_t^\Sigma \mid \Sigma \in \text{sobs}(\mathbb{S}_H)\}$. Therefore, we can assume that $\mathbb{S}\text{-tw}(G) \leq f(k)$, which provides a tree decomposition as the one in (4). In order to make progress, we need to further refine this decomposition as small treewidth modulators are not particularly helpful in finding an H -cover of small size. For this we exploit the assumption that H is a shallow-vortex minor.

To elaborate, we need some additional information, analogous to the exclusion of a planar graph in **Step 1**. This corresponds to the assumption that H is a minor of the shallow-vortex grid $\mathcal{V}_{h'}$ for some h' depending on H . One can observe that $\mathcal{V}_{3(k+1)h'}$ contains a $\mathcal{V}_{h'}$ -packing of size $(k+1)$. Therefore, the assumption that $H\text{-pack}(G) \leq k$ gives us the right to additionally assume that G also excludes the shallow-vortex minor $\mathcal{V}_{3(k+1)h'}$. Using this and the fact that $\mathbb{S}\text{-tw}(G) \leq f(k)$, we are able to further restrict the decomposition of (4). To quantify this, we introduce a new graph parameter $\mathbb{S}\text{-tw}_{\text{apex}}$ defined as follows.

$\mathbb{S}\text{-tw}_{\text{apex}}(G)$ is the minimum k for which G has a tree decomposition (T, β) where, for every $t \in V(T)$, the torso G_t contains some set A_t where $|A_t| \leq k$ and $G_t - A_t$ is embeddable in a surface in \mathbb{S} . (5)

Notice that the only difference between (4) and (5) is the measure defined on the “modulator” A_t . While in (4) it is the treewidth of the annotated graph (G_t, A_t) , in (5) it is the size of A_t . The first ingredient of our proof is that, under the absence of some shallow-vortex minor, the two parameters $\mathbb{S}\text{-tw}_{\text{apex}}$ and $\mathbb{S}\text{-tw}$ are equivalent. This is proved by combining the results of [43] with the results of [41] on the structure of the graphs excluding a shallow-vortex grid.

As a consequence, we may now assume that we have a tree decomposition (T, β) as the one in (5). This decomposition is not yet in position to play the role of the tree decomposition in **Step 2**, as its torsos may have unbounded size. To circumvent this issue, we instead prove a local structure theorem for the exclusion of $\mathfrak{D}_H \cup \{\mathcal{V}\}$, that can be extended to a global one (the desired tree decomposition), using the results of [42]. The general approach is to consider some big enough wall W_t and locally focus on a torso G_t that contains most of the essential part of W_t .



■ **Figure 4** The elementary $(h, c; k)$ -Dyck wall, where $h = 1$, $c = 1$, and $k = 6$.

Torsos with Dyck walls. According to Proposition 7, (5), and the equivalence of $\mathbb{S}\text{-tw}_{\text{apex}}$ and $\mathbb{S}\text{-tw}$, G_t comes together with a set A_t such that the graph $G'_t := G_t - A_t$ is accompanied by some Σ -embedding for a surface $\Sigma \in \mathbb{S}_H$, where H cannot be embedded. However, we require some additional infrastructure in G_t that will come in the form of a large wall-like object that is controlled by our Σ -embedding.

Notice that every adhesion $\beta(t) \cap \beta(t')$ of t defines a separation $(X_{t'}, Y_{t'})$ of $G - A_t$ of order at most 3 where $G[X_{t'} \cap Y_{t'}]$ is drawn in Σ as a clique. We fix the orientation $(X_{t'}, Y_{t'})$ such that $V(G_t) \subseteq Y_{t'}$, thereby indicating that $Y_{t'}$ is the “important” part of the separation.

Due to the results in [43], G_t contains a $(\Sigma; d)$ -Dyck wall¹⁰ D_t , which is highly linked to the wall W_t above. Here d is chosen “large enough” so as to ensure the applicability of the next steps of our proof. Also, we may assume that the “essential” part of D_t is drawn “inside” G_t in the sense that, for each $(X_{t'}, Y_{t'})$, at most one branch vertex of D_t is in $Y_{t'} \setminus X_{t'}$. The wall W_t is chosen large enough to represent some *tangle*, that is an orientation of the separations of G of some suitably bounded order. The way to algorithmically detect such a big wall W_t is given in [43].

The role of Kuratowski-connectivity. We next make some observations on how “models” of H can behave with respect to the Σ -embedding of G_t . These observations will play a key role in understanding how to “attack” and later “kill” copies of H in our graph.

The first comes from the non- Σ -embeddability property of H : “minimal” H -hosts in G , called *H -inflated copies*, cannot be entirely inside G_t , otherwise we would be able to embed H in a surface where it cannot be embedded. Another important feature comes from the fact that H is Kuratowski-connected: every H -inflated copy M in G is “well oriented” with respect to the adhesions of t in the sense that, when M traverses some adhesion $X_{t'} \cap Y_{t'} = \beta(t) \cap \beta(t')$ of G , exactly one of the two parts of M induced by $X_{t'}$ and $Y_{t'}$ should not be embeddable in the disk bounding $\beta(t) \cap \beta(t')$ with the vertices of $\beta(t) \cap \beta(t')$ on its boundary. This implies that the “non-disk-embeddable” part will always lie inside the set $X_{t'}$ of the separation $(X_{t'}, Y_{t'})$ above. Given now some adhesion $\beta(t) \cap \beta(t')$, we say that it is *H -red* if it is intersected by the (unique, due to Kuratowski-connectivity) non-disk-embeddable part of some H -inflated copy M in G . That way, it is convenient to visualize H -red adhesions as the “entrances” from which the H -inflated copies of G “invade” G_t .

Updating the Σ -embedding. From our previous observations it follows that to eliminate all copies of H locally in G_t it suffices to deal with all H -inflated copies that invade G_t through H -red adhesions. Therefore, our next objective is to update A_t , $G'_t = G_t - A_t$, and the Σ -embedding of G'_t in a way that the remaining part of G'_t will not contain any H -red adhesions, i.e., in a way that no invading H -inflated copy survives.

During our proof, this updating procedure will focus on some closed disk Δ containing some collection of H -red adhesions (these disks will be gathered together in what we call *H -red railed flat vortices*) and detect some separation (X, Y) of G where $X \setminus Y$ contains the vertices of the Dyck wall D_t and Y contains all H -red adhesions in Δ . We call such a separation a *carving separation*. Each time we find such a separation, we move $X \cap Y$ to A and also move $Y \setminus X$ “outside” G_t . As the set $X \cap Y$ adds up to the size of A we also need that $X \cap Y$ has “small” order. We refer to this operation as *taking a carving* of our Σ -embedding at the carving separation (X, Y) . When the whole procedure terminates, none of the adhesions of the updated G'_t is H -red. This implies that $(V(G_t), A_t)$ is what we call an *H -dominion* of G , that is: *if the non-disk-embeddable part of some H -inflated copy in G intersects $V(G_t)$ then it also intersects A_t .*

To achieve the previously described objective we adopt the following strategy. Recall that in the Σ -embedding of G_t , H -red adhesions are cliques of size at most three that may be drawn all around Σ . Our first step is to show that H -red adhesions can be cornered in the “interior” of less than k pairwise-disjoint territories of Σ , each maintaining a large enough “buffer” around a disk where the H -red adhesions reside. Afterwards we refine these territories in order to bound their complexity in the sense that there is no large flow in G_t

¹⁰ Here a $(\Sigma; d)$ -Dyck wall is certifying the existence of the Dyck grid \mathcal{D}_d^Σ as a minor. See Figure 4.

that crosses through these territories. Through this refinement step we obtain some additional structural information so that in the last part of the proof, these territories along with their infrastructure will allow us to finally eliminate all H -inflated copies by removing a bounded number of vertices from their interiors.

Redrawing H -inflated copies inside a railed flat vortex. To formalize the aforementioned territories that will encapsulate the H -red adhesion of our embedding, we utilize the concept of a *railed nest* $(\mathcal{C}, \mathcal{P})$ of G around some closed disk Δ^{int} of Σ . Here $\mathcal{C} = \langle C_1, \dots, C_\ell \rangle$ is a sequence of ℓ disjoint cycles of G_t , where each C_i bounds some closed disk Δ_i in Σ , where $\Delta^{\text{int}} \subseteq \Delta_1 \subsetneq \dots \subsetneq \Delta_\ell$, along with a set of paths $\mathcal{P} = \langle P_1, \dots, P_\ell \rangle$, drawn in $\Delta^{\text{ext}} := \Delta_\ell$, not traversing the interior of Δ^{int} , joining vertices of C_1 with vertices of C_ℓ , and traversing the cycles in \mathcal{C} *orthogonally*, that is $P_i \cap C_j$ is connected for every $(i, j) \in [\ell]^2$. We refer to such a railed nest, as a *railed flat vortex* and we refer to the disk Δ^{int} (resp. Δ^{ext}) as its *internal* (resp. *external disk*). Moreover, if all H -red adhesions drawn in Δ^{ext} are also drawn inside Δ^{int} , then we call it an *H -red railed flat vortex*. An important ingredient of our proof is to show that we may use the infrastructure of the cycles and the paths in $(\mathcal{C}, \mathcal{P})$ in order to *redraw* inside Δ^{ext} every H -inflated copy M that invades G_t via an H -red adhesion of Δ^{ext} . Even if the part of M that is embedded inside Δ^{ext} is not necessarily a disk embedding, we can make this redrawing possible by using disk embedability properties emerging from the Kuratowski-connectivity of H and the “linkage combing” lemma from [17, 16, 15]. We refer to this as *the redrawing lemma*.

Gathering H -adhesions in railed flat vortices. The next step of our strategy, is to corner all H -red adhesions in the interior of less than k H -red railed flat vortices. Towards this, we take advantage of the infrastructure provided by the $(\Sigma; d)$ -Dyck wall D_t . A brick of D_t is called H -red if it “contains” an H -red adhesion. More precisely, this is formalized by the notion of the *influence* of a brick which roughly corresponds to a set of H -red adhesions that are intersected or contained by a closed disk in Σ that bounds the “area” that is enclosed by the corresponding brick. This assigns each H -red adhesion to the influence of at most three neighbouring H -red bricks and defines a notion of distance between H -red adhesions expressed by the distance of the corresponding H -red bricks in D_t . Next, we prove that under this distance notion, no scattered enough set of H -red bricks of size k can exist. For this, we use the fact that each H -red brick B implies the existence of an H -inflated copy in G that, due to the aforementioned “redrawing lemma”, can be redrawn in a small radius around B . This radius is bounded but also big enough so as to permit the redrawing. Likewise, we prove that there are few H -red bricks away from the exceptional and the simple cycle of D_t (see Figure 4 for a visualization of these two cycles). Next we use a greedy procedure in order to group together this bounded number of bricks and maintain enough railed nest infrastructure around them to cluster them into less than k railed flat vortices. The construction is completed by creating two more railed flat vortices, one for the simple cycle of D_t and one for the exceptional one.

Refining H -red railed flat vortices. We are now in the position where we have defined a set of less than k many H -red railed flat vortices whose internal disks contain all H -red adhesions and whose external disks are pairwise disjoint. The next step is to further refine these flat vortices.

In our proof, we treat what is drawn in the external disk Δ^{ext} as a vortex in the classic sense and our goal is to bound their *depth*, that is to ensure that no large *transaction* goes through the society defined by each railed flat vortex. Each of them consists of a

subgraph $G_{\Delta^{\text{ext}}}$ of G (the one that is drawn in Δ^{ext}) where the vertices in the boundary of the external disk Δ^{ext} are arranged in some cyclic ordering $\Omega_{\Delta^{\text{ext}}}$. A *segment* of $\Omega_{\Delta^{\text{ext}}}$ is a set $S \subseteq V(\Omega_{\Delta^{\text{ext}}})$ such that there do not exist $s_1, s_2 \in S$ and $t_1, t_2 \in V(\Omega_{\Delta^{\text{ext}}}) \setminus S$ such that s_1, t_1, s_2, t_2 occur in Ω_{Δ} in the order listed. A *transaction* in $(G_{\Delta^{\text{ext}}}, \Omega_{\Delta^{\text{ext}}})$ is a set of pairwise disjoint paths, drawn in Δ , between two disjoint segments A, B of $\Omega_{\Delta^{\text{ext}}}$. The *depth* of $(G_{\Delta^{\text{ext}}}, \Omega_{\Delta^{\text{ext}}})$ is the maximum size of a transaction in $(G_{\Delta^{\text{ext}}}, \Omega_{\Delta^{\text{ext}}})$. Our next objective is to refine each of our H -red railed flat vortices so that, in the end, some disk $\Delta' \subseteq \Delta^{\text{ext}}$ defines a vortex $(G_{\Delta'}, \Omega_{\Delta'})$ of *bounded* depth and, moreover, the vertices in the boundary of Δ' are all connected with disjoint paths to the boundary of the external disk Δ^{ext} . We do this as follows: If there is no transaction in $(G_{\Delta^{\text{ext}}}, \Omega_{\Delta^{\text{ext}}})$ where a big part of its paths also traverse Δ^{int} , we make use of the “nest tightening”-lemma from [41] in order to either update the nest to a “tighter” one (which allows us to recurse), or find the disk Δ' claimed above, or find a small-order carving separation (X, Y) (again defined by some closed disk) at which we may take a carving of our Σ -embedding. If there is a transaction in $(G_{\Delta^{\text{ext}}}, \Omega_{\Delta^{\text{ext}}})$ where a big enough part of its paths also traverse Δ^{int} , then we use this transaction in order to split the vortex into two vortices and recurse. This split is performed using the path infrastructure offered by the transaction, along with the cycles of the railed nest and may result in either a “tighter” H -red railed flat vortex around Δ^{int} or in two H -red railed flat vortices. In both cases, this allows us to recurse. As we know by the redrawing lemma, that k such H -red railed flat vortices may give an H -packing, this procedure will end and will produce less than k H -red railed flat vortices, each with some closed disk Δ' defining a bounded depth vortex, as above.

Killing H -red flat vortices. In the next and final step we exploit all the additional structure we obtained via the refinement step and “attack” each of the obtained H -red railed flat vortices separately. For each of them we “kill” all H -red adhesions residing in its internal disk $\Delta^{\text{int}} \subseteq \Delta'$ by identifying a bounded set of vertices drawn within Δ^{int} .

Towards this, recall that the refinement step ensures that the vortex $(G_{\Delta'}, \Omega_{\Delta'})$ has bounded depth. Using a known result of [19], we construct a *bounded width linear decomposition* of $G_{\Delta'}$, that is a path decomposition $\langle X_1, X_2, \dots, X_n \rangle$ where every bag X_i contains some vertex x_i of the boundary of Δ' in a way that these x_1, \dots, x_n are the vertices of $V(\Omega_{\Delta'})$, appearing in the same order as they appear in $\Omega_{\Delta'}$. We next partition $\langle X_1, X_2, \dots, X_n \rangle$ into r segments $\{\langle X_{p_{i-1}}, \dots, X_{p_i-1}, X_{p_i} \rangle, i \in [r]\}$ each “minimally capable” to host some H -red adhesion from which an H -inflated copy invades G_t . Likewise, we find equally many H -inflated copies in G where the parts drawn inside Δ' are disjoint. Then we bound the number of these segments by proving that they may be extended to an H -packing of size r , inside Δ^{ext} . For this, we use the full power of the redrawing lemma along with the infrastructure offered by the railed nest. As long as there are less than k segments in $\{\langle X_{p_{i-1}}, \dots, X_{p_i} \rangle, i \in [r]\}$ we define a carving separation (X, Y) of G where Y contains the union of all $X_{p_{i-1}} \cup X_{p_i}$, $i \in [r]$ and $X \cap Y$ contains the union of all $(X_{p_{i-1}} \cap X_{p_i}) \cup (X_{p_i} \cap X_{p_{i+1}})$, $i \in [r]$. As the size of $X \cap Y$ depends on k and the width of the decomposition (that is bounded), we have that (X, Y) has bounded order. Therefore, we may take a carving of our Σ -embedding at the carving separation (X, Y) . When this is done for all H -red flat vortices, we know that what remains from G'_t has a Σ -embedding that has no H -red adhesions.

From local to global. Recall that all above steps were applied to an initial torso G_t and, in particular, to the corresponding Σ -embedding of $G'_t = G_t - A_t$. In the end, what we obtained is a new G'_t and A_t and a Σ -embedding of G'_t with no H -red adhesions. The elimination

of H -red adhesions was done by taking successive carvings of the Σ -embedding of G'_t at a bounded number of carving separations (X, Y) , each of bounded order. This came at some cost: By taking these carvings, we added all $X \cap Y$'s to A_t and, moreover, removed all $Y \setminus X$'s from G' . As we already mentioned above, the resulting pair $(V(G_t), A_t)$ is an H -dominion of G , which means that if the non-disk-embeddable part of some H -inflated copy in G intersects $V(G_t)$, then it also intersects A_t . At this point we should forget the initial tree decomposition and just keep in mind that we started with a wall W_t of some torso G_t and we finally computed an H -dominion (X_t, A_t) of G where X_t still maintains a big part of the Dyck grid D_t that is the “essential” part of W_t . This constitutes the proof of a “local structure theorem” that, apart from excluding the Dyck graphs in $\{\mathcal{D}_t^\Sigma \mid \Sigma \in \text{sobs}(\mathbb{S}_H)\}$, assumes that G has no H -packing of size k , and, given a big enough wall W , returns an H -dominion (X, A) of G where the “essential part” of W is intact in X . What we need now is to bring this result to the form of a global structure theorem, that is a new tree decomposition (T, β) where each node t is accompanied by a set $\alpha(t) \subseteq \beta(t)$ where $(\beta(t), \alpha(t))$ is an H -dominion of G . This decomposition may serve as the analogue of the tree decomposition in **Step 2**. For this we use an appropriate application of a recent result in [5].

From connected to disconnected. Given the decomposition (T, β) from above, we may now delete adhesions, as it was performed in **Step 2**. After this, we obtain an \mathcal{Z} -dominion (X, A) of G such that $G - X$ is H -minor-free and $|A|$ is bounded. With some more preprocessing, this decomposition may be used to obtain a separation (X, Y) of G of bounded order where $(X, X \cap Y)$ is an \mathcal{Z} -dominion and $G[Y \setminus X]$ is H -minor-free. Notice that at this point, if H is connected, then we are done. To deal with the case where H is not connected, we set up a recursive algorithm which uses the connected case as the base case and each time it is called, it is called for the union of a smaller number of connected components of H . The final outcome is an H -cover of G whose size depends single-exponentially on the size of the excluded Dyck grids from \mathfrak{D}_H and the size of the maximum H -packing in G .

4 Conclusion and open problems

Obstructions of graph classes. The (minor)-*obstruction set* of a graph class \mathcal{G} , denoted by $\text{obs}(\mathcal{G})$, consists of the minor-minimal elements of $\mathcal{G}_{\text{all}} \setminus \mathcal{G}$. Clearly $\text{obs}(\mathcal{G})$ is an antichain. Moreover, it is finite by Robertson’s & Seymour’s theorem. Obstruction sets permit the following equivalent statement of Theorem 2.

► **Theorem 8.** *Let \mathcal{Z} be an antichain in \mathbb{H}^- and let \mathcal{G} be a minor-closed graph class. \mathcal{Z} has the Erdős-Pósa property in \mathcal{G} if and only if, for every surface $\Sigma \in \text{sobs}(\mathbb{S}_{\mathcal{Z}})$, there exists an obstruction in $\text{obs}(\mathcal{G})$ which is Σ -embeddable.*

Universal obstructions. Let $p : \mathcal{G}_{\text{all}} \rightarrow \mathbb{N}$ be a minor-monotone graph parameter. We say that a set \mathfrak{H} of minor-monotone parametric graphs is a (minor-) *universal obstruction* for p if $p \sim p_{\mathfrak{H}}$ (recall (3) for the definition of $p_{\mathfrak{H}}$). Universal obstruction may serve as canonical representations of graph parameters. (For more on the foundation of universal obstructions of parameters, see [31, 32].) From this point of view, Theorem 2 can be restated follows:

► **Theorem 9.** *For every $\mathcal{Z} \in \mathbb{H}^-$, the set of parametric graphs $\mathfrak{D}_{\mathcal{Z}} = \{\mathcal{D}^\Sigma \mid \Sigma \in \text{sobs}(\mathbb{S}_{\mathcal{Z}})\} \cup \{\langle k \cdot H \rangle_{k \in \mathbb{N}} \mid H \in \mathcal{Z}\}$ is a universal obstruction for both $\text{cover}_{\mathcal{Z}}$ and $1/2\text{-pack}_{\mathcal{Z}}$.*

Given some $\mathcal{Z} \in \mathbb{H}$, for every $k \in \mathbb{N}$, we define $\mathcal{C}_k^{\mathcal{Z}} = \{G \mid \text{cover}_{\mathcal{Z}}(G) \leq k\}$. Theorem 9 (or the equivalent Theorem 8) gives us some valuable information about the obstructions in $\text{obs}(\mathcal{C}_k^{\mathcal{Z}})$, for every k .

Certainly, the simplest antichain in \mathbb{H}^- is the one consisting of the two Kuratowski graphs $\mathcal{K} = \{K_5, K_{3,3}\}$. The parameter $\text{cover}_{\mathcal{K}}$ is the *planarizer number* that is the minimum number of vertices whose removal can make a graph planar. The obstruction $\text{obs}(\mathcal{C}_k^{\mathcal{K}})$ is unknown for every positive value of k and its size is expected to grow rapidly as a function of k (see [7] for an exponential lower bound and [40] for a triply exponential upper bound). The identification of $\text{obs}(\mathcal{C}_k^{\mathcal{K}})$ is a non-trivial problem even for small values of k . In particular, it has been studied extensively for the case where $k = 1$ in [25, 28, 47]. In this direction, Mattman and Pierce conjectured that $\text{obs}(\mathcal{C}_k^{\mathcal{K}})$ contains the $Y\Delta Y$ -families of K_{n+5} and $K_{3^2, 2^n}$ and provided evidence towards this in [11]. Recently, Jobson and Kézdy identified *all* graphs in $\text{obs}(\mathcal{C}_1^{\mathcal{K}})$ of connectivity two in [20], where they also reported that $|\text{obs}(\mathcal{C}_1^{\mathcal{K}})| \geq 401$.

It is easy to see that $\{(k+1) \cdot K_5, (k+1) \cdot K_{3,3}\} \subseteq \text{obs}(\mathcal{C}_k^{\mathcal{K}})$, for every $k \in \mathbb{N}$. Our results, together with the fact that $\text{sobs}(\mathbb{S}_{\mathcal{K}}) = \{\Sigma^{(1,0)}, \Sigma^{(0,1)}\}$, provide the following extra information about $\text{obs}(\mathcal{C}_k^{\mathcal{K}})$: for every $k \in \mathbb{N}$, it contains some graph, say G_k^{t} , embeddable in the torus and some graph, say G_k^{p} , embeddable in the projective plane. Most importantly, our results indicate, that the four-member subset $\{(k+1) \cdot K_5, (k+1) \cdot K_{3,3}, G_k^{\text{t}}, G_k^{\text{p}}\}$ of $\text{obs}(\mathcal{C}_k^{\mathcal{K}})$ is sufficient to determine the approximate behaviour of the planarizer number.

Similar implications can be derived for every $\mathcal{Z} \in \mathbb{H}^-$. For instance, if \mathcal{P} is the Petersen family, we again have that $\text{sobs}(\mathbb{S}_{\mathcal{P}}) = \{\Sigma^{(1,0)}, \Sigma^{(0,1)}\}$. Therefore the parameter defined as the minimum number of vertices to remove so as to make a graph linkless, is approximately characterized by picking only nine graphs of $\text{obs}(\mathcal{C}_k^{\mathcal{P}})$, for every $k \in \mathbb{N}$.

Other examples of surface obstructions corresponding to graphs that are known to be both Kuratowski-connected and shallow-vortex minors are $\text{sobs}(\mathbb{S}_{\{K_5\}}) = \text{sobs}(\mathbb{S}_{\{K_6\}}) = \text{sobs}(\mathbb{S}_{\{M_{2n}\}}) = \{\Sigma^{(1,0)}, \Sigma^{(0,1)}\}$, where M_{2n} is the $2n$ -Möbius ladder,¹¹ for $n \in \mathbb{N}_{\geq 3}$. Two other examples are $\text{sobs}(\mathbb{S}_{\{K_{4,4}\}}) = \{\Sigma^{(1,0)}, \Sigma^{(0,2)}\}$ and $\text{sobs}(\mathbb{S}_{\{K_7\}}) = \{\Sigma^{(1,0)}\}$.

Another implication of our results is the following.

► **Theorem 10.** *For every closed and proper set of surfaces \mathbb{S} , the set of parametric graphs $\mathfrak{V}_{\mathbb{Z}} = \{\mathcal{D}^{\Sigma} \mid \Sigma \in \text{sobs}(\mathbb{S})\} \cup \{(\mathcal{V}_k)_{k \in \mathbb{N}}\}$ is a universal obstruction for $\mathbb{S}\text{-tw}_{\text{apex}}$.*

The theorem above is a direct consequence of the second step in our proof outline, that is the step “From small treewidth modulators to small size modulators”, where we obtain a local structure theorem for graphs excluding the parametric graphs in $\mathfrak{V}_{\mathbb{Z}}$, i.e., graphs where $\text{p}_{\mathfrak{V}_{\mathbb{Z}}}$ is bounded. The parameter $\mathbb{S}\text{-tw}_{\text{apex}}$ (defined in (5)) corresponds to the global version of this theorem. In particular, the equivalence between $\mathbb{S}\text{-tw}_{\text{apex}}$ and $\text{p}_{\mathfrak{V}_{\mathbb{Z}}}$ follows directly by [42, Theorem 5.18] or, alternatively, by applying [5, Theorem 6.17]. Notice that $\mathbb{S}\text{-tw}_{\text{apex}}$ can be seen as a parametric extension of graph embeddability and that the exclusion of shallow-vortex minors is pivotal for its definition. The potential algorithmic applications of $\mathbb{S}\text{-tw}_{\text{apex}}$ are open to investigate.

Notice that for both the equivalences in Theorem 9 and in Theorem 10 we have a single exponential gap which, in turn, determines the gap of our FPT-approximations. Is it possible to reduce this to a polynomial one? Certainly, this requires a polynomial dependency on k and t in Theorem 3. There are two sources of exponentiality in the proof of Theorem 3. The first is in the exclusion of the Dyck grid $\mathcal{D}_t^{\Sigma'}$, for $\Sigma' \in \text{sobs}(\mathbb{S})$ that comes from [43], where we have an exponential dependency on t . This dependency already emerges from the bounds in [19]. On the other hand the exponential dependency on k emerges from the redrawing lemma, where the exponential bound comes from the dependencies of the planar linkage

¹¹ The *Möbius ladder* M_{2n} is formed if we consider a cycle on $2n$ vertices and then connect by edges the n anti-diametrical pairs of vertices. Notice that $M_6 = K_{3,3}$. M_8 is called the *Wagner Graph*.

theorem in [2]. Avoiding these two sources of exponentiality appears to be a hard task. An alternative approach is to try instead to “enlarge” the size of the universal obstructions to obtain a polynomial parametric graph. This would also be desirable for the purposes of better FPT-approximation algorithms.

Going further than $\mathbb{K} \cap \mathbb{V}$. The central question proposed by this work is to chart the threshold of half-integrality when covering and packing graphs as minors. In this paper we resolved this question for every antichain in $\mathbb{K} \cap \mathbb{V}$. The wide open question is whether and how this can be done for more general families of antichains. For this, one needs to prove structure theorems on the exclusion of parameterized graphs of unbounded genus, as those in Figure 2. The challenges that have to be met for this, when going beyond \mathbb{K} , are different from those encountered when going beyond \mathbb{V} . We believe that the proof strategy of our paper can serve as a starting point for both directions towards the general case. The resolution of the general case is highly non-trivial and requires new tools and ideas.

We conclude with a conjecture. Our guess is that when we insist on universal obstructions of bounded genus, then we cannot go much further than the horizon of $\mathbb{K} \cap \mathbb{V}$. Let \mathbb{B} be the set of all antichains consisting of graphs where each of them can be embedded in both the torus and the projective plane. As an example, observe that $\{K_{3,4}\} \in \mathbb{B} \setminus \mathbb{K}$, while $\{K_{3,5}\} \notin \mathbb{B}$. We conjecture the following.

► **Conjecture 11.** *Let \mathcal{Z} be an antichain and let $EP_{\mathcal{Z}} : \mathcal{G}_{\text{all}} \rightarrow \mathbb{N}$ be a graph parameter such that \mathcal{Z} has the Erdős-Pósa property in a minor-closed graph class \mathcal{G} if and only if $EP_{\mathcal{Z}}$ is bounded in \mathcal{G} . Then $\mathcal{Z} \in (\mathbb{K} \cap \mathbb{V}) \cup \mathbb{B}$ if and only if there exists some $g_{\mathcal{Z}}$ such that all universal obstructions of $EP_{\mathcal{Z}}$ consist of parametric graphs of Euler genus $\leq g_{\mathcal{Z}}$.*

References

- 1 Colin C. Adams. *The knot book*. American Mathematical Society, Providence, RI, 2004. An elementary introduction to the mathematical theory of knots, Revised reprint of the 1994 original.
- 2 Isolde Adler, Stavros G. Kolliopoulos, Philipp Klaus Krause, Daniel Lokshtanov, Saket Saurabh, and Dimitrios M. Thilikos. Irrelevant vertices for the planar disjoint paths problem. *Journal of Combinatorial Theory, Series B*, 122:815–843, 2017. doi:10.1016/j.jctb.2016.10.001.
- 3 Julien Baste, Ignasi Sau, and Dimitrios M. Thilikos. A complexity dichotomy for hitting connected minors on bounded treewidth graphs: the chair and the banner draw the boundary. In *ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 951–970, 2020. doi:10.1137/1.9781611975994.57.
- 4 Hans L. Bodlaender. On disjoint cycles. *International Journal of Foundations of Computer Science*, 5(1):59–68, 1994. doi:10.1142/S0129054194000049.
- 5 Rutger Campbell, J. Pascal Gollin, Kevin Hendrey, and Sebastian Wiederrecht. Bipartite treewidth – The structure of non-zero cycles in group-labelled graphs, 2023, Manuscript submitted to SODA 2024, private communication.
- 6 Radu Curticapean and Mingji Xia. Parameterizing the permanent: Hardness for fixed excluded minors. In *Symposium on Simplicity in Algorithms (SOSA)*, pages 297–307. SIAM, 2022. doi:10.1137/1.9781611977066.23.
- 7 Michael J. Dinneen. Too many minor order obstructions (for parameterized lower ideals). *JUCS – Journal of Universal Computer Science*, 3(11):1199–1206, 1997. doi:10.3217/jucs-003-11-1199.
- 8 Walther Dyck. Beiträge zur Analysis situs. *Mathematische Annalen*, 32(4):457–512, 1888. doi:10.1007/BF01443580.
- 9 P. Erdős and L. Pósa. On independent circuits contained in a graph. *Canadian Journal of Mathematics*, 17:347–352, 1965. doi:10.4153/CJM-1965-035-8.

- 10 Michael R Fellows and Michael A Langston. Nonconstructive tools for proving polynomial-time decidability. *Journal of the ACM*, 35(3):727–739, 1988. doi:doi.org/10.1145/44483.44491.
- 11 Erica Flapan, Allison Henrich, Aaron Kaestner, and Sam Nelson. *Knots, links, spatial graphs, and algebraic invariants*, volume 689 of *Contemporary Mathematics*. American Mathematical Society, Providence, RI, 2017. doi:[10.1090/conm/689](https://doi.org/10.1090/conm/689).
- 12 Fedor V Fomin, Daniel Lokshtanov, Fahad Panolan, Saket Saurabh, and Meirav Zehavi. Hitting topological minors is FPT. In *ACM SIGACT Symposium on Theory of Computing (STOC)*, pages 1317–1326, 2020. doi:[10.1145/3357713.3384318](https://doi.org/10.1145/3357713.3384318).
- 13 Cyril Gavaille and Claire Hilaire. Minor-universal graph for graphs on surfaces, 2023. [arXiv:2305.06673](https://arxiv.org/abs/2305.06673).
- 14 Jim Geelen, Bert Gerards, Bruce Reed, Paul Seymour, and Adrian Vetta. On the odd-minor variant of hadwiger’s conjecture. *Journal of Combinatorial Theory, Series B*, 99(1):20–29, 2009. doi:[10.1016/j.jctb.2008.03.006](https://doi.org/10.1016/j.jctb.2008.03.006).
- 15 Petr A. Golovach, Giannos Stamoulis, and Dimitrios M. Thilikos. Hitting topological minor models in planar graphs is fixed parameter tractable. In *ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 931–950, 2020. doi:[10.1137/1.9781611975994.56](https://doi.org/10.1137/1.9781611975994.56).
- 16 Petr A. Golovach, Giannos Stamoulis, and Dimitrios M. Thilikos. Combing a linkage in an annulus, 2022. [arXiv:2207.04798](https://arxiv.org/abs/2207.04798).
- 17 Petr A. Golovach, Giannos Stamoulis, and Dimitrios M. Thilikos. Hitting topological minor models in planar graphs is fixed parameter tractable. *ACM Transactions on Algorithms*, 19(3):23:1–29, 2023. doi:[10.1145/3583688](https://doi.org/10.1145/3583688).
- 18 Tony Huynh, Felix Joos, and Paul Wollan. A unified Erdős-Pósa theorem for constrained cycles. *Combinatorica*, 39(1):91–133, 2019. doi:[10.1007/s00493-017-3683-z](https://doi.org/10.1007/s00493-017-3683-z).
- 19 Ken-ichi Kawarabayashi, Robin Thomas, and Paul Wollan. Quickly excluding a non-planar graph, 2021. [arXiv:2010.12397](https://arxiv.org/abs/2010.12397).
- 20 Adam S. Jobson and André E. Kézdy. All minor-minimal apex obstructions with connectivity two. *Electronic Journal of Combinatorics*, 28(1):1.23, 58, 2021. doi:[10.37236/8382](https://doi.org/10.37236/8382).
- 21 Naonori Kakimura, Ken-ichi Kawarabayashi, and Dániel Marx. Packing cycles through prescribed vertices. *Journal of Combinatorial Theory, Series B*, 101(5):378–381, 2011. doi:[10.1016/j.jctb.2011.03.004](https://doi.org/10.1016/j.jctb.2011.03.004).
- 22 Ken-ichi Kawarabayashi. Half integral packing, Erdős-Pósa-property and graph minors. In *ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1187–1196, 2007. URL: dl.acm.org/citation.cfm?id=1283383.1283511.
- 23 Ken-ichi Kawarabayashi, Yusuke Kobayashi, and Bruce A. Reed. The disjoint paths problem in quadratic time. *Journal of Combinatorial Theory, Series B*, 102(2):424–435, 2012. doi:[10.1016/j.jctb.2011.07.004](https://doi.org/10.1016/j.jctb.2011.07.004).
- 24 Ken-ichi Kawarabayashi and Stephan Kreutzer. The directed grid theorem. In *ACM Symposium on Theory of Computing (STOC)*, pages 655–664, 2015. doi:[10.1145/2746539.2746586](https://doi.org/10.1145/2746539.2746586).
- 25 Max Lipton, Eoin Mackall, Thomas W. Mattman, Mike Pierce, Samantha Robinson, Jeremy Thomas, and Ilan Weinschelbaum. Six variations on a theme: almost planar graphs. *Involve. A Journal of Mathematics*, 11(3):413–448, 2018. doi:[10.2140/involve.2018.11.413](https://doi.org/10.2140/involve.2018.11.413).
- 26 Chun-Hung Liu. Packing topological minors half-integrally. *Journal of the London Mathematical Society*, 106(3):2193–2267, 2022. doi:[10.1112/jlms.12633](https://doi.org/10.1112/jlms.12633).
- 27 Dániel Marx and Ildikó Schlotter. Obtaining a planar graph by vertex deletion. *Algorithmica*, 62(3-4):807–822, 2012. doi:[10.1007/s00453-010-9484-z](https://doi.org/10.1007/s00453-010-9484-z).
- 28 Thomas W. Mattman. Forbidden minors: finding the finite few. In *A primer for undergraduate research*, Found. Undergrad. Res. Math., pages 85–97. Birkhäuser/Springer, Cham, 2017.
- 29 Laure Morelle, Ignasi Sau, Giannos Stamoulis, and Dimitrios M. Thilikos. Faster Parameterized Algorithms for Modification Problems to Minor-Closed Classes. In Kousha Etessami, Uriel Feige, and Gabriele Puppis, editors, *50th International Colloquium on Automata, Languages, and Programming (ICALP 2023)*, volume 261 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 93:1–93:19, Dagstuhl, Germany, 2023. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:[10.4230/LIPIcs.ICALP.2023.93](https://doi.org/10.4230/LIPIcs.ICALP.2023.93).

- 30 Sergey Norin, Robin Thomas, and Hein van der Holst. On 2-cycles of graphs. *Journal of Combinatorial Theory, Series B*, 162:184–222, 2023. doi:10.1016/j.jctb.2023.06.003.
- 31 Christophe Paul, Evangelos Protopapas, and Dimitrios M. Thilikos. Graph parameters, universal obstructions, and wqo, 2023. arXiv:2304.03688.
- 32 Christophe Paul, Evangelos Protopapas, and Dimitrios M. Thilikos. Universal obstructions of graph parameters, 2023. arXiv:2304.14121.
- 33 Jean-Florent Raymond and Dimitrios M. Thilikos. Recent techniques and results on the Erdős-Pósa property. *Discrete Applied Mathematics*, 231:25–43, 2017. doi:10.1016/j.dam.2016.12.025.
- 34 Bruce Reed. Mangoes and blueberries. *Combinatorica*, 19(2):267–296, 1999. doi:10.1007/s004930050056.
- 35 Bruce Reed, Neil Robertson, Paul Seymour, and Robin Thomas. Packing directed circuits. *Combinatorica*, 16(4):535–554, 1996. doi:10.1007/BF01271272.
- 36 Neil Robertson and P. D. Seymour. Graph minors. V. Excluding a planar graph. *Journal of Combinatorial Theory, Series B*, 41(1):92–114, 1986. doi:10.1016/0095-8956(86)90030-4.
- 37 Neil Robertson and P. D. Seymour. Graph minors. XX. Wagner’s conjecture. *Journal of Combinatorial Theory, Series B*, 92(2):325–357, 2004. doi:10.1016/j.jctb.2004.08.001.
- 38 Neil Robertson, Paul Seymour, and Robin Thomas. Sach’s linkless embedding conjecture. *Journal of Combinatorial Theory, Series B*, 64(2):185–227, 1995. doi:10.1006/jctb.1995.1032.
- 39 Ignasi Sau, Giannos Stamoulis, and Dimitrios M. Thilikos. k -apices of minor-closed graph classes. II. Parameterized algorithms. *ACM Transactions on Algorithms*, 18(3):Art. 21, 30, 2022. doi:10.1145/3519028.
- 40 Ignasi Sau, Giannos Stamoulis, and Dimitrios M. Thilikos. k -apices of minor-closed graph classes. I. Bounding the obstructions. *Journal of Combinatorial Theory, Series B*, 161:180–227, 2023. doi:10.1016/j.jctb.2023.02.012.
- 41 Dimitrios M. Thilikos and Sebastian Wiederrecht. Killing a vortex. In *IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 1069–1080, 2022. doi:10.1109/FOCS54457.2022.00104.
- 42 Dimitrios M. Thilikos and Sebastian Wiederrecht. Approximating branchwidth on parametric extensions of planarity, 2023. arXiv:2304.04517.
- 43 Dimitrios M. Thilikos and Sebastian Wiederrecht. Excluding surfaces as minors in graphs, 2023. arXiv:2306.01724.
- 44 Robin Thomas. Well-quasi-ordering infinite graphs with forbidden finite planar minor. *Transactions of the American Mathematical Society*, 312(1):279–313, 1989. doi:10.2307/2001217.
- 45 Wouter Cames Van Batenburg, Tony Huynh, Gwenaél Joret, and Jean-Florent Raymond. A tight Erdős-Pósa function for planar minors. In *ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1485–1500, 2019. doi:10.1137/1.9781611975482.90.
- 46 Hein van der Holst. A polynomial-time algorithm to find a linkless embedding of a graph. *Journal of Combinatorial Theory, Series B*, 99(2):512–530, 2009. doi:10.1016/j.jctb.2008.10.002.
- 47 Yaming Yu. More forbidden minors for Wye-Delta-Wye reducibility. *Electronic Journal of Combinatorics*, 13(1):7:15, 2006. doi:10.37236/1033.