Quickly Excluding an Annotated Planar Graph

Maximilian Gorsky
Discrete Mathematics Group, Institute for Basic Science (IBS), Daejeon, South Korea

Evangelos Protopapas
Faculty of Mathematics, Informatics and Mechanics, University of Warsaw, Poland

Sebastian Wiederrecht
School of Computing, KAIST, Daejeon, South Korea

Abstract

We provide proofs certifying that the structure theorem for vertex sets of bounded bidimensionality holds with polynomial bounds. The bidimensionality of vertex sets is a common generalisation of both treewidth and the face-cover-number of vertex sets in planar graphs. As such, it plays a crucial role in extensions of Courcelle's Theorem to H-minor-free graphs. Recently, bidimensionality and similar parameters have emerged as key for extensions of known parameterized algorithms for problems defined on a terminal set R. A prominent example for such a problem is Steiner Tree, which admits efficient algorithms on planar graphs whenever R can be covered with few faces.

Key to the algorithmic applications of bidimensionality is a structure theorem that explains how a graph G can be decomposed into pieces where the behaviour of R is highly controlled. One may see this structure theorem as a rooted analogue of Robertson and Seymour's celebrated Grid Theorem. Combining recent advances in obtaining polynomial bounds in the Graph Minors framework with new techniques for handling annotated vertex sets, we show that all parameters in the structure theorem above admit polynomial bounds. As an application, we also provide a sketch showing how our techniques imply polynomial bounds for the structure theorem for graphs excluding an apex minor.

2012 ACM Subject Classification Mathematics of computing → Combinatorics; Mathematics of computing → Graph theory

Keywords and phrases Structural Graph Theory, Graph Minors, Annotated Graphs, Rooted Minors, Colorful Minors, Bidimensionality

Digital Object Identifier 10.4230/LIPIcs.ICALP.2026.99

Category Track A: Algorithms, Complexity and Games

Related Version *Full Version*: `https://arxiv.org/abs/2602.06516`

Funding *Maximilian Gorsky*: Supported by the Institute for Basic Science (IBS-R029-C1).
Evangelos Protopapas: Supported by the ERC project BUKA (n° 101126229).

1 Introduction

A central approach to dealing with computational intractability in graph problems is offered by *structural graph theory*. Through a wide range of structural notions and results, it supports the systematic design of efficient algorithms for hard problems on well-behaved graph classes. A particularly powerful toolkit from structural graph theory is the design of *graph parameters* capturing key features that facilitate the design of efficient algorithms. The algorithmic study of such graph parameters makes up a rich subfield of *parameterized algorithms*. A prime example of this interplay between structural and algorithmic graph theory is the parameter *treewidth* popularised by Robertson and Seymour [46] (see for example [2, 7, 9, 35]). While treewidth is a powerful tool for the design of parameterized algorithms for a wide range of problems, it also naturally gives rise to a problem: For which classes of problems exist structural parameters that allow for the design of parameterized algorithms *beyond* the regime of treewidth?

An *annotated graph* is a pair (G, R) where G is a graph and $R \subseteq V(G)$ is the set of *annotated vertices*. In the following, we will refer to R as the set of *red* vertices. Recently, a family of parameters has emerged that aims to target problems defined on annotated graphs such as the Steiner Tree problem [25, 8, 37, 42]. The theme of such parameters is to restrict the structural properties of the red vertices instead of the global structure of the graph. This is motivated by the observation that for many such problems, treewidth already defines their tractability horizon within minor-closed graph classes. To be more precise, the seminal Grid Theorem of Robertson and Seymour [47] says that a minor-closed graph class $\mathcal{G}$ has bounded treewidth if and only if $\mathcal{G}$ does not contain all planar graphs. Often problems like Steiner Tree are NP-hard already on planar graphs [14, 23, 34], implying that such problems are tractable on a minor-closed graph class $\mathcal{G}$ if and only if $\mathcal{G}$ has bounded treewidth. In the emerging theory of *colorful minors*, sometimes called *rooted minors*, several positive algorithmic results [25, 8, 37, 42, 29, 32] hint at a paradigm providing a way to escape such dichotomies:

Define parameters to restrict the structure relative to R instead or in addition to restricting the structure of G as a whole.

To better capture the situation where a planar graph "rooted" on a fixed set R of vertices is excluded, Thilikos and Wiederrecht defined the notion of *bidimensionality* [56]. The *bidimensionality* of an annotated graph (G, R) is the largest integer k such that there exists a minor-model[1] $\mathcal{X}$ of the $(k \times k)$-grid where $R \cap V(X) \neq \emptyset$ for all branch sets $X \in \mathcal{X}$. We call such a grid minor a *red grid*. See Figure 2 for an illustration.

The notion of bidimensionality has since found key applications in the algorithmic theory of model checking in H-minor-free graphs [55, 53]. On the structural side, Protopapas, Thilikos, and Wiederrecht [45] proved a structure theorem providing an approximate description of annotated graphs of small bidimensionality akin to Robertson and Seymour's duality between grid minors and treewidth. However, the bounds provided by Protopapas et al. for their structure theorem depend exponentially on the bidimensionality.

Our main contribution. We give an independent proof for this structure theorem providing, for the first time, *polynomial bounds* for it. Our proof is constructive, yielding a fixed-parameter tractable algorithm that either decides that the bidimensionality of (G, R) is at least k or finds the structural decomposition from [45], where all involved parameters are in $k^{\mathbf{O}(1)}$.

1.1 Our result

To state our main theorem, we require some additional definitions from Robertson and Seymour's theory of graph minors. This is because in certain situations, the number of red vertices may be unbounded *and* they might be highly connected to each other, but the way they attach to the rest of the graph is restricted in other ways. To see this, consider for

[1] A *minor-model* of a graph H in a graph G is a collection $\{G_v\}_{v \in V(H)}$ of pairwise vertex-disjoint connected subgraphs of G, called the *branch sets*, such that for all $uv \in E(H)$ there is an edge in G between $V(G_u)$ and $V(G_v)$.

example a large grid where only the very first column is red (see Figure 1). In this annotated graph the number of red vertices is unbounded, at the same time it is impossible to decompose it along small-order separations in a way that distributes the red vertices into smaller pieces. However, the entire graph is planar and all red vertices sit on a single face. In general the structure emerging from bounding the bidimensionality may be more complicated, but this example illustrates that in order to grasp the structural aspects of bidimensionality, one must account for certain topological conditions. To describe the separation into well-behaved pieces, we make use of the notion of tree-decompositions.

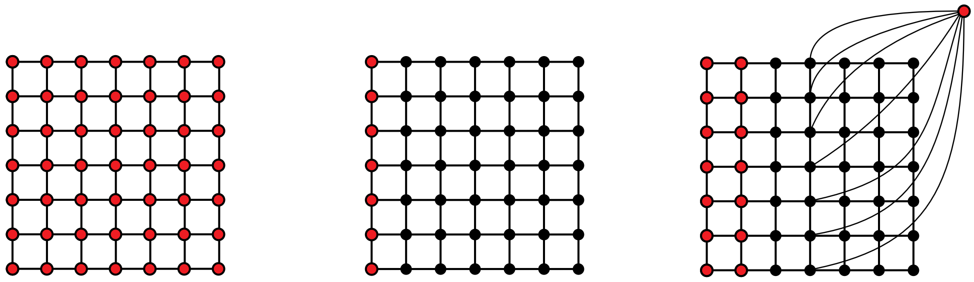

(i) (ii) (iii)

Figure 1 Diagrams of (i) the red (7×7)-grid and two examples illustrating that the structure emerging from excluding a red grid resembles the structure of excluding an arbitrary minor. (ii) A grid with its outermost column being red: no large red grid exists but at the same time the red vertices cannot be easily removed from the grid. (iii) a more complicated example where embeddability can only be achieved by deleting an apex vertex and few faces do not suffice to describe the structure of the red vertices.

A *tree-decomposition* for a graph G is a pair (T, β) such that T is a tree, $\beta \colon V(T) \to 2^{V(G)}$ assigns to each node of T a subset of the vertices of G known as a *bag*, $\bigcup_{t \in V(T)} G[\beta(t)] = G$, and for each $v \in V(G)$, the set $\{t \in V(T) \colon v \in \beta(t)\}$ is connected. The *adhesion* of (T, β) is 0 if T has only one node and $\max_{dt \in E(T)} |\beta(d) \cap \beta(t)|$ otherwise. The *width* of (T, β) is defined as $\max_{t \in V(T)} |\beta(t)| - 1$.

The topological condition we require is more complicated. Notice that in the example above, if we would also colour the second column in red, the bidimensionality of the resulting annotated graph would still be bounded by a constant, independent of the order of the grid. Moreover, even if we add one additional red vertex and make it adjacent to the entire centre column of the grid, we would only increase the bidimensionality by 1 for all such grids. But the resulting graph cannot be embedded into a surface with reasonable Euler-genus. Thus we must allow for the deletion of a small vertex set in order to reveal topological behaviour and we need to relax our condition on what it means to cover the red vertices with few faces. See Figure 1 for an illustration.

We say that a graph G has a *k-near embedding* in a surface Σ if there exists $A \subseteq V(G)$ with $|A| \leq k$, such that $G - A = G_0 \cup G_1 \cup \cdots \cup G_\ell$, $\ell \leq k$, such that G_0 has an embedding into Σ with ℓ pairwise vertex-disjoint faces F_i where for all $i \in \{1, \ldots, \ell\}$, $V(G_0) \cap V(G_i) = V(F_i)$, for $i \in \{1, \ldots, \ell\}$, the G_i's are pairwise vertex-disjoint and each such G_i has a path-decomposition[2] (P_i, β_i) of width at most k such that $V(P_i) = V(F_i)$, the vertices of P_i appear in agreement to their cyclic ordering on the boundary of F_i, and $v \in \beta_i(v)$, for all $v \in V(F_i)$. The set A is called the *apex set*, the graphs G_i, $i \in \{1, \ldots, \ell\}$, are called the *vortices*, and the sets $V(G_i) \setminus V(F_i)$ are the *interiors* of the vortices.

[2] A *path-decomposition* of a graph is a tree-decomposition (T, β) where T is a path.

In the context of annotated graphs, we need some additional information regarding the red vertices. Let (T, β) be a tree decomposition of an annotated graph (G, R). The *annotated torso* (G_t, R_t) of (G, R) at a node $t \in V(T)$ is the annotated graph obtained from $(G[\beta(t)], R)$ by first making the set $\beta(d) \cap \beta(t)$ into a clique, for every $dt \in E(T)$, and then marking all vertices in $\beta(d) \cap \beta(t)$ red, only when the part of the decomposition hanging below d contains at least one red vertex, i.e. whenever $\left(\bigcup_{d' \in V(T_d)} \beta(d')\right) \cap R \neq \emptyset$, where T_d is the unique component of $T - dt$ that contains d.

With these definitions, our main theorem reads as follows.

▶ **Theorem 1.** *There exists a function $f_1 \colon \mathbb{N} \to \mathbb{N}$ such that for all non-negative integers k, and all annotated graphs (G, R) one of the following holds:*

1. *(G, R) has bidimensionality at least k, or*
2. *(G, R) has a tree-decomposition (T, β) of adhesion at most $f_1(k)$ such that for all $t \in V(T)$, either t is a leaf with the parent d and $(\beta(t) \setminus \beta(d)) \cap R = \emptyset$, or the annotated torso (G_t, R_t) of (G, R) at t has an $f_1(k)$-near embedding in a surface of Euler-genus at most $f_1(k)$ and every vertex $v \in \beta(t) \cap R$ belongs to the apex set or the interior of some vortex.*

Moreover, $f_1(k) \in k^{\mathbf{O}(1)}$ and there exists an algorithm that finds either a $(k \times k)$-grid-minor-model witnessing that (G, R) has bidimensionality at least k, or a tree-decomposition as above in time $2^{k^{\mathbf{O}(1)}} \cdot |E(G)|^3 |V(G)| \log(|V(G)|)$.

1.2 Related work, consequences, and applications

Generalising treewidth through annotation

Notice that by our definition, a minor-model of a red $(k \times k)$-grid does not have to be a *minimal* minor-model of a $(k \times k)$-grid. That is, we explicitly allow for the red vertices to be connected to the grid in the form of dangling paths. See Figure 2 for an example. This notion of "red minor" has already been studied by a variety of authors, often under the name of "rooted minors". However, depending on the context, rooted minors are sometimes required to respect a fixed bijection between the coloured vertices of (G, R_G) and the coloured vertices of (H, R_H). For better distinction between the two concepts, we say that (H, R_H) is a *red minor* of (G, R_G) if there exists a minor-model $\{G_v\}_{v \in V(H)}$ of H in G such that $R_G \cap V(G_v) \neq \emptyset$ for all $v \in R_H$. This is in reference to the term "colorful minor" introduced by Protopapas, Thilikos, and Wiederrecht [45] which allows for vertices to carry more than one colour.

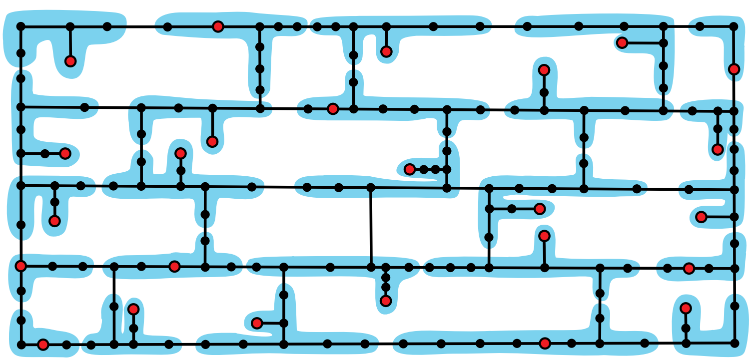

Figure 2 A model of a red (5×5)-grid.

The work of Protopapas et al. provides several structural theorems for annotated graphs excluding fundamental patterns, including a variant of Theorem 1 with exponential bounds, a structure theorem for excluding a red clique, and one for the exclusion of an *outer red grid*, i.e. an annotated graph (G, R) where G is a $(k \times k)$-grid for some k and R is the vertex set of its first column. A small outer red grid can be found in Figure 1. While Protopapas et al. investigate a variant of the outer red grid with more than one colour, the first result describing the structure of graphs excluding an outer red grid is due to Marx, Seymour, and Wollan [39]. It was later observed by Hodor, La, Micek, and Rambaud [31] that the result of Marx et al. provides a min-max characterisation of a type of modulation parameter that falls into a family of structural parametrisations started by Bulian and Dawar with their introduction of *elimination distance* [4, 5].

The *treewidth* of a graph G, denoted by $\mathsf{tw}(G)$, is the smallest integer k such that G has a tree-decomposition of width at most k.

Let G be a graph and $X \subseteq V(G)$ be a vertex set. The *torso* of G at X is the graph G_X obtained from $G[X]$ by turning the set $N_G(J) \subseteq X$ into a clique for every component J of $G - X$.

The *torso treewidth* of an annotated graph (G, R), denoted by $\mathsf{ttw}(G, R)$, is the smallest integer k such that there exists a set $X \subseteq V(G)$ with $R \subseteq X$ and $\mathsf{tw}(G_X) \leq k$.

As observed by Hodor et al. [31], the result of Marx et al. [39] implies that every annotated graph with large torso treewidth contains a big outer red grid as a red minor. One may observe a hierarchy of parameters as follows:

1. The Grid Theorem of Robertson and Seymour explains the structure of graphs excluding a planar graph as minor,
2. the theorem of Marx et al. explains the structure of annotated graphs excluding a planar graph with a single red face as a red minor, and
3. Theorem 1 explains the structure of an annotated graph excluding an arbitrary annotated planar graph as a red minor.

Among these Theorem 1 is the most general, as it in particular implies Robertson and Seymour's Grid Theorem, since the treewidth of any graph G equals the bidimensionality of $(G, V(G))$. Moreover, due to torso treewidth being sandwiched between treewidth and bidimensionality, in annotated graphs where all vertices are red all three parameters collapse into one.

An application: Excluding an apex minor

A graph H is said to be an *apex graph* if there exists some vertex $v \in V(H)$ such that $H - v$ is planar. Apex graphs themselves are very close to planar graphs. However, apex-minor-free graph classes generalise planar graphs and any graph class of bounded Euler-genus. Graphs in these classes allow for a range of interesting algorithmic results [22, 21, 30, 11, 13, 12, 24, 36]. Many of these algorithmic applications are in the realm of approximation algorithms. Indeed, as proposed by Eppstein [22, 21] and later built upon by Grohe [30] and Demain and Hajiaghayi [11], apex-minor-free graphs have bounded "local treewidth" which allows for an application of a variant of Baker's Technique [1], facilitating the design of polynomial approximation schemes (PTAS) for problems which are hard to approximate on general graphs.

Indeed, while the structure theorem for apex-minor-free graphs was known in the community, the first written proof for it can be found in the appendix of a paper by Dvořák and Thomas [20] on approximating list colourings. Since their proof builds on the original results of Robertson and Seymour, they do not come with any estimates on the bounds for the constants describing the running times of their algorithm. This issue is shared among many of the applications mentioned above and gets explicitly mentioned by Grohe [30]. Our result shows that, in the case of apex-minor-free graphs, it is possible to give explicit polynomials bounds directly through the use of modern graph minor theory.

By interpreting the red vertices in each step of the proof as the combined neighbourhood of the apex vertices produced, our techniques show that one may either find any fixed apex graph as a minor, or create a region in the surface-embedded part which avoids the neighbourhood of all apex vertices. This yields a polynomial version of Dvořák and Thomas' *local* structure theorem. The term "local" here refers to the fact that the structure theorem is stated with respect to a large wall. A local version of Theorem 1 can be found in the form of Theorem 5.1 in the full version of our article.

We say that a graph G containing a wall $W \subseteq G$ has a *weak k-near embedding centred at W* in a surface Σ if there exists $A \subseteq V(G)$ with $|A| \leq k$, such that $G - A = G_0 \cup G_1 \cup \cdots \cup G_\ell \cup J_1 \cup \cdots \cup J_h$, $\ell \leq k$ and $h \in \mathbb{N}$, such that G_0 has an embedding into Σ with $\ell + h$ faces F_i where

- the graphs $G_i - V(G_0)$ and $J_j - V(G_0)$ are vertex-disjoint for all $i \in \{1, \ldots, \ell\}$ and $j \in \{1, \ldots, h\}$,
- for all $j \in \{1, \ldots, h\}$, the face F_j contains a set S_j of at most three vertices such that all S_j's are pairwise disjoint and $V(J_j) \cap V(F_j) = S_j = V(J_j) \cap V(G_0)$, and such that, if $|S_j| = 3$, $V(F_j) = S_j$ and if $|S_j| = 2$ then the vertices of S_j are adjacent on F_j,
- for all $i \in \{h+1, \ldots, h+\ell\}$, $V(G_0) \cap V(G_i) = V(F_i)$, for $i \in \{1, \ldots, \ell\}$, the G_i's are pairwise vertex-disjoint and each such G_i has a path-decomposition (P_i, β_i) of adhesion at most k such that $V(P_i) = V(F_i)$, the vertices of P_i appear in agreement to their cyclic ordering on the boundary of F_i, and $v \in \beta_i(v)$, for all $v \in V(F_i)$, and
- W is vertex-disjoint from $\bigcup_{i=1}^{\ell} G_i$ and for each $j \in \{1, \ldots, h\}$, J_j contains at most one vertex of degree 3 from W.

The set A is called the *apex set*, the graphs G_i are called the *vortices*, and the sets $V(G_i) \setminus V(F_i)$ are the *interiors* of the vortices and the graphs J_j are called the *flaps*.

▶ **Theorem 2.** *There exist functions $g_2 \colon \mathbb{N}^2 \to \mathbb{N}$ and $f_2 \colon \mathbb{N} \to \mathbb{N}$ such that for all non-negative integers $r \geq 3$ and k, every apex graph H on at most k vertices, every graph G, and every w-wall $W \subseteq G$ with $w \geq g_2(k, r)$ one of the following holds:*

1. *G contains H as a minor, or*
2. *there exists an r-subwall $W' \subseteq W$ such that G has a weak $f_2(k)$-near embedding centred at W' such that all neighbours of the apex set are contained in the interiors of the vortices.*

Moreover, $g_2(r, k) \in (r + k)^{\mathbf{O}(1)}$, $f_2(k) \in k^{\mathbf{O}(1)}$, and there exists an algorithm that takes as input, G, W, H, and r as above and finds either a minor model of H, or an r-subwall W' together with a weak $f_2(k)$-near embedding centred at W' for G in time $\mathbf{poly}(|V(H)| + r + k)|E(G)|^3$.

It is important to stress that Theorem 2 is a consequence of our *methods* and the lemmas we develop along the way, not directly of the main statements that eventually give rise to Theorem 1.

Moreover, the reason why we only state the local variant of Dvořák and Thomas' structure theorem for apex-minor-free graphs here is three-fold.

First, almost all Robertson-Seymour-style structure theorems are proven in this way: One first proves a local variant with respect to a wall or "tangle" and then applies a well-known technique to turn the local structure theorem into a global one based on tree decompositions – see Section 6 of the fullversion of our article for this local-to-global step proving Theorem 1. Novel arguments are usually only necessary to prove the local theorems, with the local-to-global step seeing few innovations over the years.

Second, while the global theorems are better known, the local structure theorems tend to be the more important ones. Robertson and Seymour point this out in their own proof of the Graph Minor Structure Theorem [52], declaring the global structure theorem to be a "red herring".

Finally, the third reason is of a technical nature. While the process of going from local to global is by now routine, its approximate nature comes with the drawback of introducing additional vertices to the apex set. Dvořák and Thomas have laid out a technique to capture the neighbourhoods of these new apices inside additional vortices [20] (see also [41] for a more recent take on the technique). However, this step introduces additional technicalities which we believe to be outside the scope of this paper and so we leave this part to the interested readers.

The algorithmic potential of annotated graph parameters

As described above, the notion of bidimensionality for annotated vertex sets may be seen as a vast structural generalisation of the face-cover-number of annotated vertex sets in planar graphs. A wide range of problems defined on annotated graphs are known to be tractable on planar annotated graphs whenever the face-cover-number is bounded. Interesting problems that fall into this framework are the Multiway Cut problem [42] and the problem of counting perfect matchings with defects [8]. Even the Disjoint Paths problem on planar graphs exhibits strong algorithmic properties when the face-cover-number of its terminals is considered [40, 58]. A leading question in this newly arising theory of structural parameters for annotated graphs is the following:

> Given an NP-hard computational problem Π defined on annotated graphs which is tractable on planar graphs where the red vertices can be covered by a constant number of faces, which proper red-minor-closed classes of annotated graphs admit polynomial-time algorithms for Π?

A prime example of such a behaviour is the Steiner Tree problem [34]. Recently, Steiner Tree has become the subject of a streamlined investigation into the power of structural parameters for annotated graphs. Jansen and Swennenhuis [32] showed that Steiner Tree is fixed-parameter tractable when parameterized by the torso treewidth of the annotated input instance. Groenland, Nederlof, and Koana [29] have shown that excluding a red K_4-minor also gives rise to a class of polynomially solvable instances of Steiner Tree. Notice that, due to the structural result of Marx et. al, annotated graphs excluding a red K_4-minor may have unbounded torso treewidth, however, since the red K_4 is a minor of the red (3×3)-grid, excluding K_4 as a red minor implies bounded bidimensionality. Given that Steiner Tree is NP-hard on the class of all annotated planar graphs [26], it follows that it is also NP-hard on all red-minor-closed classes of unbounded bidimensionality. However, there is strong evidence, that bidimensionality might be what precisely delineates the tractability of Steiner Tree in red-minor-closed classes.

> Are there computable functions f and g such that Steiner Tree can be solved in time $f(k)n^{g(k)}$ on annotated graphs of bidimensionality at most k?

Due to a result of Krisfaludi-Bak, Nederlof, and Leeuwen [34], the dependency on k in the exponent of n in Section 1.2 is unavoidable assuming the Exponential Time Hypothesis, a feature that is shared by many problems of the same flavour. Key towards providing a positive answer to Section 1.2 is a full understanding of the structural properties of annotated graphs of low bidimensionality, which is one of the main motivations behind our work. Indeed, in case the answer to Section 1.2 is "yes", the functions f and g are likely to depend on the function from Theorem 1, further emphasising the need for constructive and good bounds.

Apart from its focal role in understanding structural parameters for annotated graphs, bidimensionality has also been observed to be a key player in recent extensions of Courcelle's Theorem [7]. Courcelle's Theorem states that the model checking problem for MSO_2-formulas – naively speaking, formulas where one is allowed to quantify over vertex sets and edge sets – is fixed-parameter tractable on graphs of bounded treewidth. In a vast generalisation of this result, Sau, Stamoulis, and Thilikos [53] have shown that in the space of minor-closed graph classes one may extend the applicability of Courcelle's Theorem beyond the scope of graphs of bounded treewidth by restricting the quantifications allowed in the formulas to vertex sets of bounded bidimensionality and certain disjoint-paths queries. Strengthenings of this "meta-algorithmic" result have been proven for annotated graphs of bounded torso treewidth and annotated graphs excluding a red clique minor by Protopapas, Thilikos, and Wiederrecht [45].

Layered parameters and the exclusion of "apex-$\mathcal{X}$" graphs

As manifested in Theorem 2, structure theory for excluding red minors in annotated graphs has direct implications for the exclusion of certain types of graphs H in the setting of graph minors. Let $\mathcal{X}$ be a minor-closed graph class. We say that a graph H is *apex-$\mathcal{X}$* if there exists a vertex $v \in V(H)$ such that $H - v \in \mathcal{X}$. In this terminology, an apex graph may also be called an apex-planar graph.

A *layering* of a graph G is a sequence $\langle L_i \rangle_{i \in \mathbb{N}}$ such that $\bigcup_{i \in \mathbb{N}} L_i = V(G)$ and for every edge $uv \in E(G)$ there is $i \in \mathbb{N}$ such that $u, v \in L_i \cup L_{i+1}$.

A layered parameter is a generalisation of a decomposition-based graph parameter evaluated over all possible decompositions and layerings for a given graph. For example, the *layered treewidth* of a graph G is the minimum value of $\max_{i \in \mathbb{N}} \max_{t \in V(T)} |L_i \cap \beta(t)|$ taken over all layerings $\langle L_i \rangle_{i \in \mathbb{N}}$ and all tree-decompositions (T, β) of G.

Recently, layered graph decompositions have emerged as a powerful tool for solving a number of long-standing open problems even beyond the scope of minor-closed graph classes and including key results in the theory of the *clustered chromatic number* [18, 17, 15, 54, 19, 16, 3, 38, 10, 31]. In many cases, the existence of such layered decompositions of small width is implied by the absence of some apex-$\mathcal{X}$ graph for a carefully chosen class $\mathcal{X}$. Indeed, Dujmović, Morin, and Wood [18] showed that a minor-closed graph class has bounded layered treewidth if and only if it excludes some apex-planar graph and their proof makes use of the structure theorem of Dvořák and Thomas. Hence, Theorem 2 has direct implications for the bounds found by Dujmović et al. [18]. To obtain a better understanding for the structure of excluding an apex-$\mathcal{X}$ graph as a minor, it is often easier to understand the structure of excluding an annotated graph (H, R) with $H \in \mathcal{X}$ as a red minor [15, 31, 6]. Our main theorem shows that, whenever $\mathcal{X}$ is a class of planar graphs, all involved bounds are polynomial.

Towards colorful minors of q-colorful graphs

The original structure theorem for annotated graphs of bounded bidimensionality due to Protopapas, Thilikos, and Wiederrecht [45] had one additional powerful feature which our polynomial version is lacking: Instead of considering only one colour, Protopapas et al.

allowed for the presence of up to q colours and considered the situation where they exclude a $(k \times k)$-grid in which every vertex carries all q colours. As a result of this generality, their bound is of the form $2^{k^{\mathbf{O}(1)} 2^{2^{\mathbf{O}(q)}}}$. The main reason for this super-exponential dependency on q is the lack of tools which are able to "homogenise" a given flat wall with respect to q distinct colours within polynomial bounds. It should be mentioned that very recently Gorsky, Seweryn, and Wiederrecht [28] introduced a technique that provides precisely such a homogenisation procedure. However, as Gorsky et al. explain in the conclusion of their paper, a variant of Theorem 1 for q colours realising polynomial bounds is still out of reach for now. The reason is that homogenising flat walls is not enough. As explained in Section 1.3, a crucial step is the homogenisation of so-called "transactions" – a notion from deep within the proof of the Graph Minor Structure Theorem (see [48, 33, 27]). A tool for efficiently dealing with the homogenisation of transactions is still missing at the time of writing.

1.3 Overview of our proof

The proof of Theorem 1 is divided into four main steps as follows.

Step 1: A variant of the Flat Wall Theorem confining all red vertices to the outside of the flat wall, which for purpose of presentation we call the *Red Flat Wall Theorem*.

Step 2: A version of the Society Classification Theorem ensuring that all new additions to the weak near embedding are free of red vertices, which we call the *Red Society Classification Theorem*.

Step 3: The local version of Theorem 1 with respect to a large wall, which we call the *Red Local Structure Theorem*.

Step 4: The globalisation of the outcome of **Step 3**, yielding Theorem 1.

The first three steps combined make up the proof of the local version of Theorem 1, that is Theorem 5.1 in the full version, while **Step 4** utilises the red local structure theorem in order to construct the tree-decomposition mentioned in Theorem 1. The proof of Theorem 5.1 in the full version inductively generates a near embedding. The proof uses the outcome of the red flat wall theorem to enter the base case of the induction and then repeatedly applies the outcome of the red society classification theorem in order to either conclude or increase the Euler-genus of the underlying surface in the near embedding.

Step 1: A flat wall

The celebrated Flat Wall Theorem of Robertson and Seymour [51] states that given any large enough wall W in a graph G, either there exists a large clique-minor in G whose model is highly connected to W, or there exists a weak near embedding of G centred at W in the sphere with a single vortex, without any bound on the adhesion of the path decomposition of the vortex. The outcome of this first step is, that in the absence of a big red grid, one may ensure that all red vertices are confined into the vortex and the apex set.

This part is relatively straightforward. We start with a large wall W_0 and apply a polynomial variant of the Flat Wall Theorem – we use the one of Gorsky et al. [27] – to obtain either a large clique or a small apex set A_1 together with a big flat subwall W_1 of W_0. A result of Protopapas, Thilikos, and Wiederrecht [45] allows us to dismiss the case where we find the clique. Otherwise we subdivide W_1 into a large number of pairwise disjoint subwalls arranged in a grid-like shape. If each of them contains a red vertex in its interior, we find a large red grid, otherwise one of them is the desired flat wall W_2 without any red vertices. See Figure 3 for an illustration.

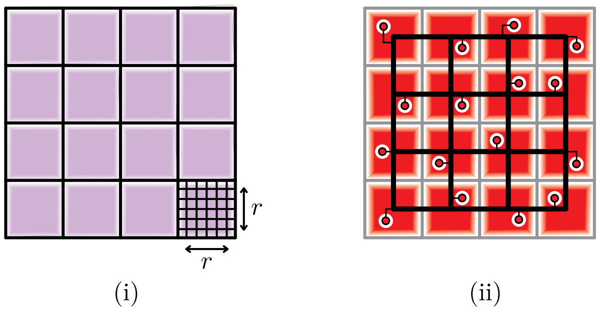

Figure 3 Diagrams of (i) a division of a $(5+4r)$-wall to either find an r-subwall without any red vertices, or a red (4×4)-grid minor and (ii) a (4×4)-grid minor in the second outcome.

Towards deducing the excluded apex-minor version of the local structure theorem, after the initial application of the Flat Wall Theorem, declare the neighbourhood of A_1 to be the set of red vertices and set up the numbers to obtain a red $(|A_1|k \times |A_1|k)$-grid as one possible outcome. In case this red grid is found, the pigeonhole principle yields a $(k\times k)$-grid rooted in the neighbourhood of a single member of A_1. Otherwise, W_2 must be flat without deleting a single apex vertex.

Step 2: Classifying a red society

The Society Classification Theorem – first proven explicitly by Kawarabayashi, Thomas, and Wollan [33] – may be seen as an extension of the Flat Wall Theorem as follows: Given a vortex surrounded by a large number of concentric cycles in the embedded part – called the *nest* – contained within a disk Δ of the surface, one may find one of four possible outcomes, see Figure 4 for an illustration:

1. A large clique-minor in G whose model is highly connected to the nest.
2. Subject to removing a small set of apices A, either
   a. a large flat crosscap transaction in Δ that traverses the interior of the vortex, and moreover is orthogonal[3] to most of the starting nest,
   b. a large flat handle transaction in Δ that traverses the interior of the vortex and moreover is orthogonal to most of the starting nest, or
   c. a weak near embedding of the part of the graph drawn in Δ, centred at the nest, with a bounded number of vortices each with a bounded adhesion path decomposition. Moreover, each of these vortices is surrounded by a large nest which is linked back to the starting nest via many disjoint paths which are orthogonal to the starting nest.

The proof of the red society classification theorem (see Theorem 4.2 in the full version) utilizes the society classification theorem above as a departure point and further refines each of its outcomes as outlined below.

Similarly to before, case (i), where a clique minor is given, is delegated to the result of [45].

[3] A *transaction* is a set of disjoint paths linking two disjoint boundary segments of Δ. We say that a transaction escaping the embedded part of our graph is *orthogonal* to a nest, if the intersection of each path in the transaction with each cycle in the nest consists of two paths – one caused by escaping the embedded part and the other by entering it again (see Figure 4 as a reference).

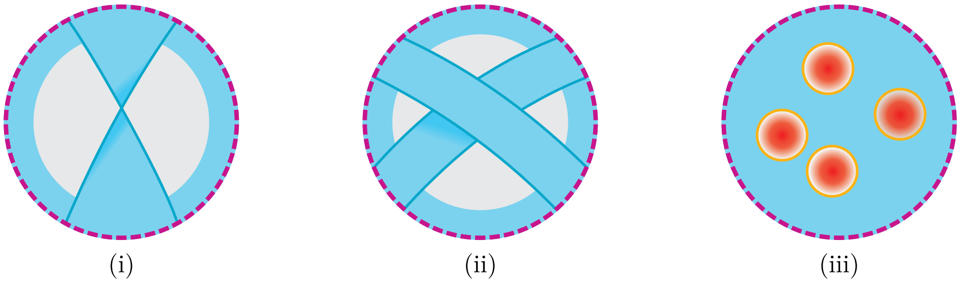

Figure 4 Diagrams of (i) a crosscap transaction, (ii) a handle transaction, and (iii) a weak near embedding in a disk with a small number of vortices.

To explain cases (ii.a) and (ii.b), we briefly recall the notion of a *flat transaction*. The goal of this outcome eventually is to augment the weak near embedding using structure extracted from the interior of the vortex. We say that a transaction is *monotone* if the order of its endpoints on one of the two segments of Δ its endpoints are placed on is mirrored by the order of its endpoints on the other segment. We define its strip of a monotone transaction as the union of those components of the graph drawn in Δ that remain after deleting the transaction and the boundary of Δ, and that have neighbours on an interior path of the transaction. A transaction is *flat (under A)* if its strip admits a weak near embedding without vortices (after deleting a set of vertices A).

As the given flat transaction is orthogonal to most of the nest, arguments analogous to those used in **Step 1** yield a flat subtransaction such that all regions between consecutive paths in the weak near embedding of the strip are homogeneous, meaning all contain a red vertex or none does. In the former case we obtain a large red grid; otherwise, the transaction is entirely free of red vertices. See Figure 5 for an illustration of the situation.

Outcome (ii.c) of the society classification theorem constitutes the main technical challenge. In existing proofs of the GMST (e.g. [48, 52, 33, 27]), refining a weak near embedding by repeatedly splitting vortices requires sacrificing part of the nest, and termination is ensured by using so-called *crooked transactions* [49]. For technical reasons, this tool is unavailable here. However, the existence of a weak near embedding as in (ii.c) allows us to avoid sacrificing cycles.

Using similar techniques as for the proof of the red flat wall theorem on the wall-like structure from (ii.c), we either find a large red grid or obtain a large nest together with many radially traversing paths orthogonal to it, such that every vortex and every red vertex is enclosed by the nest. We then view the interior of the nest as a single vortex. Either this vortex admits a path decomposition of bounded adhesion, or there exists a large transaction traversing the nest. A sequence of lemmas that follow shows that, such a transaction can always be chosen to traverse the nest interior, be orthogonal to the nest, and avoid all red vertices, similarly to cases (ii.a) and (ii.b). This relies on adapting a technique originating in [57] and further developed in [43, 44]. Crucially, this allows us to discard one side of a nest when splitting it, provided that side contains no red vertex or vortex, without reducing the nest size. It is noteworthy that this case never shows up in the actual proof of the Graph Minor Structure Theorem.

To ensure polynomial bounds, we must also control how new nests connect to previous ones. Naively splitting radial paths at each step would lead to exponential loss. Instead, we inductively maintain a *nest tree*, a hierarchical tree-like structure in which each nest

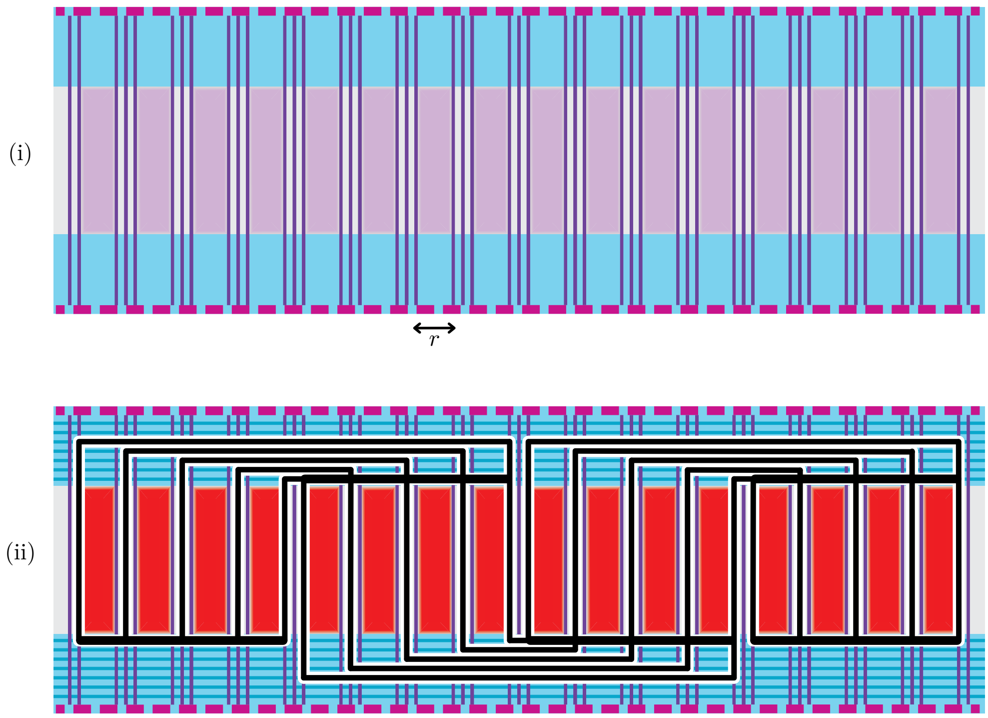

Figure 5 Diagrams of (i) a strip of a transaction divided into 16 pairwise disjoint substrips, each containing a transaction of order r and (ii) a (4×4)-grid minor in the case where all strips in (i) contain a red vertex.

is connected to its two children nests by sets of radial paths extracted from the splitting transactions, rather than from the parent nest's own radial paths. See Figure 6 for an illustration. This guarantees that no loss accumulates.

Finally, termination is ensured using the nest tree itself. If the number of leaves in the nest tree exceeds a certain quadratic threshold (in terms of the target grid size), a simple argument using Menger's Theorem and the rich structure of the nest tree yields a large red grid minor. Otherwise, we obtain only a bounded number of leaves for our nest tree, in total containing all original vortices and red vertices. Each of these leaves can now be seen as vortices admitting a path decomposition of bounded adhesion, as no large transaction traversing them exists. Again using Menger's Theorem and the structure of the nest tree we can link the nests in the leaves back to (part of) the original nest as desired.

Step 3: The local structure theorem

Using the red society classification theorem, we now inductively prove the red local structure theorem with respect to a large wall. The statement essentially reads as follows: Given an annotated graph (G, R) with a large wall W, one may find one of the four possible outcomes:

1. A small set of vertices S such that the component of $G - S$ that contains the majority of W is free of red vertices.
2. A large red clique-minor in G whose model is highly connected to W.

3. A large red grid-minor in G whose model is highly connected to W.
4. A weak near embedding of G centred at W with a bounded number of vortices each with a bounded adhesion path decomposition such that all red vertices are confined in the interior of the vortices.

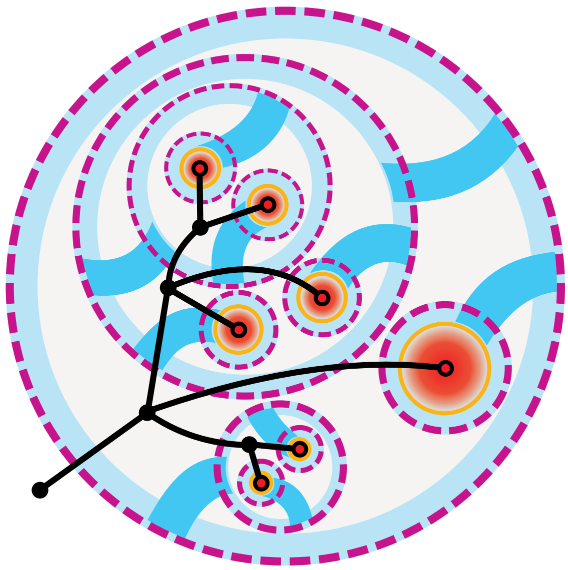

Figure 6 A diagram of the tree-like structure of nests and vortices.

The proof of the red local structure theorem proceeds by induction, with the red flat wall theorem as the base case. In the base case, the red flat wall theorem yields either a large red grid, in which case we are done, or a weak near-embedding with a single vortex that contains all red vertices in its interior. This embedding serves as the starting point of the refinement procedure.

Using the wall infrastructure provided by the flat wall theorem, we first construct the required nest and then iteratively apply the red society classification theorem to the current weak near-embedding. At each step of this process, if the outcome is a red clique or a red grid, we are done.

If instead a flat and homogeneous crosscap or handle transaction is produced, we proceed as follows. If the transaction is red, we again obtain a red grid. Otherwise, we apply tools from [33] to extend the working surface by attaching a crosscap or handle and routing most of the transaction through it. This yields a single new vortex cell for further refinement. The new vortex is disjoint from the added crosscap or handle, contains all red vertices in its interior, and contains a sufficiently large nest which is defined by combining a part of the crosscap or handle transaction with the original nest, allowing the induction to continue.

In the remaining case, the red society classification theorem produces a weak near-embedding in which the vortex is replaced by a bounded number of vortices, each admitting a path decomposition of bounded adhesion and together containing all red vertices in their interiors. By stitching this embedding along the boundary of the vortex to the already embedded part of the graph, we obtain the final outcome and conclude the proof.

Step 4: Local to global

Finally, we are able to use the local structure theorem of **Step 3** that in the absence of a large red grid, finds the desired tree-decomposition of Theorem 1.

As we have previously discussed this part follows in a straightforward manner, by employing a fairly standard technique originating from [50] that allows to turn the red local structure theorem into the desired global theorem based on tree-decompositions.

References

1 Brenda S. Baker. Approximation algorithms for NP-complete problems on planar graphs. *Journal of the ACM*, 41(1):153–180, January 1994. doi:10.1145/174644.174650.

2 Hans L. Bodlaender. Classes of graphs with bounded tree-width. Technical Report RUU-CS-86-22, Utrecht University, December 1986.

3 Marthe Bonamy, Nicolas Bousquet, Louis Esperet, Carla Groenland, Chun-Hung Liu, François Pirot, and Alexander Scott. Asymptotic dimension of minor-closed families and Assouad–Nagata dimension of surfaces. *Journal of the European Mathematical Society*, 26(10):3739–3791, May 2023. doi:10.4171/jems/1341.

4 Jannis Bulian and Anuj Dawar. Graph isomorphism parameterized by elimination distance to bounded degree. In *Parameterized and exact computation*, volume 8894 of *Lecture Notes in Comput. Sci.*, pages 135–146. Springer, Cham, 2014. doi:10.1007/978-3-319-13524-3_12.

5 Jannis Bulian and Anuj Dawar. Fixed-parameter tractable distances to sparse graph classes. *Algorithmica*, 79(1):139–158, 2017. doi:10.1007/s00453-016-0235-7.

6 Quentin Claus, Jędrzej Hodor, Gwenaël Joret, and Pat Morin. Excluding an apex-forest or a fan as quickly as possible. *arXiv preprint arXiv:2602.03833*, 2026. doi:10.48550/arXiv.2602.03833.

7 Bruno Courcelle. The monadic second-order logic of graphs. I. Recognizable sets of finite graphs. *Information and Computation*, 85(1):12–75, March 1990. doi:10.1016/0890-5401(90)90043-H.

8 Radu Curticapean. Counting Matchings with k Unmatched Vertices in Planar Graphs. In *24th Annual European Symposium on Algorithms (ESA 2016)*, volume 57 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 33:1–33:17. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2016. doi:10.4230/LIPIcs.ESA.2016.33.

9 Marek Cygan, Fedor V Fomin, Łukasz Kowalik, Daniel Lokshtanov, Dániel Marx, Marcin Pilipczuk, Michał Pilipczuk, and Saket Saurabh. *Parameterized Algorithms*. Springer International Publishing : Imprint: Springer, Cham, 1st ed. 2015 edition, 2015. doi:10.1007/978-3-319-21275-3.

10 Clément Dallard, Martin Milanič, Andrea Munaro, and Shizhou Yang. Layered tree-independence number and clique-based separators. *arXiv preprint*, 2025. doi:10.48550/arXiv.2506.12424.

11 Erik D. Demaine and MohammadTaghi Hajiaghayi. Equivalence of local treewidth and linear local treewidth and its algorithmic applications. In *Proceedings of the Fifteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 840–849. ACM, New York, 2004.

12 Erik D. Demaine, MohammadTaghi Hajiaghayi, and Ken-ichi Kawarabayashi. Approximation algorithms via structural results for apex-minor-free graphs. In *Automata, languages and programming. Part I*, volume 5555 of *Lecture Notes in Comput. Sci.*, pages 316–327. Springer, Berlin, 2009. doi:10.1007/978-3-642-02927-1_27.

13 Feodor F. Dragan, Fedor V. Fomin, and Petr A. Golovach. A PTAS for the sparsest spanners problem on apex-minor-free graphs. In *Mathematical Foundations of Computer Science 2008*, volume 5162 of *Lecture Notes in Comput. Sci.*, pages 290–298. Springer, Berlin, 2008. doi:10.1007/978-3-540-85238-4_23.

14 S. E. Dreyfus and R. A. Wagner. The Steiner problem in graphs. *Networks*, 1:195–207, 1971/72. doi:10.1002/net.3230010302.

15 Vida Dujmović, David Eppstein, Gwenaël Joret, Pat Morin, and David R. Wood. Minor-Closed Graph Classes with Bounded Layered Pathwidth. *SIAM Journal on Discrete Mathematics*, 34(3):1693–1709, January 2020. `doi:10.1137/18M122162X`.

16 Vida Dujmović, Louis Esperet, Pat Morin, Bartosz Walczak, and David R. Wood. Clustered 3-colouring graphs of bounded degree. *Combin. Probab. Comput.*, 31(1):123–135, 2022. `doi:10.1017/s0963548321000213`.

17 Vida Dujmović and Fabrizio Frati. Stack and queue layouts via layered separators. *J. Graph Algorithms Appl.*, 22(1):89–99, 2018. `doi:10.7155/jgaa.00454`.

18 Vida Dujmović, Pat Morin, and David R. Wood. Layered separators in minor-closed graph classes with applications. *J. Combin. Theory Ser. B*, 127:111–147, 2017. `doi:10.1016/j.jctb.2017.05.006`.

19 Vida Dujmović, Pat Morin, and Céline Yelle. Two results on layered pathwidth and linear layouts. *J. Graph Algorithms Appl.*, 25(1):43–57, 2021. `doi:10.7155/jgaa.00549`.

20 Zdeněk Dvořák and Robin Thomas. List-coloring apex-minor-free graphs, December 2016. `arXiv:1401.1399`.

21 D. Eppstein. Diameter and Treewidth in Minor-Closed Graph Families. *Algorithmica*, 27(3):275–291, June 2000. `doi:10.1007/s004530010020`.

22 David Eppstein. Subgraph isomorphism in planar graphs and related problems. *J. Graph Algorithms Appl.*, 3:no. 3, 27, 1999. `doi:10.7155/jgaa.00014`.

23 Ranel E. Erickson, Clyde L. Monma, and Arthur F. Veinott, Jr. Send-and-split method for minimum-concave-cost network flows. *Math. Oper. Res.*, 12(4):634–664, 1987. `doi:10.1287/moor.12.4.634`.

24 Fedor V. Fomin, Daniel Lokshtanov, Dániel Marx, Marcin Pilipczuk, Michał Pilipczuk, and Saket Saurabh. Subexponential parameterized algorithms for planar and apex-minor-free graphs via low treewidth pattern covering. In *57th Annual IEEE Symposium on Foundations of Computer Science—FOCS 2016*, pages 515–524. IEEE Computer Soc., Los Alamitos, CA, 2016. `doi:10.1109/FOCS.2016.62`.

25 Greg N. Frederickson. Planar graph decomposition and all pairs shortest paths. *J. Assoc. Comput. Mach.*, 38(1):162–204, 1991. `doi:10.1145/102782.102788`.

26 M. R. Garey and D. S. Johnson. The rectilinear Steiner tree problem is NP-complete. *SIAM J. Appl. Math.*, 32(4):826–834, 1977. `doi:10.1137/0132071`.

27 Maximilian Gorsky, Michał T. Seweryn, and Sebastian Wiederrecht. Polynomial Bounds for the Graph Minor Structure Theorem, April 2025. `doi:10.48550/arXiv.2504.02532`.

28 Maximilian Gorsky, Michał T. Seweryn, and Sebastian Wiederrecht. The price of homogeneity is polynomial, February 2026. `doi:10.48550/arXiv.2602.01882`.

29 Carla Groenland, Jesper Nederlof, and Tomohiro Koana. A polynomial time algorithm for Steiner tree when terminals avoid a rooted K_4-minor. In *19th International Symposium on Parameterized and Exact Computation*, volume 321 of *LIPIcs. Leibniz Int. Proc. Inform.*, pages Art. No. 12, 17. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024. `doi:10.4230/lipics.ipec.2024.12`.

30 Martin Grohe. Local Tree-Width, Excluded Minors, and Approximation Algorithms. *Combinatorica*, 23(4):613–632, December 2003. `doi:10.1007/s00493-003-0037-9`.

31 Jędrzej Hodor, Hoang La, Piotr Micek, and Clément Rambaud. Quickly excluding an apex-forest, April 2025. `doi:10.48550/arXiv.2404.17306`.

32 Bart M. P. Jansen and Céline M. F. Swennenhuis. Steiner tree parameterized by multiway cut and even less. In *32nd annual European Symposium on Algorithms*, volume 308 of *LIPIcs. Leibniz Int. Proc. Inform.*, pages Art. No. 76, 16. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024. `doi:10.4230/lipics.esa.2024.76`.

33 Ken-ichi Kawarabayashi, Robin Thomas, and Paul Wollan. Quickly excluding a non-planar graph, January 2021. `arXiv:2010.12397`.

34 Sándor Kisfaludi-Bak, Jesper Nederlof, and Erik Jan van Leeuwen. Nearly ETH-tight algorithms for planar Steiner tree with terminals on few faces. *ACM Trans. Algorithms*, 16(3):Art. 28, 30, 2020. doi:10.1145/3371389.

35 Tuukka Korhonen. A Single-Exponential Time 2-Approximation Algorithm for Treewidth. *SIAM Journal on Computing*, pages FOCS21–174–FOCS21–194, November 2023. doi:10.1137/22M147551X.

36 Tuukka Korhonen, Wojciech Nadara, Michał Pilipczuk, and Marek Sokołowski. Fully dynamic approximation schemes on planar and apex-minor-free graphs. In *Proceedings of the 2024 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 296–313, Philadelphia, PA, January 2024. Society for Industrial and Applied Mathematics. doi:10.1137/1.9781611977912.12.

37 Robert Krauthgamer, James R. Lee, and Havana Rika. Flow-cut gaps and face covers in planar graphs. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 525–534. SIAM, Philadelphia, PA, 2019. doi:10.1137/1.9781611975482.33.

38 Chun-Hung Liu and David R. Wood. Clustered coloring of graphs with bounded layered treewidth and bounded degree. *European J. Combin.*, 122:Paper No. 103730, 7, 2024. doi:10.1016/j.ejc.2023.103730.

39 Daniel Marx, Paul Seymour, and Paul Wollan. Rooted grid minors. *Journal of Combinatorial Theory, Series B*, 122:428–437, January 2017. doi:10.1016/j.jctb.2016.07.003.

40 Frédéric Mazoit. A single exponential bound for the redundant vertex Theorem on surfaces, September 2013. doi:10.48550/arXiv.1309.7820.

41 Laure Morelle, Evangelos Protopapas, Dimitrios M Thilikos, and Sebastian Wiederrecht. Excluding pinched spheres. *arXiv preprint arXiv:2506.14421*, 2025. URL: https://arxiv.org/abs/2506.14421.

42 Sukanya Pandey and Erik Jan van Leeuwen. Planar multiway cut with terminals on few faces. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2032–2062. [Society for Industrial and Applied Mathematics (SIAM)], Philadelphia, PA, 2022. doi:10.1137/1.9781611977073.81.

43 Christophe Paul, Evangelos Protopapas, Dimitrios M. Thilikos, and Sebastian Wiederrecht. Obstructions to Erdös-Pósa Dualities for Minors. In *2024 IEEE 65th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 31–52, Chicago, IL, USA, October 2024. IEEE. doi:10.1109/FOCS61266.2024.00013.

44 Christophe Paul, Evangelos Protopapas, Dimitrios M. Thilikos, and Sebastian Wiederrecht. The local structure theorem for graph minors with finite index, 2025. doi:10.48550/arXiv.2507.02769.

45 Evangelos Protopapas, Dimitrios M. Thilikos, and Sebastian Wiederrecht. Colorful Minors, July 2025. doi:10.48550/arXiv.2507.10467.

46 Neil Robertson and Paul D Seymour. Graph minors. II. Algorithmic aspects of tree-width. *Journal of Algorithms*, 7(3):309–322, September 1986. doi:10.1016/0196-6774(86)90023-4.

47 Neil Robertson and Paul D Seymour. Graph minors. V. Excluding a planar graph. *Journal of Combinatorial Theory, Series B*, 41(1):92–114, August 1986. doi:10.1016/0095-8956(86)90030-4.

48 Neil Robertson and Paul D Seymour. Graph minors. IX. Disjoint crossed paths. *Journal of Combinatorial Theory, Series B*, 49(1):40–77, June 1990. doi:10.1016/0095-8956(90)90063-6.

49 Neil Robertson and Paul D Seymour. Graph minors. VIII. A Kuratowski theorem for general surfaces. *Journal of Combinatorial Theory, Series B*, 48(2):255–288, April 1990. doi:10.1016/0095-8956(90)90121-F.

50 Neil Robertson and Paul D Seymour. Graph Minors. X. Obstructions to Tree-Decomposition. *Journal of Combinatorial Theory, Series B*, 52(2):153–190, July 1991. doi:10.1016/0095-8956(91)90061-N.

51 Neil Robertson and Paul D Seymour. Graph Minors. XIII. The Disjoint Paths Problem. *Journal of Combinatorial Theory, Series B*, 63(1):65–110, January 1995. doi:10.1006/jctb.1995.1006.

52 Neil Robertson and Paul D Seymour. Graph Minors. XVI. Excluding a non-planar graph. *Journal of Combinatorial Theory, Series B*, 89(1):43–76, September 2003. doi:10.1016/S0095-8956(03)00042-X.

53 Ignasi Sau, Giannos Stamoulis, and Dimitrios M. Thilikos. Parameterizing the quantification of CMSO: model checking on minor-closed graph classes. In *Proceedings of the 2025 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 3728–3742. SIAM, Philadelphia, PA, 2025. doi:10.1137/1.9781611978322.124.

54 Alex Scott and David R. Wood. Better bounds for poset dimension and boxicity. *Trans. Amer. Math. Soc.*, 373(3):2157–2172, 2020. doi:10.1090/tran/7962.

55 Sebastian Siebertz and Alexandre Vigny. Advances in algorithmic meta theorems. In *44th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science*, volume 323 of *LIPIcs. Leibniz Int. Proc. Inform.*, pages Art. No. 2, 29. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024. doi:10.4230/lipics.fsttcs.2024.2.

56 Dimitrios M. Thilikos and Sebastian Wiederrecht. The Graph Minor Structure Theorem through Bidimensionality, February 2024. arXiv:2306.01724.

57 Dimitrios M. Thilikos and Sebastian Wiederrecht. Killing a Vortex. *Journal of the ACM*, 71(4):1–56, August 2024. doi:10.1145/3664648.

58 Hilde Verbeek. Disjoint paths and directed Steiner tree on planar graphs with terminals on few faces. Master's thesis, Utrecht University, 2022.