Treewidth is Polynomial in Maximum Degree on Graphs Excluding a Planar Induced Minor

Édouard Bonnet
Univ Lyon, CNRS, ENS de Lyon, Université Claude Bernard Lyon 1, LIP UMR5668, France

Jędrzej Hodor
Theoretical Computer Science Department, Faculty of Mathematics and Computer Science and Doctoral School of Exact and Natural Sciences, Jagiellonian University, Kraków, Poland

Tuukka Korhonen
Department of Informatics, University of Bergen, Norway

Tomáš Masařík
Institute of Informatics, Faculty of Mathematics, Informatics and Mechanics, University of Warsaw, Warszawa, Poland

Abstract

A graph G contains a graph H as an induced minor if H can be obtained from G by vertex deletions and edge contractions. We show that for every k-vertex planar graph H, every graph G excluding H as an induced minor has treewidth at most $\Delta(G)^{2^{O(k)}}$ where $\Delta(G)$ denotes the maximum degree of G. Previously, Korhonen [JCTB '23] has shown the upper bound of $k^{O(1)}2^{\Delta(G)^5}$ whose dependence in $\Delta(G)$ is exponential. More precisely, we show that every graph G excluding as induced minors a k-vertex planar graph and a q-vertex graph has treewidth at most $k^{O(1)} \cdot \Delta(G)^{f(q)}$ with $f(q) = 2^{O(q)}$. A direct consequence of our result is that for every hereditary graph class $\mathcal{C}$, if graphs of $\mathcal{C}$ have treewidth bounded by a function of their maximum degree, then they in fact have treewidth polynomial in their maximum degree.

Funding *Édouard Bonnet*: was supported by the ANR projects TWIN-WIDTH (ANR-21-CE48-0014) and Digraphs (ANR-19-CE48-0013).
Jędrzej Hodor: was partially supported by a Polish Ministry of Education and Science grant (Perły Nauki; PN/01/0265/2022).
Tuukka Korhonen: was supported by the Research Council of Norway via the project BWCA (grant no. 314528).
Tomáš Masařík: was supported by Polish National Science Centre SONATA-17 grant number 2021/43/D/ST6/03312.

Acknowledgements We thank Eunjung Kim and Jana Masaříková for discussions at an earlier stage of the project. This work was done at the 2023 Structural Graph Theory workshop held at the IMPAN Będlewo Conference Center (Poland) in September. This workshop was a part of STRUG: Stuctural Graph Theory Bootcamp, funded by the "Excellence initiative – research university (2020-2026)" of University of Warsaw. We are grateful to both the organizers and the participants of the workshop for creating an enjoyable work environment.

1 Introduction

A graph G contains a graph H as a *minor* if H can be obtained from G by vertex deletions, edge deletions, and edge contractions. The notion of *induced minor* is defined similarly except edge deletions are disallowed. The celebrated Grid Minor theorem [25, 26] implies that graphs without large grid minors have low treewidth. What can be said about the treewidth of graphs solely excluding grids as *induced* minor? Their treewidth can be arbitrarily large, as exemplified by cliques. However, a notable result by Korhonen is that their treewidth can be upperbounded by a function of their maximum degree $\Delta(\cdot)$.

▶ **Theorem 1** ([18]). *Every graph G excluding a fixed k-vertex planar graph as an induced minor has treewidth at most $k^\gamma 2^{\Delta(G)^5}$ for some universal constant γ.*

In this paper, we show that the exponential dependence in $\Delta(G)$ is not necessary.

▶ **Theorem 2.** *Every graph G excluding a fixed k-vertex planar graph as an induced minor has treewidth at most $\Delta(G)^{f(k)}$ with $f(k) = 2^{O(k)}$.*

The next step is to confine the dependence in k (at present, in the exponent of $\Delta(G)$) to a mere multiplicative factor and show a treewidth upper bound of $g(k)\Delta(G)^{O(1)}$ for some function g. We in fact prove a stronger statement than Theorem 2, resolving the next step for graphs excluding a (non-planar) graph as an induced minor.

▶ **Theorem 3.** *Every graph G excluding as induced minors a k-vertex planar graph and a q-vertex graph has treewidth at most $k^{O(1)} \cdot \Delta(G)^{f(q)}$ with $f(q) = 2^{O(q)}$.*

For instance, Theorem 3 implies that the string graphs that exclude a fixed k-vertex planar graph have treewidth $(k\Delta(G))^{O(1)}$. Indeed, the class of *all* string graphs exclude the 1-subdivision of K_5 (a 15-vertex graph) as an induced minor.

Our tools combine well with classes of graphs that admit a product structure. More precisely, we prove the following.

▶ **Theorem 4.** *Let H be a graph of treewidth at most t, and P be a path. Let G be a subgraph of $H \boxtimes P$ excluding a k-vertex planar graph as an induced minor. Then the treewidth of G is at most $k^{O(1)} \cdot t^{O(1)} \cdot \Delta(G)^{O(1)}$.*

Note that there are graph classes that admit a product structure of the form as in Theorem 4, and yet possess every graph as induced minor; see [11]. Thus Theorem 4 is not a special case of Theorem 3. For instance, Theorem 4 implies that ℓ-planar graphs[1] that exclude a fixed k-vertex planar graph as an induced minor have treewidth $(k\ell\Delta(G))^{O(1)}$.

A dependence in $\Delta(G)$ is necessary. There are subgraphs of the strong product of a path with a star (hence a graph H of treewidth 1) avoiding a planar induced minor, but whose treewidth is a growing function of the number of vertices. Take the $n \times n$ grid, remove the "vertical" edges, and add in each "column" a vertex adjacent to every vertex in the column; see Figure 1. This construction found by Pohoata [23], and rediscovered by Davies [8], has treewidth $\Theta(n)$ but avoids the 5×5 grid as an induced minor. The figure is a proof-by-picture that these graphs are indeed subgraphs of strong products of a path and a star.

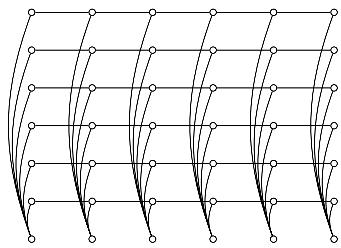

Figure 1 The Pohoata–Davies 6×6 grid.

[1] those graphs that can be drawn in the plane such that every edge is intersected by at most ℓ other edges

Chudnovsky [4, Open problem 4.1] asks if, when $\Delta(G) = O(\log |V(G)|)$, the treewidth of graphs excluding the $k \times k$ grid as an induced minor and the biclique $K_{t,t}$ as a subgraph is $O_{t,k}(\log |V(G)|)$. Our results partially answer this question without requiring the absence of $K_{t,t}$ subgraph: The treewidth of these graphs is at most polylogarithmic.

At first sight, Chudnovsky's question centered around forbidden induced subgraphs may look somewhat different from the setting of Theorem 3. The two statements match since forbidding large cliques and bicliques as induced subgraphs is, by Ramsey's theorem [24], equivalent to excluding large bicliques as subgraphs, and forbidding a subdivision of a large wall or the line graph of a subdivision of a large wall as an induced subgraph is the same as excluding a large grid as an induced minor. Another simplifying feature of working with induced minors rather than induced subgraphs is that excluding as induced minor a large grid, or a large wall, or a planar graph of large treewidth are all equivalent.

The motivation behind the $\Delta(G) = O(\log |V(G)|)$ condition in Chudnovsky's question is that the treewidth could in principle be logarithmic in $|V(G)|$ as well. This would yield polynomial-time algorithms for several problems including Max Independent Set. We come slightly short of proving it, but Theorem 2 implies a quasipolynomial-time algorithm for Max Independent Set on graphs of logarithmic degree and excluding a fixed grid as an induced minor. Let us insist that we do not need to assume the absence of some biclique subgraph.

It is possible (and believed) that graphs G excluding a k-vertex planar graph as an induced minor have treewidth $g(k)\Delta(G)$, for some function g. This also is motivated by fast algorithms for Max Independent Set, as it would imply a subexponential-time algorithm running in $2^{\tilde{O}_k(\sqrt{|V(G)|})}$. Dallard, Milanič, and Štorgel [5] even ask whether a (quasi)polynomial-time algorithm always exists in the absence of a fixed planar induced minor. After Korhonen [18] gave the first (very slightly) subexponential algorithm, Korhonen and Lokshtanov [19] provided an algorithm running in time $2^{\tilde{O}_k(|V(G)|^{2/3})}$, which extends to the case when the forbidden induced minor is non-planar. There have been several recent developments in (quasi)polynomial algorithms for Max Independent Set on graphs excluding a planar induced minor [1, 2, 6, 7, 13, 15, 16, 22], some phrased in terms of forbidden induced subgraphs instead.

Let us conclude by explicitly mentioning the potential next improvements to Theorem 2 by increasing difficulty.

▶ **Question 1.** *Does every graph G excluding a fixed k-vertex planar graph as an induced minor have, for some function f, treewidth at most $f(k)\Delta^{k^{O(1)}}$? treewidth at most $f(k)\Delta^{O(1)}$? treewidth at most $f(k)\Delta$?*

We note that Gartland and Lokshtanov [14] conjecture the following, which would in particular imply a positive answer to every case of the above question.

▶ **Conjecture 5** (Gartland–Lokshtanov)**.** *There is a function $f : \mathbb{N} \to \mathbb{N}$ such that every graph excluding a fixed k-vertex planar graph as an induced minor has a balanced separator dominated by at most $f(k)$ vertices.*

Fully spelled out, the conjecture says that for every G excluding a k-vertex planar graph as an induced minor, there is a set $D \subseteq V(G)$ of size at most $f(k)$ such that $G - N[D]$ has no connected component of size larger than $|V(G)|/2$. In particular, these graphs would have balanced separators of size $f(k)(\Delta(G) + 1)$, known to imply treewidth $O(f(k)\Delta(G))$ [12]. If true, by a simple win-win argument, Max Independent Set could be solved in time $2^{\tilde{O}_k(\sqrt{n})}$ on n-vertex graphs excluding a k-vertex planar graph as an induced minor.

Outline of the proof of Theorem 3. Let G be any graph excluding a k-vertex planar graph and a q-vertex graph (possibly the same graph, and $k = q$) as induced minors. We want to upperbound the treewidth of G, denoted by $\mathrm{tw}(G)$, by a polynomial function of its maximum degree $\Delta(G)$. Our general plan is to find, for some functions $f, g : \mathbb{N} \to \mathbb{N}$, an induced subgraph G' of G retaining treewidth at least $\mathrm{tw}(G)^\beta / \Delta(G)^{f(q)}$, for some constant[2] $\beta > 0$, with maximum degree at most $g(q)$. If we get such an induced subgraph G', we can advantageously apply Theorem 1 to G', which, like its induced supergraph G, excludes a k-vertex planar graph as an induced minor. We thus get that the treewidth of G' is at most $k^{O(1)} 2^{g(q)^5}$. Importantly we avoided here the dependence in $\Delta(G)$ (in the exponent). This translates to an upper bound of $k^{O(1)} 2^{g(q)^5/\beta} \Delta(G)^{f(q)/\beta}$ for the treewidth of G. Our main theorem is then a consequence of the existence of G' for f, g both having a single-exponential growth.

Let us then describe how we find G'. The first ingredient is a classical result by Klein–Plotkin–Rao [17]. Originally and in the context of approximation algorithms, the result is (informally) expressed as follows. For every graph excluding a biclique minor (but more generally, excluding any fixed minor), there is an arbitrarily small fraction of its edges whose removal gives rise to connected components of bounded weak diameter. This means that in every connected component of the resulting graph, every pair of vertices are close together when measured in the original graph. Following Lee [21], who observed that the result extends to graphs merely excluding a 1-subdivided biclique as an *induced* minor (but more generally, any induced minor), we reprove this fact in the language of so-called *clustered edge-colorings*: Every graph G excluding a q-vertex graph as an induced minor can be 2^{q+1}-edge-colored such that every monochromatic connected component has weak diameter $q^{O(1)}$, and hence size at most $\Delta(G)^{q^{O(1)}}$. The latter quantity is called the *clustering* of this edge-coloring.

The second ingredient is to utilize a contraction–uncontraction technique of Korhonen [18] (albeit in an abstract and more general framework) on each color class of the clustered edge-coloring. This goes as follows. By contracting each connected component of one color class F to a single vertex, the treewidth is divided by at most $\Delta(G)^{q^{O(1)}}$ (the maximum size of such connected components). Hence, by a celebrated result of Chekuri and Chuzhoy [3], the contracted graph admits a subgraph H with $\mathrm{tw}(H) \geqslant \mathrm{tw}(G)/(\Delta(G)^{q^{O(1)}} \mathrm{polylog}\ \mathrm{tw}(G))$ and maximum degree at most 3. Now reverting the contraction and going back to G, it is easy to build an induced subgraph of G with at most three edges of F incident to every vertex that can recreate H as a minor, and hence has at least its treewidth. After the 2^{q+1} successive iterations of the contraction–uncontraction technique, the extracted induced subgraph has maximum degree bounded by $3 \cdot 2^{q+1}$ and treewidth at least $\mathrm{tw}(G)^{.99}/\Delta(G)^{2^{O(q)}}$, as required in the first paragraph of the outline, which contains the end of the argument.

Organization of the paper. In Section 2 we recall the relevant graph-theoretic definitions and notations. In Section 3 we abstract out one iteration of the contraction–uncontraction technique, and in Section 4 we link it to clustered edge-colorings and repeatedly apply it. In Section 5 we reprove Klein–Plotkin–Rao's result when excluding an *induced* minor (rather than a minor). In Section 6 we wrap everything up and prove Theorem 3, our main result. In Section 7 we reduce the search for clustered edge-colorings to edge-colorings all color classes of which have bounded treewidth. We use that, in Section 8, to bound the treewidth of classes excluding a planar induced minor and admitting a product structure, via the tools developed in the paper, thereby establishing Theorem 4.

---

[2] Actually our proof gets β arbitrarily close to 1.

2 Preliminaries

If $i \leqslant j$ are two integers, we denote by $[i, j]$ the set of integers $\{i, i+1, \dots, j-1, j\}$, and by $[i]$, the set $[1, i]$. All our graphs are finite and simple (i.e., do not have parallel edges or self-loops). We denote by $V(G)$ and $E(G)$ the set of vertices and edges of a graph G, respectively, so $G = (V(G), E(G))$. For $S \subseteq V(G)$, the *subgraph of G induced by S*, denoted $G[S]$, is obtained by removing from G all the vertices that are not in S (together with their incident edges). We say that S is *connected* if $G[S]$ has a single connected component. Then $G - S$ is a short-hand for $G[V(G) \setminus S]$. We denote by $N_G(v)$ and $N_G[v]$, the open, respectively closed, neighborhood of v in G. For each $S \subseteq V(G)$, we set $N_G(S) := (\bigcup_{v \in S} N_G(v)) \setminus S$ and $N_G[S] := N_G(S) \cup S$. We denote by $\Delta(G)$ the maximum degree of a graph G, and by $\mathrm{tw}(G)$, its treewidth.

We say that a graph H is an *induced minor* of a graph G if H can be obtained from G after vertex deletions and edge contractions, where the contraction of edge uv unifies u, v in a single vertex adjacent to every vertex that is a neighbor of u or of v. We insist that our edge contractions keep the graph simple: we ignore self-loops, and retain at most one edge between a fixed pair of vertices. A *minor* is the same except edge deletions are further allowed. An *induced minor model* (resp. *minor model*) of H in G is a collection $\{B_v \ : \ v \in V(H)\}$ of pairwise vertex-disjoint connected subsets of $V(G)$, often called *branch sets*, such that for every $v, w \in V(H)$, $B_v \cup B_w$ is connected in G if and only if $vw \in E(H)$ (resp. for every $vw \in E(H)$, $B_v \cup B_w$ is connected). Then, it can be observed that H is a(n induced) minor of G if and only if G admits a(n induced) minor model of H.

For a positive integer s and a graph G, the *s-subdivision of G* is the graph G with each edge subdivided s times, denoted by s-subd(G). The *diameter of G* is defined as $\max_{u,v \in V(G)} d_G(u, v)$, where $d_G(u, v)$ is the number of edges in a shortest path between u and v. The *weak diameter of S in G* for $S \subseteq V(G)$ is defined as $\max_{u,v \in S} d_G(u, v)$.

If $\mathcal{P} = \{P_1, \dots, P_h\}$ is a partition of $V(G)$ for some graph G, then by $G/\mathcal{P}$, we denote the graph with vertex set $\mathcal{P}$ such that $P, P' \in \mathcal{P}$ are connected by an edge whenever there is an edge in G connecting a vertex in P with a vertex in P'. The *BFS layering of a connected graph G from a vertex $v \in V(G)$* is the partition $\mathcal{L} = \{L_0, L_1, \dots, L_s\}$ of G, where $L_i = \{u \in V(G) : d_G(v, u) = i\}$, called the *$i$-th layer of $\mathcal{L}$*, for each $i \in [0, s]$, and s is the maximum integer such that there exists $u \in V(G)$ with $d_G(v, u) = s$. In particular $L_0 = \{v\}$. A *BFS layering of a graph G* is the BFS layering of G from some vertex $v \in V(G)$. Sometimes we abbreviate it simply to *BFS of G*.

Let $\mathcal{L} = \{L_0, L_1, \dots, L_s\}$ be a BFS of G from some $v \in V(G)$. For each vertex $u \in L_i$ with $i \in [s]$, we pick an arbitrary neighbor of u in L_{i-1} and call it the parent of u. Now, let F be the set of edges consisting of all vertex–parent relations. The graph $(V(G), F)$ is a tree. We consider it to be rooted in v, and we call it a *BFS tree of $\mathcal{L}$*. A path is *vertical* in a rooted tree if it connects a vertex with one of its ancestors or descendants.

3 Contraction–uncontraction technique

We will need a *treewidth sparsifier*, i.e., the extraction of a subcubic subgraph of large treewidth in a graph of larger treewidth. We could here use the Grid Minor theorem [26], but the following result of Chekuri and Chuzhoy provides a better lower bound in the resulting treewidth.

▶ **Theorem 6** ([3]). *There is a constant $\delta > 0$ such that every graph of treewidth k admits a subcubic subgraph of treewidth at least $k/\log^{\delta} k$.*

The next lemma abstracts out the contraction–uncontraction technique of the third author which, in [18], is specifically used over radius-2 balls.

▶ **Lemma 7.** *Let p be a positive integer, G be a graph, and $F \subseteq E(G)$ be such that every connected component of the graph $(V(G), F)$ has at most p vertices. Then, G admits an induced subgraph G' such that*
- *in G', every vertex is incident to at most three edges of $F \cap E(G')$, and*
- *$tw(G') \geqslant tw(G)/(p \log^\delta tw(G))$, with δ the constant of Theorem 6.*

Proof. Let $\mathcal{P}$ be the partition $\{P_1, \dots, P_h\}$ of $V(G)$ into the vertex sets of the connected components of $(V(G), F)$. It follows that $|P_i| \leqslant p$ for every $i \in [h]$. In particular, $\text{tw}(G/\mathcal{P}) \geqslant \text{tw}(G)/p$. Indeed, a tree-decomposition of $G/\mathcal{P}$ of width at most $\text{tw}(G)/p - 1$ could be turned into a tree-decomposition of G of width at most $\text{tw}(G)/p \cdot \max_{i \in [h]} |P_i| - 1 \leqslant \text{tw}(G) - 1$, simply by flattening the parts of $\mathcal{P}$ in each bag, leading to a contradiction. On the other hand, $\text{tw}(G/\mathcal{P}) \leqslant \text{tw}(G)$ since $G/\mathcal{P}$ is obtained from G by performing edge contractions, as each P_i is connected.

By Theorem 6 applied to $G/\mathcal{P}$, there is a subcubic subgraph H of $G/\mathcal{P}$ with

$$\text{tw}(H) \geqslant \frac{\text{tw}(G/\mathcal{P})}{\log^\delta \text{tw}(G/\mathcal{P})} \geqslant \frac{\text{tw}(G)}{p \log^\delta \text{tw}(G)}.$$

We now build an induced subgraph G' of G having H as a minor (hence at least its treewidth) such that every vertex of G' is incident to at most three edges of F. As H is subcubic, each $P \in V(H)$ is incident to at most three edges of H. From each $P \in V(H)$, let us keep a minimal subset $P' \subseteq P$ such that $G[P']$ is connected and $G[\bigcup_{P \in V(H)} P']/\mathcal{P}'$ still contains H as a subgraph, where $\mathcal{P}' := \{P' : P \in V(H)\}$. By *minimal* we mean that for each $P' \in V(H)$, the removal of any vertex in P' breaks one of the latter conditions.

Note that each $P' \in V(H)$ comprises up to three *terminals* realizing the up-to-three edges in H, plus a minimal subset connecting these three terminals in P. Therefore, if P' would contain a vertex v with more than three neighbors in P', we could delete one of its neighbors by taking shortest paths from v to the terminals in $G[P']$ and deleting a neighbor not used in these shortest paths. This implies that every vertex of P' is incident to at most three edges of F in $G[P']$, since no edge of F can have exactly one endpoint in P.

Thus we set $G' := G[\bigcup_{P' \in V(H)} P']$, and get $\text{tw}(G') \geqslant \text{tw}(H) \geqslant \text{tw}(G)/(p \log^\delta \text{tw}(G))$. ◀

4 Clustered edge-coloring

In light of Lemma 7 and our plan to find an induced subgraph whose maximum degree is low but whose treewidth is still high, we wish to edge-partition our graph into a bounded number of graphs with no large connected component. This is referred to as *clustered edge-coloring*. We note that the vertex kind, *clustered coloring* has been more thoroughly explored. An edge-coloring has *clustering p* if every monochromatic connected component has at most p vertices. For instance, an edge-coloring with clustering 2 is a proper edge-coloring.

▶ **Lemma 8.** *Let p be a positive integer, G be a graph, and $F_1, \dots, F_h$ be the color classes of an edge-coloring of G with clustering p. Then, G admits an induced subgraph G' such that*
- *$\Delta(G') \leqslant 3h$, and*
- *$tw(G') \geqslant tw(G)/(p \log^\delta tw(G))^h$, with δ the constant of Theorem 6.*

Proof. Set $G_0 := G$. For every $i \in [h]$ going from 1 to h, let G_i be the induced subgraph of G_{i-1} obtained by applying Lemma 7 with edge subset $F := F_i \cap E(G_{i-1})$. We then define

G' as G_h. By the first item of Lemma 7, every vertex of G' has at most three incident edges in $F_i \cap E(G')$, hence has degree at most $3h$. The second item readily follows from that of Lemma 7. ◀

5 Klein–Plotkin–Rao iterated breadth-first searches

We will now use a classical result of Klein, Plotkin, and Rao [17], or rather its proof. The result and its follow-ups are phrased in terms of fairly small edge cuts disconnecting pieces of small weak diameter in proper minor-closed classes. A first relevant observation, made by Lee [21], is that the result actually holds in the more general context of classes excluding an induced minor. We will extract a somewhat different-looking statement than the original Klein–Plotkin–Rao formulation.

For a fixed integer h (which we will set to 2), consider the following recursive process applied to a graph G. For each connected component H of G, pick an arbitrary vertex $v \in V(H)$ and compute the BFS layering $\{L_0, L_1, \dots, L_s\}$ of H from v. For each $i \in [0, s-h+1]$, recurse in $G[\bigcup_{i \leqslant j < i+h} L_j]$. We call this process *iterated BFSes of width h to property $\mathcal{Q}$* if we stop each recursion branch when the induced subgraph considered in this branch has property $\mathcal{Q}$ (possibly as part of the initial graph G). The *depth* of the iterated BFSes is defined as the maximum recursion depth. To make the process deterministic, one can imagine the vertices of G being labeled from 1 to $|V(G)|$, and picking the vertex to start each BFS as the vertex of the smallest possible label. Thus, for every connected induced subgraph G' of G, we may speak of *the BFS* of G'. The proof of the Klein–Plotkin–Rao theorem establishes that in proper minor-closed classes the iterated BFSes of constant width to constant weak diameter have constant depth. We later turn these iterated BFSes into clustered edge-colorings, where the clustering depends on the maximum degree. In a rooted tree (such as BFS trees), we call *vertical path* any path between a pair of ancestor–descendant.

As the authors of [17] showed their result when excluding a minor rather than an induced minor, we will reproduce their proof with some light (not to say *minor*) adjustments. Something similar has been done by Lee [21], in a somewhat different language (see [20, Theorem 4.2]).

We start with a key lemma of [17] on the control given by two nested BFSes.

▶ **Lemma 9.** *Fix a positive integer h. Let G be a connected graph, and let $\mathcal{L}$ be the BFS of G from some $v \in V(G)$. Let G' be a subgraph of G induced by h consecutive layers of $\mathcal{L}$, and let $\mathcal{L}' = \{L'_0, L'_1, \dots, L'_s\}$ be a BFS of G'.*

Then, for every $w \in L'_\ell$, the path P from v to w in the BFS tree of $\mathcal{L}$ can only intersect layers of $\mathcal{L}'$ whose indices are in $[\ell-h+1, \ell+h-1]$, and only the last h vertices of P can be in $V(G')$.

Proof. The path P contains at most one vertex in each layer of $\mathcal{L}$. Since G' is contained in h layers of $\mathcal{L}$, P can only intersect $V(G')$ at its last h vertices. Moreover, the graph $P' = P \cap G'$ is connected. Next, suppose to the contrary that there is $u \in L_t \cap V(P)$ with $t \notin [\ell-h+1, \ell+h-1]$. Since P' is connected and $|V(P')| \leqslant h$, we have $d_{G'}(u, w) \leqslant h-1$, which contradicts the definition of a BFS layering. ◀

We now adapt [17] to work for induced minors.

▶ **Lemma 10** (Closely following [17]). *Let $p \geqslant 2$ and $q \geqslant 1$ be integers, and G be a graph excluding the 1-subdivision of $K_{p,q}$ as an induced minor. Then, the iterated BFSes of width h to weak diameter at most $d := ((8h+2)q + 4h + 6)(p-1)$ of G has depth at most $q+1$.*

Proof. Assume, for the sake of contradiction, that there is a (connected) graph G_{q+1} at recursion depth $q+1$ whose weak diameter is more than d (thereby pushing the recursion depth to at least $q+2$). We show how to build a model of an induced minor of 1-subd($K_{p,q}$) in G_1, a connected component of G. Let $G_1, G_2, \ldots, G_q, G_{q+1}$ be the (connected) graphs on the path from G_1 to G_{q+1} of the branching forest of the iterated BFSes (rooted at the connected components of G). Thus, for every $i \in [q]$, G_{i+1} is a connected component of a subgraph of G_i induced by h consecutive layers of the BFS of G_i. For each $i \in [q]$, let $\mathcal{T}_i$ be the BFS tree of the BFS of G_i.

We maintain the invariant that for every $i \in [q+1]$, G_i has an induced minor model of 1-subd($K_{p,q+1-i}$). Moreover, if $i > 1$, then denoting by $A_1(i), A_2(i), \ldots, A_p(i)$ the branch sets of the "left" side of the biclique, there is, for every $j \in [p]$, a vertical path $P_j(i)$ in $\mathcal{T}_{i-1}$ of length $4h+2$ with the following properties:

(1) $P_j(i)$ intersects exactly one branch set, $A_j(i)$;
(2) the deeper (in $\mathcal{T}_{i-1}$) endpoint $p_j(i)$ of $P_j(i)$ is in $A_j(i)$, and the other endpoint is not;
(3) the paths $P_1(i), P_2(i), \ldots, P_p(i)$ are pairwise vertex-disjoint and non-adjacent;
(4) no vertex of $P_j(i)$ is adjacent to a branch set B of the model for $B \neq A_j(i)$;
(5) each branch set $A_j(i)$ has size at most $(4h+1)(q+1-i)+1$; and
(6) every pair $u \in A_j(i), v \in A_{j'}(i)$ with $j \neq j'$ is at distance at least $(8h+2)(i-1)+4h+6$ in G_1 (or equivalently in G).

Note that if we show the invariant in G_{q+1}, and the induction step from G_{i+1} to G_i, the mere fact that G_1 admits 1-subd($K_{p,q}$) as an induced minor is our desired contradiction.

Base case. The process starts in G_{q+1}, which we assumed of weak diameter larger than d. We want to find an induced minor model of $K_{p,0}$ in G_{q+1}, and paths $P_1(q+1)$, $P_2(q+1), \ldots, P_p(q+1)$ in G_q satisfying the invariant. As G_{q+1} has weak diameter more than $d = ((8h+2)q+4h+6)(p-1)$ and is connected, it has in particular p vertices $p_1(q+1), \ldots, p_p(q+1)$ pairwise at distance at least $(8h+2)q+4h+6$ in G_1 (or G). We set $A_j(q+1) := \{p_j(q+1)\}$ for every $j \in [p]$. We choose $P_j(q+1)$ as the vertical path in $\mathcal{T}_q$ of length $4h+2$ with one endpoint being $p_j(q+1)$. Note that such a vertical path exists, otherwise $p_j(q+1)$ is at distance at most $4h+1$ from the root of $\mathcal{T}_q$. This would imply that every other $p_{j'}(q+1)$ (and at least one exists since $p \geqslant 2$) is at distance at most $5h$ from the root of $\mathcal{T}_q$ (recall indeed that every vertex of G_{q+1} is contained within h consecutive BFS layers of $\mathcal{T}_q$). In turn, this would make the pair $p_j(q+1), p_{j'}(q+1)$ at distance at most $9h+1 < (8h+2)q+4h+6$, a contradiction.

Since if $j \neq j'$, then the distance between $p_j(q+1)$ and $p_{j'}(q+1)$, is at least $(8h+2)q + 4h+6 > 8h+5$, one can see that the paths $P_j(q+1)$ are pairwise vertex-disjoint and non-adjacent. It is then clear that all the conditions of the invariant are met.

Induction. Let $\mathcal{B}_i$ be the induced minor model of 1-subd($K_{p,q+1-i}$) in G_i with $i > 1$, and $P_1(i), \ldots, P_p(i)$ be the vertical paths in $\mathcal{T}_{i-1}$ satisfying the invariant at step i. Let $A_1(i), \ldots, A_p(i)$ be the branch sets of $\mathcal{B}_i$ corresponding to the vertices on the "left" side of the biclique, and let $B_1, \ldots, B_{q+1-i}$ be the branch sets corresponding to the vertices on its "right" side. Finally, let $\mathcal{Z}_i \subset \mathcal{B}_i$ be the collection of the branch sets corresponding to the subdivision vertices.

Construction of the new model. The construction of $\mathcal{B}_{i-1}$, the induced minor model of 1-subd($K_{p,q+2-i}$) in G_{i-1}, is quite simple. We augment $\mathcal{B}_i$ in the following way. For every $j \in [p]$, let z_j be the endpoint of $P_j(i)$ that is *not* in $A_j(i)$. Let X be the union of the paths in $\mathcal{T}_{i-1}$ between each z_j and the root of $\mathcal{T}_{i-1}$.

We use X to build a branch set corresponding to the new vertex on the "right" side.

Namely, we set $B_{q+2-i} := X \setminus \{z_j \mid j \in [p]\}$. The branch sets corresponding to the subdivision vertices adjacent to the new vertex are the singletons $\{z_j\}$. Finally, we augment each $A_j(i)$ by adding to it the path $P_j(i)$. Specifically, for each $j \in [p]$, we define $A_j(i-1) := A_j(i) \cup P_j(i) \setminus \{z_j\}$. This finishes the construction of

$$\mathcal{B}_{i-1} := \{A_1(i-1), \ldots, A_p(i-1)\} \cup \{B_1, \ldots, B_{q+1-i}, B_{q+2-i}\} \cup \mathcal{Z}_i \cup \{\{z_1\}, \ldots, \{z_p\}\}.$$

We refer the reader to the representation of G_{i-1} and $\mathcal{B}_{i-1}$ in Figure 2.

$\mathcal{B}_{i-1}$ is an appropriate induced minor model. Let us first check that $\mathcal{B}_{i-1}$ is indeed an induced minor model of 1-subd$(K_{p,q+2-i})$ in G_{i-1}. Clearly, each set in $\mathcal{B}_{i-1}$ induces a connected subgraph of G_{i-1}. By Lemma 9, each path $P_j(i)$ has no more than its at most h last vertices in $V(G_i)$ (where the *first* vertex of the path is the one closer to the root of $\mathcal{T}_{i-1}$ in $\mathcal{T}_{i-1}$). The rest of its vertices are in $V(G_{i-1}) \setminus V(G_i)$. Thus, one can observe that every set among $B_{q+2-i}, \{z_1\}, \ldots, \{z_p\}$ is non-adjacent to every branch set of $\mathcal{B}_i$. By definition, B_{q+2-i} is adjacent to each of $\{z_1\}, \ldots, \{z_p\}$. By the invariant, all the paths $P_1(i), \ldots, P_p(i)$ are far away from each other, hence, the sets $\{z_1\}, \ldots, \{z_p\}$ are pairwise non-adjacent.

To conclude, it suffices to argue that B_{q+2-i} does not touch $A_j(i-1)$ for any $j \in [p]$, and that the rest of the adjacencies of $A_j(i-1)$ are as desired. It is clear that $A_j(i-1)$ is adjacent to $\{z_j\}$. By the invariant, $P_j(i)$ is only adjacent to one branch set of $\mathcal{B}_i$, namely, $A_j(i)$. Finally, an edge between B_{q+2-i} and $A_j(i-1)$ (or a common vertex) would create a path of length at most $1 + (4h+1) + (5h+1) = 9h+3$ between two sets $A_j(i), A_{j'}(i)$ for some $j' \neq j$, which contradicts the last item of the invariant at step i. The existence of such a path would indeed be guaranteed by the fact that vertices $p_j(i)$ and $p_{j'}(i)$ lie within h consecutive layers of BFS tree $\mathcal{T}_{i-1}$ (this explains the $5h+1$ as $h+4h+1$). In particular, $A_j(i-1)$ is not adjacent to neither B_{q+2-i} nor $\{z_{j'}\}$ for any $j' \neq j$. This concludes the proof that $\mathcal{B}_{i-1}$ is an induced minor model of 1-subd$(K_{p,q+2-i})$ in G_{i-1}.

Construction of the vertical paths. When $i = 2$, this finishes the induction step. Otherwise, we shall define the paths $P_1(i-1), \ldots, P_p(i-1)$ required by the invariant. First, we argue that the graph G_{i-1} is deep enough in $\mathcal{T}_{i-2}$. Let r be such that G_{i-1} is a connected component contained in the h consecutive layers of $\mathcal{T}_{i-2}$ whose index is in $[r, r+h-1]$. It follows that there is a path of length at most $2r + 2h - 2$ between any vertices in the branch sets of the model $\mathcal{B}_i$. By the last item of the invariant at step i, this implies that $2r + 2h - 2 \geqslant (8h+2)(i-1) + 4h + 6 \geqslant 12h + 8$. In particular, $r \geqslant 4h+2$. For every $j \in [p]$, let $p_j(i-1)$ be the vertex in the middle of $P_j(i)$. Since $r \geqslant 4h+2$, we can choose $P_j(i-1)$ as the unique vertical path of $\mathcal{T}_{i-2}$ whose length is $4h+2$ and whose extremity further away from the root of $\mathcal{T}_{i-2}$ is $p_j(i-1)$. See Figure 2 for an illustration.

Checking that all the items of the invariant hold (per paragraph: 1–2, 3, 4, 5–6). The first two items of the invariant hold by Lemma 9 combined with the fact that no path of length at most $(2h+1) + h + (2h+1) = 5h+2$ can exist between two distinct $A_j(i)$ and $A_{j'}(i)$; this is why we set the path length to $4h+2$, and why the next paths target the middle point of the previous paths. (This trick is called *moat* in [17].)

The third item holds, since if two distinct paths $P_j(i-1), P_{j'}(i-1)$ would touch (i.e., intersect or be adjacent), then a path of length at most $(2h+1) + (4h+2) + 1 + (4h+2) + (2h+1) = 12h+7$ would link a vertex of $A_j(i)$ to a vertex of $A_{j'}(i)$. Similarly, no path $P_j(i-1)$ is adjacent to a branch set $A_{j'}(i-1)$ with $j \neq j'$, as this would create a path of length at most $(2h+1) + (4h+2) + (4h+2)$ between $A_j(i)$ and $A_{j'}(i)$. Finally, no path $P_j(i-1)$ can touch X. Indeed, by Lemma 9 (in G_{i-1}), if X would touch a path $P_j(i-1)$ (for some $j \in [p]$), there would be a path of length at most $h+1$ between X and $p_j(i-1)$ in G_i. Let then x be a vertex in $X \cap N(P_j(i-1))$ linked to $p_j(i-1)$ by a path of length

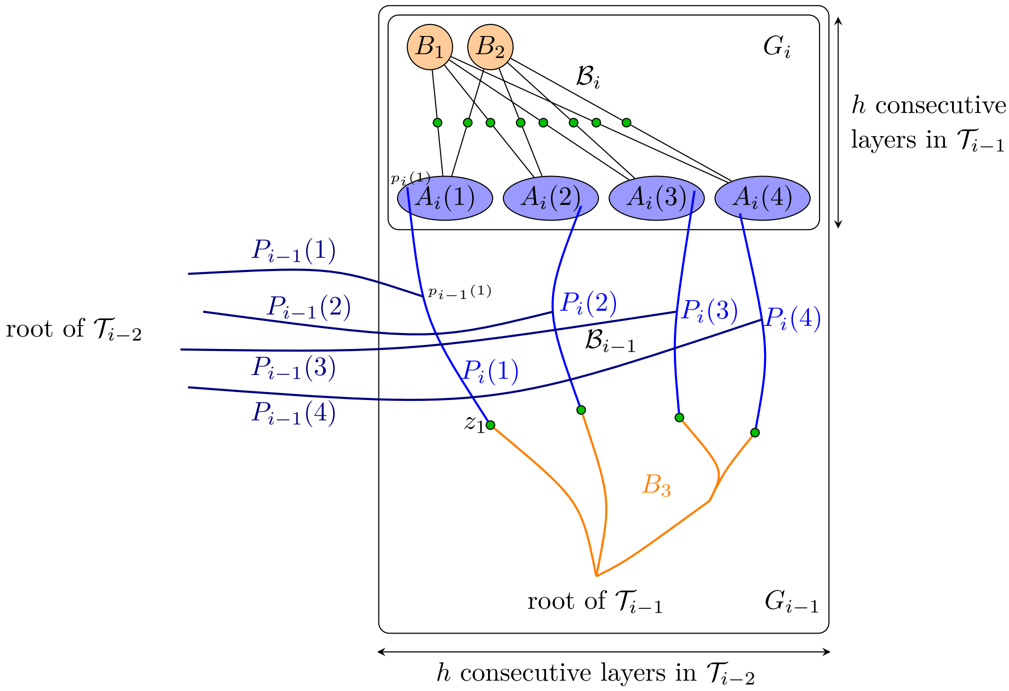

Figure 2 Illustration of the various labels used in the proof, and of how the induced minor model $\mathcal{B}_i$ in G_i (here at step $i = q-1$) is turned into the induced minor model $\mathcal{B}_{i-1}$ in G_{i-1}: The orange part of $\mathcal{T}_{i-1}$ gives a new branch set B_{q+2-i}, the vertices z_j (in green) are the new "subdivision" vertices, each path $P_j(i)$ (in lighter blue) is absorbed in $A_j(i)$, while the paths $P_j(i-1)$ (in darker blue) provide the next vertical paths.

at most $h+1$. In $\mathcal{T}_{i-1}$, vertex x has to be on a path to $z_{j'}$ with $j' \in [p] \setminus \{j\}$. But this contradicts that the distances of $p_j(i)$ and $p_{j'}(i)$ to the root of $\mathcal{T}_{i-1}$ differ by at most h (as these vertices lie within h consecutive layers of $\mathcal{T}_{i-1}$). Indeed $p_j(i)$ is at distance at most $(h+1)+(2h+1) = 3h+2$ from x, while $p_{j'}(i)$ is at distance at least $4h+2$ from x. Hence, the fourth item holds.

The fifth item holds since every $A_j(i)$ was added at most $4h+1$ vertices to become $A_j(i-1)$. For the same reason, sets $A_j(i-1)$ can only get closer to each other by $2 \cdot (4h+1) = 8h+2$, which confirms the last item. ◀

▶ **Lemma 11.** *Let $p \geqslant 2$ and $q \geqslant 1$ be integers, and G be a graph excluding the 1-subdivision of $K_{p,q}$ as an induced minor. Then G has a 2^{q+1}-edge-coloring with clustering $\Delta(G)^{18(q+1)(p-1)}$.*

Proof. Consider the iterated BFSes of width 2 to weak diameter at most $18(q+1)(p-1)-1$ of G. By Lemma 10 it has depth at most $q+1$. We show the present lemma by induction on the depth. If the depth is 0, the current vertex subset X has weak diameter at most $18(q+1)(p-1)-1$. Therefore $|X| \leqslant \Delta(G)^{18(q+1)(p-1)}$. We thus give the same color to all edges of $G[X]$.

If the depth is $d \geqslant 1$, then let $\{L_0, L_1, L_2, \ldots, L_s\}$ be the BFS layering of the current graph. At each new graph, connected component of $G[L_0 \cup L_1]$, $G[L_1 \cup L_2]$, $G[L_2 \cup L_3]$, …, $G[L_{s-1} \cup L_s]$, the recursion depth is at most $d-1$. By induction, every such induced subgraph can be 2^{d-1}-edge-colored with clustering $\Delta(G)^{18(q+1)(p-1)}$.

Observe that the colorings of the connected components of $G[L_0 \cup L_1], G[L_2 \cup L_3], \ldots$ define a 2^{d-1}-edge-coloring with clustering $\Delta(G)^{18(q+1)(p-1)}$ of $G[L_0 \cup L_1] \uplus G[L_2 \cup L_3] \uplus \ldots$ as a whole. The same applies to $G[L_1 \cup L_2], G[L_3 \cup L_4], \dots$ This yields for the current graph an edge-coloring with clustering $\Delta(G)^{18(q+1)(p-1)}$ using at most $2^{d-1} + 2^{d-1} = 2^d$ colors. Indeed, all its edges are covered by $G[L_0 \cup L_1] \uplus G[L_2 \cup L_3] \uplus \dots$ and $G[L_1 \cup L_2] \uplus G[L_3 \cup L_4] \uplus \dots$ Edges that appear in both graphs can take any of the two colors they are assigned to.

We thus get a 2^{q+1}-edge-coloring with clustering $\Delta(G)^{18(q+1)(p-1)}$ for G. ◀

6 Putting things together

We first show an upper bound on the treewidth of graphs excluding a grid as an induced minor and admitting edge-colorings with few colors and moderately large clustering.

▶ **Lemma 12.** *Every graph G excluding a k-vertex planar graph as an induced minor and admitting an h-edge-coloring with clustering $c > 0$ has treewidth at most $k^{O(1)} 2^{O(h^5 + h \log c)}$.*

Proof. By Lemma 8, G admits an induced subgraph G' of maximum degree at most $3h$ and

$$\mathrm{tw}(G)/(c \log^\delta \mathrm{tw}(G))^h \leqslant \mathrm{tw}(G'),$$

with δ the constant of Theorem 6.

As G excludes a k-vertex planar graph as an induced minor, so does G'. Thus by Theorem 1,

$$\mathrm{tw}(G') \leqslant k^\gamma 2^{\Delta(G')^5} \leqslant k^\gamma 2^{243h^5},$$

for some universal constant γ.

From the two previous inequalities, we get that

$$\mathrm{tw}(G)/(c \log^\delta \mathrm{tw}(G))^h \leqslant k^\gamma 2^{243h^5}.$$

If $\log^{\delta h} \mathrm{tw}(G) \leqslant \sqrt{\mathrm{tw}(G)}$, we get that

$$\mathrm{tw}(G) \leqslant c^{2h} k^{2\gamma} 2^{2 \cdot 243h^5} = k^{O(1)} 2^{O(h^5 + h \log c)},$$

as claimed. If instead $\log^{\delta h} \mathrm{tw}(G) > \sqrt{\mathrm{tw}(G)}$, the statement of the lemma also holds, as then $\mathrm{tw}(G) = 2^{O(h^2)}$. ◀

We obtain our main result as a consequence of Lemma 12.

▶ **Theorem 3.** *Every graph G excluding as induced minors a k-vertex planar graph and a q-vertex graph has treewidth at most $k^{O(1)} \cdot \Delta(G)^{f(q)}$ with $f(q) = 2^{O(q)}$.*

Proof. As G excludes a q-vertex graph as an induced minor, it also excludes as such the 1-subdivision of $K_{q^2,q}$, which admits every q-vertex graph as an induced minor. We can assume that $q \geqslant 2$ (hence $q^2 \geqslant 2$) as otherwise G is edgeless, hence has treewidth 0. Thus, by Lemma 11, G admits a 2^{q+1}-edge-coloring with clustering $\Delta(G)^{18(q+1)(q^2-1)}$. We can assume that $\Delta(G) \geqslant 2$, as otherwise G has treewidth $\leqslant 1$. By Lemma 12,

$$\mathrm{tw}(G) \leqslant k^{O(1)} 2^{O(2^{5(q+1)} + 2^{q+1}(q+1)(q^2-1) \log \Delta(G))} = k^{O(1)} \Delta(G)^{f(q)},$$

with $f(q) = 2^{O(q)}$, as claimed. ◀

7 Clusters of bounded treewidth

The goal of this section is to relax the notion of *clustering* of edge-colorings so that Lemma 12 still holds. Namely, we now allow clusters to be arbitrarily large, however, we want their treewidth to be bounded. This can be converted into an edge-coloring (still with few colors) with bounded clustering. Indeed, we show that a graph G of bounded treewidth admits a 3-edge-coloring with clustering $f(\Delta(G))$.

The main tool we plan to use is the notion of *tree-partitions* of graphs. A pair $(T, \{B_x : x \in V(T)\})$ is a tree-partition of a graph G if T is a tree, and $\{B_x : x \in V(T)\}$ is a partition of $V(G)$ such that for every $uv \in E(G)$, there exist a pair $x, y \in V(T)$ of equal or adjacent vertices such that $u \in B_x, v \in B_y$. The *width* of a tree-partition $(T, \{B_x : x \in V(T)\})$ is defined as the maximum cardinality of an element of $\{B_x : x \in V(T)\}$. The *tree-partition width* of a graph G, denoted $\text{tpw}(G)$, is the minimum width of a tree-partition of G. An anonymous referee of [9] showed that every graph G has tree-partition width of at most $24\text{tw}(G)\Delta(G)$ (see also [27, 10]).

▶ **Lemma 13.** *Every graph G admits a 3-edge-coloring with clustering $tpw(G)(\Delta(G)+1)$, which is in particular $O(tw(G)\Delta(G)^2)$.*

Proof. Let $(T, \{B_x : x \in V(T)\})$ be a tree-partition of G of width w. We root T at an arbitrary vertex r. Assign to each vertex $x \in V(T)$ its distance to r in T, denoted $\text{depth}(x)$. We define a coloring $\text{col} : E(G) \to \{0, 1, 2\}$ as follows. Let $uv \in E(G)$, and let $x \in V(T), y \in N_T[x]$ be such that $u \in B_x$ and $v \in B_y$. Without loss of generality assume that $\text{depth}(x) \leqslant \text{depth}(y)$. If $x \neq y$, then we set $\text{col}(uv) = \text{depth}(x) \bmod 2$, and otherwise, we set $\text{col}(uv) = 2$. Every monochromatic connected component of color 2 is contained in a single part B_x for some $x \in V(T)$, and so, its cardinality is at most w. On the other hand, for every monochromatic connected component of color 0 or 1, there exists a single part B_x for some $x \in V(T)$ such that every edge in the component is incident to a vertex in B_x. It follows that the size of this monochromatic component is at most $w \cdot (\Delta(G) + 1)$. ◀

Now, we state and prove a relaxed version of Lemma 12.

▶ **Lemma 14.** *Suppose graph G excludes as an induced minor a k-vertex planar graph and admits an edge-coloring col_1 with color classes $F_1, \dots, F_h$ such that for each $i \in [h]$, the graph $(V(G), F_i)$ has treewidth at most t. Then the treewidth of G is at most $k^{O(1)}t^{O(h^5)}\Delta(G)^{O(h^5)}$.*

Proof. By Lemma 13, for each $i \in [h]$, the graph $(V(G), F_i)$ admits a 3-edge-coloring with clustering $O(t\Delta(G)^2)$. Since $\{F_1, \dots, F_h\}$ is a partition of $V(G)$, the above edge-colorings give a 3-edge-coloring col_2 of $E(G)$. Consider the product edge-coloring col of col_1 and col_2 of $E(G)$, that is, $\text{col}(e) = (\text{col}_1(e), \text{col}_2(e))$ for every $e \in E(G)$. Observe that col uses at most $3h$ colors and has clustering $O(t\Delta(G)^2)$. Finally, by Lemma 12, we obtain

$$\text{tw}(G) \leqslant k^{O(1)} \cdot (t\Delta(G)^2)^{O(h^5)} = k^{O(1)}t^{O(h^5)}\Delta(G)^{O(h^5)},$$

as claimed. ◀

8 Product structure

The *strong product of graphs H_1 and H_2*, denoted by $H_1 \boxtimes H_2$, is the graph with vertex set $V(H_1) \times V(H_2)$ such that there is an edge $(u, v)(u', v')$ whenever either $u = u'$ and $vv' \in E(H_2)$, or $uu' \in E(H_1)$ and $v = v'$, or $uu' \in E(H_1)$ and $vv' \in E(H_2)$. We prove the following theorem.

▶ **Theorem 4.** *Let H be a graph of treewidth at most t, and P be a path. Let G be a subgraph of $H \boxtimes P$ excluding a k-vertex planar graph as an induced minor. Then the treewidth of G is at most $k^{O(1)} \cdot t^{O(1)} \cdot \Delta(G)^{O(1)}$.*

Proof. We claim that $H \boxtimes P$ admits a 3-edge-coloring such that if F is any of its color classes, then the graph $(V(H \boxtimes P), F)$ has treewidth at most $2t$. First, note that this suffices to prove the theorem. Indeed, we can restrict this edge-coloring to G and apply Lemma 14 with $h = 3$ to end the proof.

Let us justify the initial claim. We construct a coloring $\text{col} : E(H \boxtimes P) \to \{0, 1, 2\}$. Let $P = v_1 v_2 \dots v_m$. We set the color of each edge $(u, v)(u', v')$ such that $v = v'$ to 2, and each edge $(u, v)(u', v')$ such that $v \neq v'$ to $i \bmod 2$, where i the positive integer satisfying $\{v, v'\} = \{v_i, v_{i+1}\}$. The graph G restricted to edges of color 2 is simply a disjoint union of copies of H, hence, it has treewidth at most t. On the other hand, the graph G restricted to edges of color 0 or 1 is a disjoint union of copies of the graph $H' = H \boxtimes K_2$. Thus $\text{tw}(H') \leqslant 2\text{tw}(H) \leqslant 2t$, which ends the proof. ◀

References

1 Marthe Bonamy, Édouard Bonnet, Hugues Déprés, Louis Esperet, Colin Geniet, Claire Hilaire, Stéphan Thomassé, and Alexandra Wesolek. Sparse graphs with bounded induced cycle packing number have logarithmic treewidth. In Nikhil Bansal and Viswanath Nagarajan, editors, *Proceedings of the 2023 ACM-SIAM Symposium on Discrete Algorithms, SODA 2023, Florence, Italy, January 22–25, 2023*, pages 3006–3028. SIAM, 2023. doi:10.1137/1.9781611977554.ch116.

2 Édouard Bonnet, Julien Duron, Colin Geniet, Stéphan Thomassé, and Alexandra Wesolek. Maximum independent set when excluding an induced minor: $K_1 + tK_2$ and $tC_3 \uplus C_4$. *CoRR*, abs/2302.08182, to appear at ESA 2023. arXiv:2302.08182.

3 Chandra Chekuri and Julia Chuzhoy. Degree-3 treewidth sparsifiers. In Piotr Indyk, editor, *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015, San Diego, CA, USA, January 4–6, 2015*, pages 242–255. SIAM, 2015. doi:10.1137/1.9781611973730.19.

4 Maria Chudnovsky, Neeldhara Misra, Daniël Paulusma, Oliver Schaudt, and Akanksha Agrawal. Vertex Partitioning in Graphs: From Structure to Algorithms (Dagstuhl Seminar 22481). *Dagstuhl Reports*, 12(11):109–123, 2023. doi:10.4230/DagRep.12.11.109.

5 Clément Dallard, Martin Milanič, and Kenny Štorgel. Treewidth versus clique number. I. graph classes with a forbidden structure. *SIAM J. Discret. Math.*, 35(4):2618–2646, 2021. doi:10.1137/20M1352119.

6 Clément Dallard, Martin Milanič, and Kenny Štorgel. Treewidth versus clique number. III. tree-independence number of graphs with a forbidden structure. *CoRR*, abs/2206.15092, 2022. arXiv:2206.15092.

7 Clément Dallard, Martin Milanič, and Kenny Štorgel. Treewidth versus clique number. II. tree-independence number, 2021. doi:10.48550/ARXIV.2111.04543.

8 James Davies. Oberwolfach report 1/2022. doi:10.4171/OWR/2022/1., 2022.

9 Guoli Ding and Bogdan Oporowski. Some results on tree decomposition of graphs. *Journal of Graph Theory*, 20(4):481–499, 1995. doi:10.1002/jgt.3190200412.

10 Marc Distel and David R. Wood. Tree-partitions with small bounded degree trees, 2023. arXiv:2210.12577.

11 Vida Dujmović, Pat Morin, and David R. Wood. Graph product structure for non-minor-closed classes. *Journal of Combinatorial Theory, Series B*, 162:34–67, 2023. doi:10.1016/j.jctb.2023.03.004.

12 Zdenek Dvorák and Sergey Norin. Treewidth of graphs with balanced separations. *J. Comb. Theory, Ser. B*, 137:137–144, 2019. URL: https://doi.org/10.1016/j.jctb.2018.12.007, doi:10.1016/J.JCTB.2018.12.007.

13 Peter Gartland and Daniel Lokshtanov. Independent set on P_k-free graphs in quasi-polynomial time. In Sandy Irani, editor, *61st IEEE Annual Symposium on Foundations of Computer Science, FOCS 2020, Durham, NC, USA, November 16–19, 2020*, pages 613–624. IEEE, 2020. doi:10.1109/FOCS46700.2020.00063.

14 Peter Gartland and Daniel Lokshtanov. private communication, 2023.

15 Peter Gartland, Daniel Lokshtanov, Tomáš Masařík, Marcin Pilipczuk, Michał Pilipczuk, and Paweł Rzążewski. Maximum weight independent set in graphs with no long claws in quasi-polynomial time. *CoRR*, abs/2305.15738, 2023. arXiv:2305.15738.

16 Peter Gartland, Daniel Lokshtanov, Marcin Pilipczuk, Michał Pilipczuk, and Paweł Rzążewski. Finding large induced sparse subgraphs in $C_{>t}$-free graphs in quasipolynomial time. In Samir Khuller and Virginia Vassilevska Williams, editors, *STOC '21: 53rd Annual ACM SIGACT Symposium on Theory of Computing, Virtual Event, Italy, June 21–25, 2021*, pages 330–341. ACM, 2021. doi:10.1145/3406325.3451034.

17 Philip N. Klein, Serge A. Plotkin, and Satish Rao. Excluded minors, network decomposition, and multicommodity flow. In S. Rao Kosaraju, David S. Johnson, and Alok Aggarwal, editors, *Proceedings of the Twenty-Fifth Annual ACM Symposium on Theory of Computing, May 16–18, 1993, San Diego, CA, USA*, pages 682–690. ACM, 1993. doi:10.1145/167088.167261.

18 Tuukka Korhonen. Grid induced minor theorem for graphs of small degree. *Journal of Combinatorial Theory, Series B*, 160:206–214, 2023. doi:10.1016/j.jctb.2023.01.002.

19 Tuukka Korhonen and Daniel Lokshtanov. Induced-minor-free graphs: Separator theorem, subexponential algorithms, and improved hardness of recognition. *CoRR*, abs/2308.04795, 2023. arXiv:2308.04795.

20 James R. Lee. Separators in region intersection graphs. *CoRR*, abs/1608.01612, 2016. arXiv:1608.01612.

21 James R. Lee. Separators in region intersection graphs. In Christos H. Papadimitriou, editor, *8th Innovations in Theoretical Computer Science Conference, ITCS 2017, January 9–11, 2017, Berkeley, CA, USA*, volume 67 of *LIPIcs*, pages 1:1–1:8. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2017. doi:10.4230/LIPIcs.ITCS.2017.1.

22 Marcin Pilipczuk, Michal Pilipczuk, and Pawel Rzążewski. Quasi-polynomial-time algorithm for independent set in P_t-free graphs via shrinking the space of induced paths. In Hung Viet Le and Valerie King, editors, *4th Symposium on Simplicity in Algorithms, SOSA 2021, Virtual Conference, January 11–12, 2021*, pages 204–209. SIAM, 2021. doi:10.1137/1.9781611976496.23.

23 Andrei Cosmin Pohoata. Unavoidable induced subgraphs of large graphs. 2014.

24 Frank P. Ramsey. On a problem of formal logic. In *Proc. London Math. Soc. series 2*, volume 30 of *264–286*, 1930.

25 Neil Robertson and Paul 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.

26 Neil Robertson, Paul D. Seymour, and Robin Thomas. Quickly excluding a planar graph. *Journal of Combinatorial Theory, Series B*, 62(2):323–348, 1994. doi:10.1006/jctb.1994.1073.

27 David R. Wood. On tree-partition-width. *European Journal of Combinatorics*, 30(5):1245–1253, 2009. Part Special Issue on Metric Graph Theory. doi:10.1016/j.ejc.2008.11.010.