INDUCED SUBGRAPHS AND TREE DECOMPOSITIONS XII. GRID THEOREM FOR PINCHED GRAPHS

BOGDAN ALECU[**¶], MARIA CHUDNOVSKY[*II], SEPEHR HAJEBI[§], AND SOPHIE SPIRKL[§‖]

Abstract. Given $c \in \mathbb{N}$, we say a graph G is *c-pinched* if G does not contain an induced subgraph consisting of c cycles, all going through a single common vertex and otherwise pairwise disjoint and with no edges between them. What can be said about the structure of c-pinched graphs?

For instance, 1-pinched graphs are exactly graphs of treewidth 1. However, bounded treewidth for $c > 1$ is immediately seen to be a false hope because complete graphs, complete bipartite graphs, subdivided walls and line graphs of subdivided walls are all examples of 2-pinched graphs with arbitrarily large treewidth. There is even a fifth obstruction for larger values of c, discovered by Pohoata and later independently by Davies, consisting of 3-pinched graphs with unbounded treewidth and no large induced subgraph isomorphic to any of the first four obstructions.

We fuse the above five examples into a grid-type theorem fully describing the unavoidable induced subgraphs of pinched graphs with large treewidth. More precisely, we prove that for *every* $c \in \mathbb{N}$, a c-pinched graph G has large treewidth if and only if G contains one of the following as an induced subgraph: a large complete graph, a large complete bipartite graph, a subdivision of a large wall, the line graph of a subdivision of a large wall, or a large graph from the Pohoata-Davies construction. Our main result also generalizes to an extension of pinched graphs where the lengths of excluded cycles are lower-bounded.

1. Introduction

1.1. **Background.** The set of all positive integers is denoted by $\mathbb{N}$. Graphs in this paper have finite vertex sets, no loops and no parallel edges. Let G be a graph. For $X \subseteq V(G)$, we denote by $G[X]$ the subgraph of G induced by X, and by $G \setminus X$ the induced subgraph of G obtained by removing X. We use induced subgraphs and their vertex sets interchangeably. For graphs G and H, we say G *contains* H if G has an induced subgraph isomorphic to H, and we say G is *H-free* if G does not contain H. A class of graphs is *hereditary* if it is closed under isomorphism and taking induced subgraphs.

Date: 21st January, 2025.

[*] Princeton University, Princeton, NJ, USA.

[**] School of Computing, University of Leeds, Leeds, UK.

[§] Department of Combinatorics and Optimization, University of Waterloo, Waterloo, Ontario, Canada.

[II] Supported by NSF-EPSRC Grant DMS-2120644 and by AFOSR grant FA9550-22-1-0083.

[¶] Supported by DMS-EPSRC Grant EP/V002813/1.

[‖] We acknowledge the support of the Natural Sciences and Engineering Research Council of Canada (NSERC), [funding reference number RGPIN-2020-03912]. Cette recherche a été financée par le Conseil de recherches en sciences naturelles et en génie du Canada (CRSNG), [numéro de référence RGPIN-2020-03912]. This project was funded in part by the Government of Ontario.

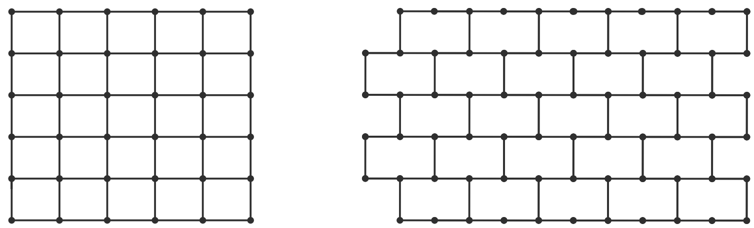

Figure 1. The 6-by-6 square grid (left) and the 6-by-6 wall $W_{6\times 6}$ (right).

The *treewidth* of a graph G (denoted by $\mathrm{tw}(G)$) is the smallest $w \in \mathbb{N}$ for which there exists a tree T as well as an assignment $(T_v : v \in V(G))$ of non-empty subtrees of T to the vertices of G with the following specifications.

(T1) For every edge $uv \in V(G)$, T_u and T_v share at least one vertex.
(T2) For every $x \in V(T)$, there are at most $w + 1$ vertices $v \in V(G)$ for which $x \in V(T_v)$.

As one of the most extensively studied graph invariants, the enduring interest in treewidth is partly explained by its role in the development of Robertson and Seymour's graph minors project, as well as the vast range of nice structural [15] and algorithmic [5] properties of graphs of small treewidth.

Graphs of large treewidth have also been a central topic of research for several decades. Usually, it is most desirable to certify large treewidth in a graph G by means of a well-understood "obstruction" which still has relatively large treewidth, and which lies in G under a certain containment relation. The cornerstone result in this category is the so-called *Grid Theorem* of Robertson and Seymour [15], Theorem 1.1 below, which says that under two of the most studied graph containment relations, namely the graph minor relation and the subgraph relation, the only obstructions to bounded treewidth are the "basic" ones: the t-by-t square grid for minors, and subdivisions of the t-by-t hexagonal grid for subgraphs. The t-by-t hexagonal grid is also known as the *t-by-t wall*, denoted by $W_{t\times t}$ (see Figure 1, and also [3] for full definitions).

Theorem 1.1 (Robertson and Seymour [15]). *For every $t \in \mathbb{N}$, every graph of sufficiently large treewidth contains the t-by-t square grid as a minor, or equivalently, a subdivision of $W_{t\times t}$ as a subgraph.*

It is therefore tempting to inquire about an analogue of Theorem 1.1 for another standard graph containment relation: induced subgraphs. The basic obstructions in this case already suggest that a more involved grid-type theorem is to be expected: complete graphs, complete bipartite graphs, subdivided walls, and line graphs of subdivided walls are all examples of induced-subgraph-minimal graphs with large treewidth. It is convenient to group all these graphs together. Given $t \in \mathbb{N}$, we say a graph H is a *t-basic obstruction* if H is isomorphic to one of the following: the complete graph K_t, the complete bipartite graph $K_{t,t}$, a subdivision of $W_{t\times t}$, or the line graph of a subdivision of $W_{t\times t}$, where the *line graph* $L(F)$ of a graph F is the graph with vertex set $E(F)$, such that two vertices of $L(F)$ are adjacent if and only

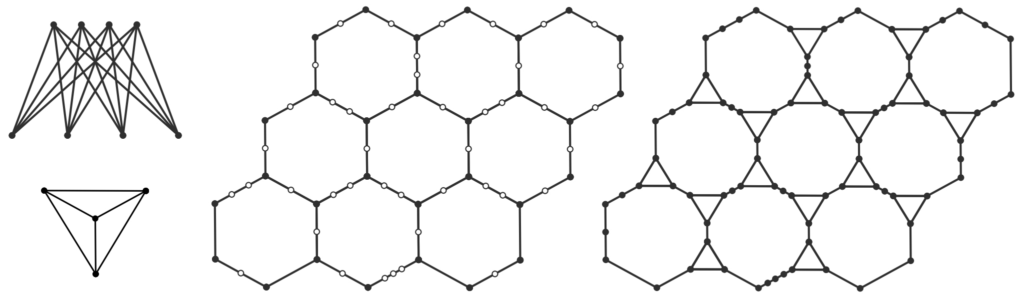

Figure 2. The 4-basic obstructions.

if the corresponding edges of F share an end (see Figure 2). We say a graph G is *t-clean* if G does not contain a t-basic obstruction (as an induced subgraph). A graph class $\mathcal{G}$ is *clean* if for every $t \in \mathbb{N}$, there is a constant $w(t) \in \mathbb{N}$ (depending on $\mathcal{G}$) for which every t-clean graph in $\mathcal{G}$ has treewidth at most $w(t)$. Since the basic obstructions have unbounded treewidth (K_{t+1}, $K_{t,t}$, subdivisions of $W_{t\times t}$ and line graphs of subdivisions of $W_{t\times t}$ are all known to have treewidth t), it follows that for every hereditary class of bounded treewidth, there exists some $t \in \mathbb{N}$ such that every graph in the class is t-clean. The converse would be a particularly nice grid-type theorem for induced subgraphs: *every hereditary class is clean.* This, however, is now known to be far from true, thanks to the numerous constructions [6, 7, 13, 16] of graphs with arbitrarily large treewidth which are t-clean for small values of t (and we will take a closer look at the one from [7, 13] in a moment).

On the other hand, there are several hereditary classes that are known to be clean for highly non-trivial reasons. As a notable example, Korhonen [11] proved that every graph class of bounded maximum degree is clean, settling a conjecture from [1]:

Theorem 1.2 (Korhonen [11])**.** *For every $d \in \mathbb{N}$, the class of all graphs with maximum degree at most d is clean.*

One possible attempt at generalizing Theorem 1.2 is to look for clean classes under weaker assumptions than bounded maximum degree. For instance, bounded maximum degree is equivalent to excluding a fixed star as a subgraph, and in a recent joint work with Abrishami, we extended Theorem 1.2 to graphs that exclude a fixed subdivided star as an induced subgraph. In fact, we proved:

Theorem 1.3 (Abrishami, Alecu, Chudnovsky, Hajebi, Spirkl [2])**.** *Let H be a graph. Then the class of all H-free graphs is clean if and only if every component of H is a subdivided star.*

Another natural candidate for a configuration forcing large-degree vertices consists of several cycles sharing a single vertex. For $c \in \mathbb{N}$, let us say a graph G is *c-pinched* if G does not contain c induced cycles, all going through a common vertex and otherwise pairwise disjoint and anticomplete (for disjoint subsets X, Y of vertices in a graph G, we say that X is *anticomplete* to Y if no edges between X and Y are present in G, and that X is *complete*

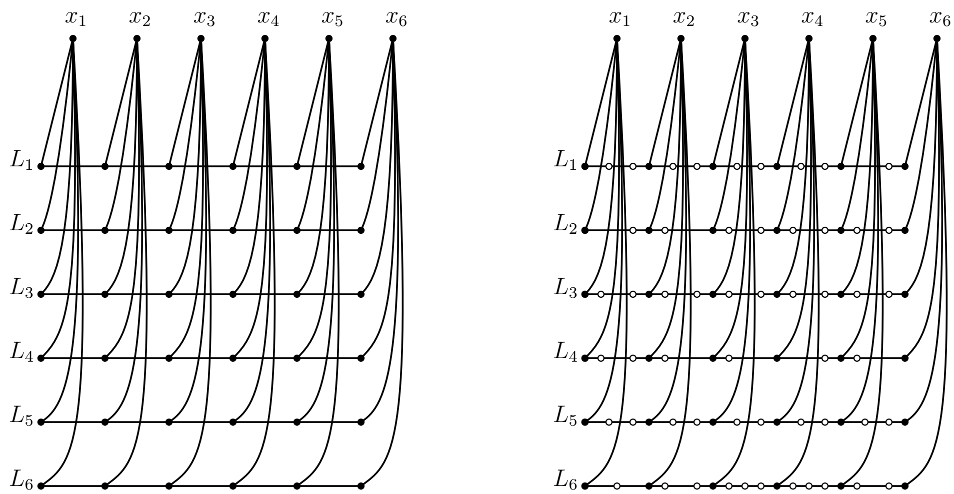

FIGURE 3. The graph PD_6 (left) and an expansion of PD_6 (right).

to Y if all edges with an end in X and an end in Y are present in G). Note that 1-pinched graphs are forests, which are the only graphs with treewidth 1. For $c = 2$, it is easily seen that all basic obstructions are 2-pinched, and we will show that they are the only representatives of large treewidth in 2-pinched graphs:

Theorem 1.4. *The class of all* 2*-pinched graphs is clean.*

One may then ask if the class of c-pinched graphs is clean for all $c \in \mathbb{N}$. But this is false, as shown by a 10-year-old construction due to Pohoata [13], also re-discovered recently by Davies [7], which we describe below.

Given an integer k, we write $\mathbb{N}_k$ to denote the set of all positive integers less than or equal to k (so we have $\mathbb{N}_k = \varnothing$ if and only if $k \leq 0$). For $s \in \mathbb{N}$, let PD_s be the graph whose vertex set can be partitioned into a stable set $S = \{x_1, \ldots, x_s\}$, and s pairwise disjoint induced paths $L_1, \ldots, L_s$ with no edges between them, such that the following hold.

- (PD1) For every $i \in \mathbb{N}_s$, L_i has length $s - 1$ (and so exactly s vertices).
- (PD2) For every $i \in \mathbb{N}_s$, the vertices in the interior of L_i may be enumerated from one end to the other as u^i_1-$\cdots$-u^i_s such that for every $j \in \mathbb{N}_s$, x_j has **exactly** one neighbor in $V(L_i)$, namely u^i_j.

See Figure 3. It is straightforward to show that for every $s \in \mathbb{N}$, PD_s is a 4-clean, 3-pinched graph of treewidth at least s. More generally, as we will prove in Theorem 3.1, the same holds for every graph obtained from PD_s by subdividing the edges of $L_1, \ldots, L_s$ arbitrarily. We refer to these graphs as *expansions* of PD_s (see Figure 3). It follows that expansions of the Pohoata-Davies graphs are “non-basic” obstructions to bounded treewidth in pinched graphs. Strikingly, the converse turns out to hold, too. We prove that:

Theorem 1.5. *For all $c, s, t \in \mathbb{N}$, every c-pinched graph of sufficiently large treewidth contains either a t-basic obstruction or an expansion of PD_s.*

Since the basic obstructions and the graphs PD_s have arbitrarily large treewidth, Theorem 1.5 provides a full grid-type theorem for the class of c-pinched graphs for all $c \in \mathbb{N}$.

More generally, our main result in this paper, Theorem 3.2, renders a complete description of the induced subgraph obstructions to bounded treewidth in the class of (c, h)-*pinched* graphs for all $c, h \in \mathbb{N}$, that is, graphs containing no c induced cycles each of length at least $h + 2$, all going through a common vertex and otherwise pairwise disjoint and anticomplete (so a graph is c-pinched if and only if it is $(c, 1)$-pinched). Indeed, the strengthening is direct enough that Theorem 1.5 is the special case of Theorem 3.2 where $h = 1$. Note also that no expansion of the graph PD_4 is 2-pinched. Therefore, Theorem 1.5 implies Theorem 1.4.

Let us remark that grid-type theorems involving non-basic obstructions, such as Theorem 1.5 (or Theorem 3.2, rather), are as of yet quite rare. Indeed, the only other examples we are aware of are the analogous result for "c-perforated" graphs – which we proved recently [4] – and an result from [10] concerning the class of "circle graphs." In fact, the proof of Theorem 3.2 bears a close resemblance to that of the main result of [4], and crucially builds on some tools developed there.

To elaborate, let $\mathcal{G}$ be the non-clean class for which we wish to prove a grid-type theorem. Roughly speaking, the idea is to break the proof into two steps: first, we show that every t-clean graph in the class with sufficiently large treewidth must contain an "approximate version" of the non-basic obstruction we are looking for, and second, we perform further analysis on the approximate version in pursuit of the exact one. Luckily, in the case of pinched graphs, the approximate and the exact non-basic obstructions are actually quite close. Note that the (expansions of) Pohoata-Davies graphs consist of a stable set and a number of pairwise disjoint and anticomplete paths such that every vertex in the stable set has a neighbor in every path. We call such a configuration in a graph G a "constellation" (a notion which was also used in [4] as an approximate non-basic obstruction for perforated graphs).

As for the present paper, our first goal is to show that for all $c, t, h \in \mathbb{N}$, every t-clean (c, h)-pinched graph of sufficiently large treewidth contains a huge constellation. This involves a useful result from an earlier paper in this series [2] concerning the "local connectivity" in clean classes, accompanied by a collection of Ramsey-type arguments to tidy up an induced subgraph of G with high local connectivity. The second step then is to turn a constellation into an expansion of a Pohoata-Davies-like graph. To that end, for every path L in the "path side" of the constellation, we consider the intersection graph I of the minimal subpaths of L containing all neighbors of each vertex in the "stable set side" S. Provided that S is large enough, I contains either a big stable set or a big clique. In the former case, on L, the neighbors of the vertices from the stable set do not interlace, and the resulting "alignment" of vertices according to their neighbors on L signals the emergence of a Pohoata-Davies-like structure. In the latter case, several vertices in S turn out to have neighbors in several pairwise disjoint and anticomplete subpaths of L. This eventually yields c induced cycles with a vertex in common and otherwise pairwise disjoint and anticomplete, a contradiction.

We take the two steps above in the reverse order in Sections 4 and 5, respectively. In the next section, we discuss the connectivity result from [2] (together with its bells and whistles). Section 3 introduces a variety of notions from [4] which we use in this paper, and also features the statement of our main result, Theorem 3.2, of which we give a complete proof in Section 6.

2. Blocks

We begin with a couple of definitions. Let $G = (V(G), E(G))$ be a graph. For an induced subgraph H of G and a vertex $x \in V(G)$, we denote by $N_H(x)$ the set of all neighbors of x in H, and write $N_H[x] = N_H(x) \cup \{x\}$. A *stable set in* G is a set of pairwise non-adjacent vertices. A *path in* G is an induced subgraph of G which is a path. If P is a path in G, we write $P = p_1\text{-}\cdots\text{-}p_k$ to mean that $V(P) = \{p_1, \ldots, p_k\}$ and p_i is adjacent to p_j if and only if $|i - j| = 1$. We call the vertices p_1 and p_k the *ends of* P and write $\partial P = \{p_1, p_k\}$. The *interior of* P, denoted P^*, is the set $P \setminus \partial P$. For $x, y \in V(P)$, we denote by $x\text{-}P\text{-}y$ the subpath of P with ends x, y. The *length* of a path is its number of edges. Similarly, a *cycle in* G is an induced subgraph of G that is a cycle. If C is a cycle in G, we write $C = c_1\text{-}\cdots\text{-}c_k\text{-}c_1$ to mean that $V(C) = \{c_1, \ldots, c_k\}$ and c_i is adjacent to c_j if and only if $|i - j| \in \{1, k - 1\}$. The *length* of a cycle is also its number of edges. For a collection $\mathcal{P}$ of paths in G, we adopt the notations $V(\mathcal{P}) = \bigcup_{P \in \mathcal{P}} V(P)$, $\mathcal{P}^* = \bigcup_{P \in \mathcal{P}} P^*$ and $\partial \mathcal{P} = \bigcup_{P \in \mathcal{P}} \partial P$.

Let $k \in \mathbb{N}$ and let G be a graph. A *k-block* in G is a pair $(B, \mathcal{P})$ where $B \subseteq V(G)$ with $|B| \geq k$ and $\mathcal{P} : \binom{B}{2} \to 2^{V(G)}$ is map such that $\mathcal{P}_{\{x,y\}} = \mathcal{P}(\{x, y\})$, for each 2-subset $\{x, y\}$ of B, is a set of at least k pairwise internally disjoint paths in G from x to y. We say that $(B, \mathcal{P})$ is *strong* if for all distinct 2-subsets $\{x, y\}, \{x', y'\}$ of B, we have $V(\mathcal{P}_{\{x,y\}}) \cap V(\mathcal{P}_{\{x',y'\}}) = \{x, y\} \cap \{x', y'\}$; that is, each path $P \in \mathcal{P}_{\{x,y\}}$ is disjoint from each path $P' \in \mathcal{P}_{\{x',y'\}}$, except P and P' may share an end. In [2], with Abrishami we proved the following:

Theorem 2.1 (Abrishami, Alecu, Chudnovsky, Hajebi, Spirkl [2])**.** *For all $k, t \in \mathbb{N}$, there is a constant $\xi = \xi(k, t) \in \mathbb{N}$ such that for every t-clean graph G of treewidth more than ξ, there is a strong k-block in G.*

The following result, which guarantees that the graphs we work with exclude "short subdivisions" of large complete graphs, paves the way for our application of Theorem 2.1. Recall that a *subdivision* of graph H is a graph H' obtained from H by replacing the edges of H with pairwise internally disjoint paths of non-zero length between the corresponding ends. Let $r \in \mathbb{N} \cup \{0\}$. An *$(\leq r)$-subdivision* of H is a subdivision of H in which the path replacing each edge has length at most $r + 1$.

Theorem 2.2 (Dvořák, see Theorem 6 in [8]; Lozin and Razgon, see Theorem 3 in [12])**.** *For every graph H and all $d \in \mathbb{N} \cup \{0\}$ and $t \in \mathbb{N}$, there is a constant $m = m(H, d, t) \in \mathbb{N}$ with the following property. Let G be a graph with no induced subgraph isomorphic to a subdivision of H. Assume that G contains a $(\leq d)$-subdivision of K_m as a subgraph. Then G contains either K_t or $K_{t,t}$.*

3. Bundles and constellations

In this section we state our main result, Theorem 3.2. We start with a few definitions that first appeared in [4].

Let G be a graph and let $l \in \mathbb{N}$. By an *l-polypath in G* we mean a set $\mathcal{L}$ of l pairwise disjoint paths in G. We say $\mathcal{L}$ is *plain* if every two distinct paths $L, L' \in \mathcal{L}$ are anticomplete in G. Also, two polypaths $\mathcal{L}$ and $\mathcal{L}'$ in G are said to be *disentangled* if $V(\mathcal{L}) \cap V(\mathcal{L}') = \varnothing$.

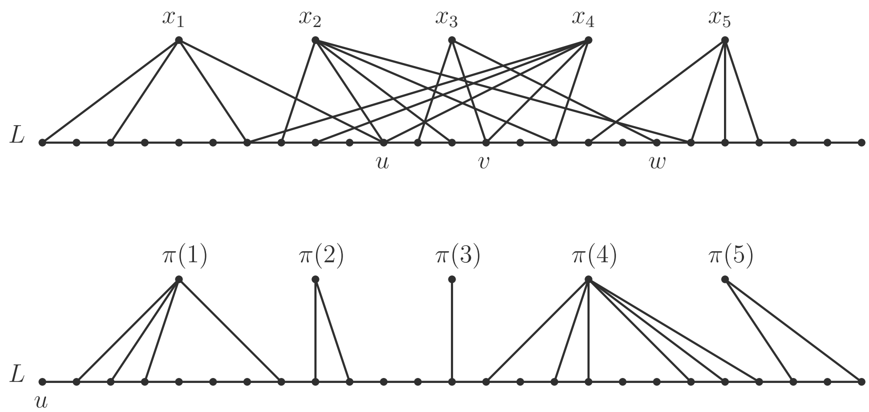

FIGURE 4. Top: a $(5,1)$-constellation $\mathfrak{c}$ with $S_\mathfrak{c} = \{x_1, x_2, x_3, x_4, x_5\}$ and $L_\mathfrak{c} = L$. Note that $\mathfrak{c}$ is 3-meager (with u being the only vertex in L with three neighbors in $S_\mathfrak{c}$) and 6-hollow (with the x_3-gap v-L-w of length five being the longest). Bottom: a 5-alignment.

For $s \in \mathbb{N}$, an (s,l)-*bundle in* G is a pair $\mathfrak{b} = (S_\mathfrak{b}, \mathcal{L}_\mathfrak{b})$ where $S_\mathfrak{b} \subseteq V(G)$ with $|S_\mathfrak{b}| = s$ and $\mathcal{L}_\mathfrak{b}$ is an l-polypath in G (note that $S_\mathfrak{b}$ and $V(\mathcal{L}_\mathfrak{b})$ are not necessarily disjoint). If $l = 1$, say $\mathcal{L}_\mathfrak{b} = \{L_\mathfrak{b}\}$, we also denote the $(s,1)$-bundle $\mathfrak{b}$ by the pair $(S_\mathfrak{b}, L_\mathfrak{b})$. Given an (s,l)-bundle $\mathfrak{b}$ in G, we write $V(\mathfrak{b}) = S_\mathfrak{b} \cup V(\mathcal{L}_\mathfrak{b})$, and for every $L \in \mathcal{L}_\mathfrak{b}$, we denote by $\mathfrak{b}_L$ the $(s,1)$-bundle $(S_\mathfrak{b}, L)$. Also, we say that $\mathfrak{b}$ is *plain* if the l-polypath $\mathcal{L}_\mathfrak{b}$ is plain. For two bundles $\mathfrak{b}$ and $\mathfrak{b}'$ in a graph G, we say $\mathfrak{b}$ and $\mathfrak{b}'$ are *disentangled* if $V(\mathfrak{b}) \cap V(\mathfrak{b}') = \varnothing$.

An (s,l)-*constellation in* G is an (s,l)-bundle $\mathfrak{c} = (S_\mathfrak{c}, \mathcal{L}_\mathfrak{c})$ in G such that $S_\mathfrak{c}$ is a stable set (of cardinality s) in $G \setminus V(\mathcal{L}_\mathfrak{c})$, and every $s \in S_\mathfrak{c}$ has a neighbor in every path $L \in \mathcal{L}_\mathfrak{c}$.

Let $\mathfrak{c}$ be an $(s,1)$-constellation in a graph G. For a vertex $x \in S_\mathfrak{c}$, by an x-*gap in* $\mathfrak{c}$ we mean a path P in $L_\mathfrak{c}$ (possibly of length zero) where x is adjacent to the ends of P and anticomplete to P^*. For $d \in \mathbb{N}$, we say $\mathfrak{c}$ is d-*hollow* if for every $x \in S_\mathfrak{c}$, every x-gap in $\mathfrak{c}$ has length less than d. Also, we say $\mathfrak{c}$ is d-*meager* if every vertex in $L_\mathfrak{c}$ is adjacent to at most d vertices in $S_\mathfrak{c}$ (see Figure 4). In general, for an (s,l)-constellation $\mathfrak{c}$, we say $\mathfrak{c}$ is d-*hollow* (d-*meager*) if for every $L \in \mathcal{L}_\mathfrak{c}$, $\mathfrak{c}_L$ is d-hollow (d-meager).

For $s \in \mathbb{N}$, an s-*alignment* in G is a triple (S, L, π) where (S, L) is an $(s,1)$-constellation in G and $\pi : \mathbb{N}_s \to S$ is a bijection such that for some end u of L, the following holds.

(AL) For all $i, j \in \mathbb{N}_s$ with $i < j$, every neighbor $v_i \in L$ of $\pi(i)$ and every neighbor $v_j \in L$ of $\pi(j)$, the path in L from u to v_j contains v_i in its interior. In other words, traversing L starting at u, all neighbors of $\pi(i)$ appear before all neighbors of $\pi(j)$.

See Figure 4. In particular, (S, L) is 1-meager.

For $h, s \in \mathbb{N}$, an (s,h)-*array* in a graph G is a plain, h-hollow (s,s)-constellation $\mathfrak{a}$ in G which satisfies the following.

(AR) There exists a bijection $\pi : \mathbb{N}_s \to S_\mathfrak{a}$ such that for every $L \in \mathcal{L}_\mathfrak{a}$, $(S_\mathfrak{a}, L, \pi)$ is an s-alignment in G.

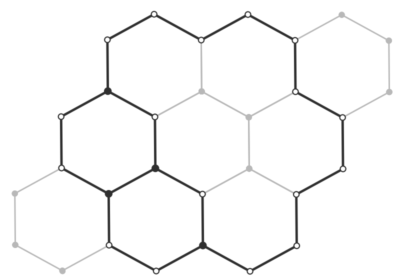

FIGURE 5. An induced subgraph of $W_{4\times 4}$ isomorphic to a subdivision of K_4 with exactly one unsubdivided edge.

See Figure 6. It is readily observed that for $X \subseteq V(G)$, if there is an $(s,1)$-array $\mathfrak{a}$ in G with $V(\mathfrak{a}) = X$, then $G[X]$ contains an expansion of the graph PD_s, and if $G[X]$ is isomorphic to an expansion of PD_s for some $X \subseteq V(G)$, then there is an $(s,1)$-array $\mathfrak{a}$ in G with $V(\mathfrak{a}) = X$. Moreover, we have:

Theorem 3.1. *Let $h, s \in \mathbb{N}$ and let G be a graph. Let $\mathfrak{a}$ be an (s,h)-array in G. Then $G[V(\mathfrak{a})]$ is a 4-clean, $(3,h)$-pinched graph of treewidth at least s.*

Proof. Let $J = G[V(\mathfrak{a})]$. Note that J contains $K_{s,s}$ as a minor (by contracting each path $L \in \mathcal{L}_{\mathfrak{a}}$ into a vertex), which implies that $\mathrm{tw}(J) \geq s$. For the rest of the proof, let $\pi : \mathbb{N}_s \to S_{\mathfrak{a}}$ be the bijection satisfying (AR).

Assume that J has an induced subgraph isomorphic to a subdivision of $W_{4\times 4}$ or the line graph of a subdivision of $W_{4\times 4}$. Note that $W_{4\times 4}$ has an induced subgraph isomorphic to a subdivision of K_4 with exactly one unsubdivided edge (see Figure 5). Thus, J has an induced subgraph W that is isomorphic to either a subdivision of K_4 with at most one unsubdivided edge or the line graph of a subdivision of K_4 with at most one unsubdivided edge. In either case, it is straightforward to observe that for every vertex $w \in V(W)$, the graph $W \setminus N_W[w]$ is connected. On the other hand, W is an induced subgraph of J with $\mathrm{tw}(W) \geq 3$ (because W contains K_4 as a minor). Consequently, we have $|S_{\mathfrak{a}} \cap V(W)| \geq 3$; say $w_1, w_2, w_3 \in S_{\mathfrak{a}} \cap V(W)$ such that $\pi(w_1) < \pi(w_2) < \pi(w_3)$. But now $W \setminus N_W[w_2]$ is disconnected, a contradiction. We deduce that J has no induced subgraph isomorphic to a subdivision of $W_{4\times 4}$ or the line graph of a subdivision of $W_{4\times 4}$. Moreover, it is easy to check that J is K_4-free and $K_{2,3}$-free. Hence, J is 4-clean.

It remains to show that J is $(3,h)$-pinched. Suppose for a contradiction that there are three cycles C_1, C_2, C_3 in J of length at least $h+2$ with $C_1 \cap C_2 \cap C_3 = \{x\}$ and $C_1 \setminus \{x\}, C_2 \setminus \{x\}, C_3 \setminus \{x\}$ are pairwise disjoint and anticomplete. Since x has degree at least six in J, it follows that $x \in S_{\mathfrak{a}}$. Also, since $\mathfrak{a}$ is h-hollow, it follows that there is no cycle of length at least $h+2$ in $J[V(\mathcal{L}_{\mathfrak{a}} \cup \{x\})]$. Thus, for every $l \in \{1,2,3\}$, we may pick $x_l \in C_l \cap (S_{\mathfrak{a}} \setminus \{x\}) \neq \varnothing$. By symmetry, we may assume that there are distinct $l_1, l_2 \in \{1,2,3\}$ for which $\pi^{-1}(x)$ is smaller than $\pi^{-1}(x_{l_1})$ and $\pi^{-1}(x_{l_2})$. In particular, we may assume without loss of generality that $\pi^{-1}(x) < \pi^{-1}(x_1) < \pi^{-1}(x_2)$. But then $x, x_2 \in C_2$ are in different components of $J \setminus N_J[x_1]$ (see Figure 6), which violates the fact that $C_1 \setminus \{x\}$ and $C_2 \setminus \{x\}$ are disjoint and anticomplete. This completes the proof of Theorem 3.1. ■

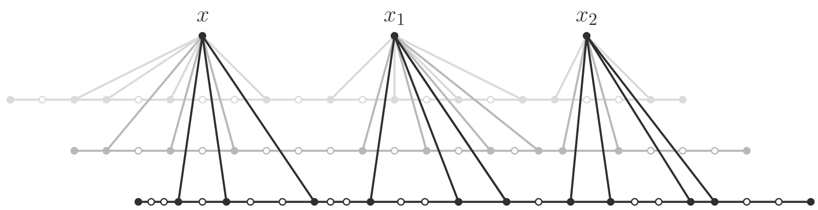

FIGURE 6. A $(3,4)$-array (vertex labels x, x_1 and x_2 are relevant in the proof of Theorem 3.1).

Now we can state the main result of this paper:

Theorem 3.2. *For all $c, h, s, t \in \mathbb{N}$, there is a constant $\tau = \tau(c, h, s, t) \in \mathbb{N}$ such that for every t-clean (c, h)-pinched graph G of treewidth more than τ, there is an (s, h)-array in G.*

In view of Theorem 3.1, Theorem 3.2 yields, for all $c, h \in \mathbb{N}$, a complete description of unavoidable induced subgraphs of (c, h)-pinched graphs with large treewidth. Moreover, as mentioned above, there is an $(s, 1)$-array in a graph G if and only if G contains an expansion of PD_s. So Theorem 1.5 is the special case of Theorem 3.2 for $h = 1$. We also point out that for $s \geq 4$ and $h \in \mathbb{N}$, if a graph G is $(2, h)$-pinched, then there is no (s, h)-array in G. Hence, Theorem 1.5 implies a strengthening of Theorem 1.4, that for every $h \in \mathbb{N}$, the class of all $(2, h)$-pinched graphs is clean.

4. Dealing with plain constellations

In this section, we prove the following:

Theorem 4.1. *For all $c, h, s, t \in \mathbb{N}$, there are constants $\sigma = \sigma(c, h, s, t) \in \mathbb{N}$ and $\lambda = \lambda(c, h, s, t) \in \mathbb{N}$ with the following properties. Let G be a (c, h)-pinched graph which does not K_t and $K_{t,t}$. Assume that there exists a plain (σ, λ)-constellation $\mathfrak{c}$ in G. Then there is an (s, h)-array in G.*

We need a couple of lemmas, beginning with the following. Although we have proved a similar result in [4], we include the proof here as it is short.

Lemma 4.2. *Let $a, d, s, l \in \mathbb{N}$ and let G be a graph. Assume that there exists a d-meager $(a^{l-1}(s + d(l - 1)), 1)$-constellation (S_0, L_0) in G. Then one of the following holds.*

(a) *There exists an a-alignment (S, L, π) in G with $S \subseteq S_0$ and $L \subseteq L_0$.*
(b) *There exists a plain (s, l)-constellation $\mathfrak{c}$ in G such that $S_{\mathfrak{c}} \subseteq S_0$ and $L \subseteq L_0$ for every $L \in \mathcal{L}_{\mathfrak{c}}$. In particular, $\mathfrak{c}$ is also d-meager.*

Proof. For fixed a, d and s, we proceed by induction on l. Note that if $l = 1$, then (S_0, L_0) is an $(s, 1)$-constellation in G satisfying Lemma 4.2(b). Thus, we may assume that $l \geq 2$.

Let u, v be the ends of L_0. For every vertex $x \in S_0$, traversing L_0 from u to v, let u_x and v_x be the first and the last neighbor of x in L_0 (possibly $u_x = v_x$), and let $L_x = u_x\text{-}L_0\text{-}v_x$. Let Y be a subset of S_0 for which the paths $\{L_y : y \in Y\}$ are pairwise disjoint, such that $|Y|$ is as large as possible, and subject to this, $\sum_{y \in Y} |u\text{-}L_0\text{-}u_y|$ is as large as possible.

Clearly, if $|Y| \geq a$, then Lemma 4.2(a) holds. Therefore, we may assume that $|Y| < a$. Let $W = \{u_y : y \in Y\}$; then we have $|W| < a$. We claim that:

(1) *For every $x \in S_0$, we have $L_x \cap W \neq \varnothing$.*

Suppose for a contradiction that for some $x \in S_0$, we have $L_x \cap W = \varnothing$. Then $x \notin Y$. By the maximality of Y, there exists $y \in Y$ such that $L_x \cap L_y \neq \varnothing$. Since $L_x \cap W = \varnothing$, it follows that $u_y \notin L_x$ (and so $u_x \in L_y$), and $L_x \cap L_{y'} = \varnothing$ for all $y' \in Y \setminus \{y\}$. In particular, we have $|u\text{-}L_0\text{-}u_x| > |u\text{-}L_0\text{-}u_y|$, and the paths $\{L_{y'} : y' \in (Y \setminus \{y\}) \cup \{x\}\}$ are pairwise disjoint. But now $(Y \setminus \{y\}) \cup \{x\}$ is a better choice than Y, a contradiction. This proves (1).

Since $|S_0| = a^{l-1}(s + (l-1)d)$ and $|W| < a$, it follows from (1) that there exists $A \subseteq S_0$ with $|A| = a^{l-2}(s+d(l-1))$ and a vertex $w \in W$, such that for every $x \in A$, we have $w \in L_x$. On the other hand, since (S_0, L_0) is d-meager, there are at most d vertices in A which are adjacent to w, and so

$$|A \setminus N_A(w)| \geq a^{l-2}(s + d(l-1)) - d \geq a^{l-2}(s + d(l-2)) > 0.$$

It follows that $L_0 \setminus \{w\}$ has two components, say L_1 and L_2, and there exists $B \subseteq A$ with $|B| = a^{l-2}(s + d(l-2))$ such that every vertex in B has a neighbor in L_1 and a neighbor in L_2. It follows that (B, L_1) and (B, L_2) are both d-meager $(a^{l-2}(s+d(l-2)), 1)$-constellations in G. From the induction hypothesis applied to (B, L_1), we deduce that either there exists an a-alignment (S, L) in G with $S \subseteq B \subseteq S_0$ and $L \subseteq L_1 \subseteq L_0$, or there exists a plain $(s, l-1)$-constellation $\mathfrak{c}_1$ in G such that $S_{\mathfrak{c}_1} \subseteq B \subseteq S_0$ and $L \subseteq L_1 \subseteq L_0$ for every $L \in \mathcal{L}_{\mathfrak{c}_1}$. In the former case, Lemma 4.2(a) holds, as required. In the latter case, $\mathfrak{c} = (S_{\mathfrak{c}_1}, \mathcal{L}_{\mathfrak{c}_1} \cup \{L_2\})$ is a plain (s, l)-constellation in G such that $S_{\mathfrak{c}} = S_{\mathfrak{c}_1} \subseteq S_0$ and $L \subseteq L_1 \cup L_2 \subseteq L_0$ for every $L \in \mathcal{L}_{\mathfrak{c}}$, and so Lemma 4.2(b) holds. This completes the proof of Lemma 4.2. ■

Indeed, for pinched graphs, it turns out that Lemma 4.2(a) is the only possible outcome:

Lemma 4.3. *Let $a, c, d, h \in \mathbb{N}$ and let G be a (c, h)-pinched graph. Assume that there exists a d-meager $(a^{2cdh-1}d(h + 2cdh - 1), 1)$-constellation (S_0, L_0) in G. Then there exists an a-alignment (S, L, π) in G with $S \subseteq S_0$ and $L \subseteq L_0$.*

Proof. Suppose not. Then applying Lemma 4.2 to (S_0, L_0), it follows that there exists a plain $(dh, 2cdh)$-constellation $\mathfrak{c}$ in G such that $S_{\mathfrak{c}} \subseteq S_0$ and $L \subseteq L_0$ for every $L \in \mathcal{L}_{\mathfrak{c}}$. In particular, $\mathfrak{c}$ is d-meager. Let u be an end of L_0. Let $\mathcal{L}_{\mathfrak{c}} = \{L_i : i \in \mathbb{N}_{2cdh}\}$ and let u_i, v_i be the ends of L_i for each $i \in \mathbb{N}_{2cdh}$, such that traversing L_0 starting at u, the vertices $u_1, v_1, \ldots, u_{2cdh}, v_{2cdh}$ appear on L_0 in this order. For each $i \in \mathbb{N}_{2cdh-1}$, let v_i' be the neighbor of v_i in $L_0 \setminus L_i$. Since $\mathfrak{c}$ is plain, it follows that $v_i\text{-}L_0\text{-}u_{i+1}$ is a path of length at least two in L_0 whose interior is disjoint from $\{u_i : i \in \mathbb{N}_{2cdh}\} \cup \{v_i : i \in \mathbb{N}_{2cdh}\}$ and contains v_i'.

For every $i \in \mathbb{N}_{2cdh}$, since $\mathfrak{c}$ is a constellation, every vertex in $S_{\mathfrak{c}}$ has a neighbor in L_i. Let R_i be the shortest path in L_i containing v_i such that every vertex in $S_{\mathfrak{c}}$ has a neighbor in R_i. Since $\mathfrak{c}$ is d-meager and $|S_{\mathfrak{c}}| = dh$, it follows that $|R_i| \geq h$ for all $i \in \mathbb{N}_{2cdh}$. Let w_i be the end of R_i distinct from v_i. Then the minimality of R_i implies that there exists a vertex $x_i \in S_{\mathfrak{c}}$ that is adjacent to w_i and anticomplete to $R_i \setminus \{w_i\}$. Since $|S_{\mathfrak{c}}| = dh$, it follows that there exist $x \in S_0$ as well as two disjoint c-subsets $\{i_k : k \in \mathbb{N}_c\}$ and $\{j_k : k \in \mathbb{N}_c\}$ of $\mathbb{N}_{2cdh}$, such that $i_1 < j_1 < \cdots < i_c < j_c$ and we have $x_{i_k} = x_{j_k} = x$ for all $k \in \mathbb{N}_c$.

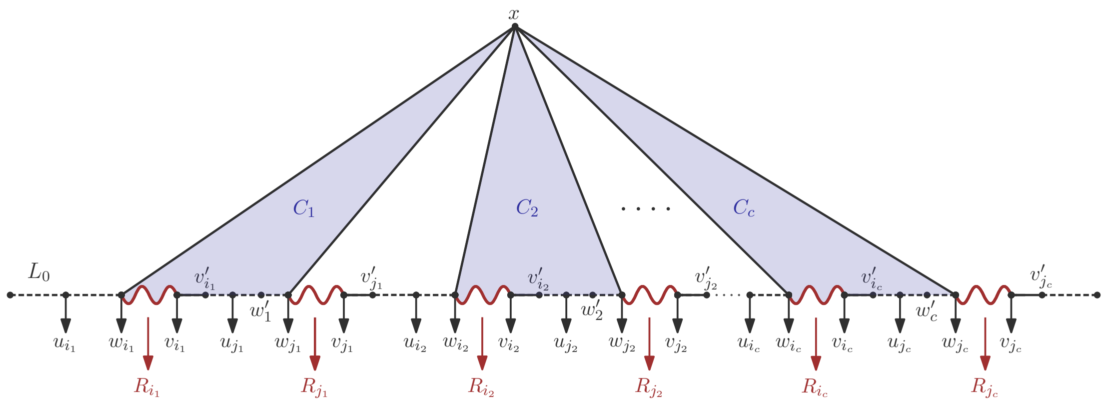

Figure 7. Proof of Lemma 4.3 (dashed lines depict paths of arbitrary length, and highlighted paths have length at least h).

For each $k \in \mathbb{N}_c$, traversing v'_{i_k}-L_0-w_{j_k} starting at v'_{i_k}, let w'_k be the first neighbor of x in v'_{i_k}-L_0-w_{j_k}. Since x is adjacent to w_{j_k}, it follows that w'_k exists (See Figure 7). Let $C_k = x$-w_{i_k}-L_0-w'_k-x. Then C_k is a cycle of length at least $h+2$ in G, as $R_k \cup \{v'_{i_k}\} \subseteq C_k$ and $|R_k| \geq h$. Now $C_1, \ldots, C_c$ are c cycles of length at least $h+2$ in G with $C_1 \cap \cdots \cap C_c = \{x\}$. Also, since $C_1 \setminus \{x\}, \ldots, C_c \setminus \{x\} \subseteq L$ are contained in pairwise distinct component of $L_0 \setminus \{v'_{j_k} : k \in \mathbb{N}_c\}$, it follows that $C_1 \setminus \{x\}, \ldots, C_c \setminus \{x\} \subseteq L$ are pairwise disjoint and anticomplete in G. This violates the assumption that G is (c,h)-pinched, hence completing the proof of Lemma 4.3. ■

We also need the following quantified version of Ramsey's Theorem. This has appeared in several references; see, for instance, [9].

Theorem 4.4 (Ramsey [14], see also [9])**.** *For all $c, s \in \mathbb{N}$, every graph G on at least c^s vertices contains either a clique of cardinality c or a stable set of cardinality s.*

From Theorem 4.4, we deduce that:

Lemma 4.5. *Let $l, s, t \in \mathbb{N}$, let G be a graph and let $\mathfrak{c}$ be a $(s, l+(st)^t)$-constellation in G. Then one of the following holds.*

(a) *G contains K_t or $K_{t,t}$.*
(b) *There exists $\mathcal{L} \subseteq \mathcal{L}_\mathfrak{c}$ with $|\mathcal{L}| = l$ such that $(S_\mathfrak{c}, \mathcal{L})$ is a t-meager (s,l)-constellation in G.*

Proof. Suppose that Lemma 4.5(a) does not hold. For every $L \in \mathcal{L}_\mathfrak{c}$, let t_L be the largest number in $\mathbb{N}_s$ for which some vertex $u_L \in L$ has at least t_L neighbors in $S_\mathfrak{c}$. It follows that:

(2) *We have $|\{L \in \mathcal{L}_\mathfrak{c} : t_L \geq t\}| < (st)^t$.*

Suppose not. Let $\mathcal{S} \subseteq \{L \in \mathcal{L}_\mathfrak{c} : t_L \geq t\}$ with $|\mathcal{S}| = (st)^t$. Then for every $L \in \mathcal{S}$, we may choose $u_L \in L$ and $T_L \subseteq S_\mathfrak{c}$ such that $|T_L| = t$ and u_L is complete to T_L in G. Since $|S_\mathfrak{c}| = s$, it follows that there exist $T \subseteq S_\mathfrak{c}$ and $\mathcal{T} \subseteq \mathcal{S}$ with $|T| = t$ and $|\mathcal{T}| = t^t$, such that for every $L \in \mathcal{T}$, we have $T_L = T$. Let $U = \{u_L : L \in \mathcal{T}\}$. Then T and U are disjoint and complete in

G. Also, since G is K_t-free and $|U| = t^t$, it follows from Theorem 4.4 that there is a stable set $U' \subseteq U$ in G with $|U'| = t$. But then $G[T \cup U']$ is isomorphic to $K_{t,t}$, a contradiction. This proves (2).

Now the result is immediate from (2) and the fact that $|\mathcal{L}_\mathfrak{c}| = l + t^t s^t$. This completes the proof of Lemma 4.5. ■

We are now ready to prove Theorem 4.1.

Proof of Theorem 4.1. We claim that:

$$\sigma = \sigma(c, h, s, t) = s^{2cht-1} t(h + 2cht - 1);$$

$$\lambda = \lambda(c, h, s, t) = cs(s!)\sigma^s + (\sigma t)^t$$

satisfy Theorem 4.1. Let G be a t-clean (c, h)-pinched graph which does not contain K_t and $K_{t,t}$. Let $\mathfrak{c}$ be a plain (σ, λ)-constellation in G. By the choice of σ and λ, we may apply Lemma 4.5 to $\mathfrak{c}$, and deduce that Lemma 4.5(b) holds, that is, there exists $\mathcal{L}_0 \subseteq \mathcal{L}_\mathfrak{c}$ with $|\mathcal{L}_0| = cs(s!)\sigma^s$ such that for every $L \in \mathcal{L}_0$, $(S_\mathfrak{c}, L)$ is a t-meager, plain $(\sigma, 1)$-constellation in G. In particular, since $\sigma = s^{2cht-1}t(h + 2cht - 1)$, we may apply Lemma 4.3 to $(S_\mathfrak{c}, L)$ to show that:

(3) *For every $L \in \mathcal{L}_0$, there exists an s-alignment (S_L, Q_L, π_L) in G with $S_L \subseteq S_\mathfrak{c}$ and $Q_L \subseteq L$.*

Note that $|S_L| = s$ for all $L \in \mathcal{L}_0$. Recall also that $|S_\mathfrak{c}| = \sigma$ and $|\mathcal{L}_0| = cs(s!)\sigma^s$. This, along with a pigeon-hole argument, implies immediately that:

(4) *There exists $S \subseteq S_\mathfrak{c}$ with $|S| = s$, and $\mathcal{L} \subseteq \mathcal{L}_0$ with $|\mathcal{L}| = cs$ and $\pi : \mathbb{N}_s \to S$ such that for every $L \in \mathcal{L}$, (S, Q_L, π) is an s-alignment in G.*

We further deduce that:

(5) *There exists $\mathcal{S} \subseteq \mathcal{L}$ with $|\mathcal{S}| = s$ such that for every $L \in \mathcal{S}$, the $(s, 1)$-constellation (S, Q_L) is h-hollow.*

To see this, since $|S| = s$ and $|\mathcal{L}| = cs$, it suffices to show that for every $x \in S$, there are fewer than c paths $L \in \mathcal{L}$ for which there is an x-gap in (S, Q_L) of length at least h. Suppose for a contradiction that for some $x \in S$, there are c distinct paths $L_1, \dots, L_c \in \mathcal{L}$ such that for every $i \in \mathbb{N}_c$, there is an x-gap P_i in (S, Q_{L_i}) of length at least h; let y_i, z_i be the ends of P_i. It follows that $C_i = x\text{-}y_i\text{-}P_i\text{-}z_i\text{-}x$ is a cycle of length at least $h + 2$ in G. Now $C_1, \dots, C_c$ are c cycles of length at least $h+2$ in G with $C_1 \cap \dots \cap C_c = \{x\}$. Moreover, since $\mathfrak{c}$ is plain, $C_1 \setminus \{x\}, \dots, C_c \setminus \{x\} \subseteq V(\mathcal{L})$ are pairwise disjoint and anticomplete in G. This violates the assumption that G is (c, h)-pinched, and so proves (5).

Let $\mathfrak{a} = (S, \mathcal{S})$ where S comes from (4) and $\mathcal{S}$ comes from (5). Also, let π be as in (4). It follows from (4) and (5) that $\mathfrak{a}, \pi$ satisfies (AR). Hence, $\mathfrak{a}$ is an (s, h)-array in G. This completes the proof of Theorem 4.1. ■

5. Obtaining a plain constellation

This section contains the main ingredient of the proof of Theorem 3.2:

Theorem 5.1. *For all $c, h, l, s, t \in \mathbb{N}$, there is a constant $\Omega = \Omega(c, h, l, s, t) \in \mathbb{N}$ with the following property. Let G be a (c, h)-pinched graph. Assume that there exists a strong Ω-block in G. Then one of the following holds.*

(a) *G contains either K_t or $K_{t,t}$.*
(b) *There exists a plain (s, l)-constellation in G.*

Our road to the proof of Theorem 5.1 passes through a number of definitions and lemmas from [4], beginning with two useful Ramsey-type results for polypaths:

Lemma 5.2 (Alecu, Chudnovsky, Hajebi, Spirkl; see Lemma 5.2 in [4])**.** *For all $b, f, g, m, n, t \in \mathbb{N}$ and $s \in \mathbb{N} \cup \{0\}$, there is a constant $\beta = \beta(b, f, g, m, n, s, t) \in \mathbb{N}$ with the following property. Let G be a graph and let $\mathfrak{B}$ be a collection of β pairwise disentangled $(b, 2b(g-1) + f)$-bundles in G. Then one of the following holds.*

(a) *G contains either K_t or $K_{t,t}$.*
(b) *There exists $\mathfrak{N} \subseteq \mathfrak{B}$ with $|\mathfrak{N}| = n$ as well as $S \subseteq \bigcup_{\mathfrak{b} \in \mathfrak{B} \setminus \mathfrak{N}} S_{\mathfrak{b}}$ with $|S| = s$, such that for every $\mathfrak{b} \in \mathfrak{N}$, there exists $\mathcal{G}_{\mathfrak{b}} \subseteq \mathcal{L}_{\mathfrak{b}}$ with $|\mathcal{G}_{\mathfrak{b}}| = g$ for which $(S, \mathcal{G}_{\mathfrak{b}})$ is an (s, g)-constellation in G.*
(c) *There exists $\mathfrak{M} \subseteq \mathfrak{B}$ with $|\mathfrak{M}| = m$ as well as $\mathcal{F}_{\mathfrak{b}} \subseteq \mathcal{L}_{\mathfrak{b}}$ with $|\mathcal{F}_{\mathfrak{b}}| = f$ for each $\mathfrak{b} \in \mathfrak{M}$, such that for all distinct $\mathfrak{b}, \mathfrak{b}' \in \mathfrak{M}$, $S_{\mathfrak{b}}$ is anticomplete to $S_{\mathfrak{b}'} \cup V(\mathcal{F}_{\mathfrak{b}'})$ in G.*

Lemma 5.3 (Alecu, Chudnovsky, Hajebi, Spirkl; see Lemma 5.3 in [4])**.** *For all $a, g \in \mathbb{N}$, there is a constant $\varphi = \varphi(a, g) \in \mathbb{N}$ with the following property. Let G be a graph and let $\mathcal{F}_1, \ldots, \mathcal{F}_a$ be a collection of a pairwise disentangled φ-polypaths in G. Then for every $i \in \mathbb{N}_a$, there exists a g-polypath $\mathcal{G}_i \subseteq \mathcal{F}_i$, such that for all distinct $i, i' \in \mathbb{N}_a$, either $V(\mathcal{G}_i)$ is anticomplete to $V(\mathcal{G}_{i'})$ in G, or for every $L \in \mathcal{G}_i$ and every $L' \in \mathcal{G}_{i'}$, there is an edge in G with an end in L and an end in L'.*

We continue with a few definitions from [4]. Let G be a graph, let $d \in \mathbb{N} \cup \{0\}$ and let $r \in \mathbb{N}$. For $X \subseteq V(G)$, by a *(d, r)-patch for X in G* we mean a $(1, r)$-bundle $\mathfrak{p}$ in G where:

(P1) $S_{\mathfrak{p}} \subseteq V(G) \setminus V(\mathcal{L}_{\mathfrak{p}})$;
(P2) every path $L \in \mathcal{L}_{\mathfrak{p}}$ has length at least d; and
(P3) for every $L \in \mathcal{L}_{\mathfrak{p}}$, one may write $\partial L = \{x_L, y_L\}$ such that $L \cap X = \{x_L\}$ and $N_L(S_{\mathfrak{p}}) = \{y_L\}$.

Also, by a *(d, r)-match for X in G* we mean an r-polypath $\mathcal{M}$ in G such that

(M1) every path $L \in \mathcal{M}$ has length at least d; and
(M2) $V(\mathcal{M}) \cap X = \partial \mathcal{M}$.

See Figure 8. We also need:

Lemma 5.4 (Alecu, Chudnovsky, Hajebi, Spirkl; see Lemma 7.2 in [4])**.** *For all $d, l, m, r, r', s \in \mathbb{N}$, there is a constant $\eta = \eta(d, l, m, r, r', s) \in \mathbb{N}$ with the following property. Let G be a graph, let $X \subseteq V(G)$ and let $\mathfrak{p}$ be a (d, η)-patch for X in G. Then one of the following holds.*

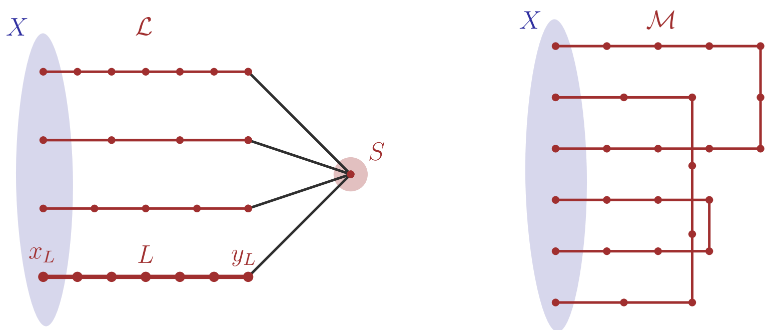

FIGURE 8. A $(3,6)$-patch for X (left) and a $(7,3)$-match for X (right).

(a) *There exists a plain (s,l)-constellation in G.*
(b) *G contains a $(\leq d+1)$-subdivision of K_m as a subgraph.*
(c) *There exists a plain $(2(d+1),r)$-match $\mathcal{M}$ for X in G such that $V(\mathcal{M}) \subseteq V(\mathfrak{p})$.*
(d) *There exists a plain (d,r')-patch $\mathfrak{q}$ for X in G such that $V(\mathfrak{q}) \subseteq V(\mathfrak{p})$.*

We can now prove the main result of this section:

Proof of Theorem 5.1. Let H be the unique graph (up to isomorphism) consisting of c induced cycles of length $h+2$, all sharing a common vertex and otherwise pairwise disjoint and anticomplete. Let $m = m(H, h-1, t)$ be as in Theorem 2.2 (note that m only depends on c, h, t).

Let $\gamma = (c(h+2))^{2chl-1}l(h+2chl-1)$; then $\gamma > 1$ (in fact, it holds that $\gamma \geq 6$). Let $\varphi = \varphi(c, s(\gamma+2l-2)^l)$ be as in Lemma 5.3. Let

$$\beta = \beta(1, \varphi, l, c, 1, s, t)$$

be as in Lemma 5.2, and let

$$\eta = \eta(h-1, l, m, c, 2(l-1)+\varphi, s)$$

be as in Lemma 5.4. We prove that

$$\Omega = \Omega(c, h, l, s, t) = \max\{m^{\beta+1}, \eta\}$$

satisfies Theorem 5.1.

Let G be a (c,h)-pinched graph such that there is a strong Ω-block in G. Then G does not contain a subdivision of H. Suppose for a contradiction that G contains neither K_t nor $K_{t,t}$, and there is no plain (s,l)-constellation in G. By Theorem 2.2 and the choice of m, G contains no subgraph isomorphic to a $(\leq h-1)$-subdivision of K_m. It is convenient to sum up all this in one statement:

(6) *The following hold.*

- *G does not contain K_t or $K_{t,t}$.*
- *There is no plain (s,l)-constellation in G.*
- *G does not contain a subgraph isomorphic to a $(\leq h-1)$-subdivision of K_m.*

Let B be a strong Ω-block in G and for every 2-subset $\{x,y\}$ of B, let $\mathcal{P}_{\{x,y\}}$ be the corresponding set of Ω paths in G from x to y. Let J be a graph with vertex set B such that two distinct vertices $x,y\in B$ are adjacent in J if and only if there exists $P_{\{x,y\}}\in\mathcal{P}_{\{x,y\}}$ of length at most h. Since $|V(J)|=|B|\geq m^{\beta+1}$, it follows from Theorem 4.4 that J contains either clique C on m vertices, or a stable set on $\beta+1$ vertices. In the former case, the union of the paths $P_{\{x,y\}}$ for all distinct $x,y\in C$ forms a subgraph of G isomorphic to a $(\leq h-1)$-subdivision of K_m, which violates the third bullet in (6). Therefore, J contains a stable set on $\beta+1$ vertices. This, along with the choice of Ω, implies that one may choose $\beta+1$ distinct vertices $x,y_1,\dots,y_\beta\in V(G)$ as well as, for each $i\in\mathbb{N}_\beta$, a collection $\mathcal{P}_i$ of η pairwise internally disjoint paths in G from x to y_i, such that:

- for every $i\in\mathbb{N}_\beta$, every path $P\in\mathcal{P}_i$ has length at least $h+1\geq 2$; and
- $V(\mathcal{P}_1)\setminus\{x\},\dots,V(\mathcal{P}_\beta)\setminus\{x\}$ are pairwise disjoint in G.

For each $i\in\mathbb{N}_\beta$, let $\mathcal{L}_i=\mathcal{P}_i^*$, and for every $L\in\mathcal{L}_i$, let x_L and y_L be the (unique) neighbors of x and y_i in L, respectively. Thus, we have $\partial L=\{x_L,y_L\}$ (note that x_L,y_L might be the same, but they are distinct from both x and y_i). Let $X_i=\{x_L : L\in\mathcal{L}_i\}$. It follows that $\mathcal{L}_i$ is an η-polypath in G, and

- $y_i\in V(G)\setminus V(\mathcal{L}_i)$;
- every path $L\in\mathcal{L}_i$ has length at least $h-1$; and
- $L\cap X_i=\{x_L\}$ and $N_L(y_i)=\{y_L\}$ for every $L\in\mathcal{L}_i$.

Therefore, by (P1), (P2) and (P3), for every $i\in\mathbb{N}_\beta$, the $(1,\eta)$-bundle $\mathfrak{p}_i=(\{y_i\},\mathcal{L}_i)$ is an $(h-1,\eta)$-patch for X_i in G. In addition, we show that:

(7) *For every $i\in\mathbb{N}_\beta$, there is a plain $(h-1,2(l-1)+\varphi)$-patch $\mathfrak{q}_i$ for X_i in G with $V(\mathfrak{q}_i)\subseteq V(\mathfrak{p}_i)\subseteq V(G)\setminus\{x\}$.*

By the choice of η, we can apply Lemma 5.4 to X_i and $\mathfrak{p}_i$. Note that Lemma 5.4(a) and Lemma 5.4(b) violate the second and the third bullets of (6), respectively. Assume that Lemma 5.4(c) holds. Then there is a plain $(2h,c)$-match $\mathcal{M}_i$ for X_i in G with $V(\mathcal{M}_i)\subseteq V(\mathfrak{p}_i)\subseteq V(G)\setminus\{y_i\}$. In particular, for every $M\in\mathcal{M}_i$, $C_M=M\cup\{x\}$ is a cycle of length at least $2h+2$ in G. But now $\{C_M : M\in\mathcal{M}_i\}$ is a collection of c cycles of length at least $2h+2$ in G where $\bigcap_{M\in\mathcal{M}_i}C_M=\{x\}$ and the sets $\{C_M\setminus\{x\} : M\in\mathcal{M}\}$ are pairwise disjoint and anticomplete in G, violating the assumption that G is (c,h)-pinched. So Lemma 5.4(d) holds. This proves (7).

For each $i\in\mathbb{N}_\beta$, let $\mathfrak{q}_i$ be as in (7) and let $S_{\mathfrak{q}_i}=\{z_i\}$ (note that z_i may or may not belong to $X_i\subseteq N_G(x)$). Now, $\mathfrak{B}=\{\mathfrak{q}_1,\dots,\mathfrak{q}_\beta\}$ is a collection of β pairwise disentangled plain $(1,2(l-1)+\varphi)$-bundles in G. Given the choice of β, we can apply Lemma 5.2 to $\mathfrak{B}$. Note that Lemma 5.2(a) directly violates the first bullet of (6). Also, if Lemma 5.2(b) holds, then there is an (s,l)-constellation $\mathfrak{c}$ in G with $\mathcal{L}_{\mathfrak{c}}\subseteq\mathcal{L}_{\mathfrak{q}_i}$ for some $i\in\mathbb{N}_\beta$; in particular, $\mathfrak{c}$ is plain. But this violates the second bullet of (6). It follows that Lemma 5.2(c) holds, that is, there exists $I \subseteq \mathbb{N}_\beta$ with $|I| = c$ as well as $\mathcal{F}_i \subseteq \mathcal{L}_{\mathfrak{q}_i}$ with $|\mathcal{F}_i| = \varphi$ for each $i \in I$, such that for all distinct $i, i' \in I$, z_i is anticomplete to $V(\mathcal{F}_{i'}) \cup \{z_{i'}\}$ in G. We further deduce that:

(8) *There are distinct $i, i' \in I$ for which there exist $s(\gamma + 2l - 2)^l$-polypaths $\mathcal{G}_i \subseteq \mathcal{F}_i$ and $\mathcal{G}_{i'} \subseteq \mathcal{F}_{i'}$ such that for every $L \in \mathcal{G}_i$ and every $L' \in \mathcal{G}_{i'}$, there is an edge in G with an end in L and an end in L'.*

Suppose not. Then by the choice of φ, we may apply Lemma 5.3 to the φ-polypaths $\{\mathcal{F}_i : i \in I\}$ and deduce that for every $i \in I$, there exists a $s(\gamma + 2l - 2)^l$-polypath $\mathcal{G}_i \subseteq \mathcal{F}_i$, such that for all distinct $i, i' \in \mathbb{N}_a$, the sets $V(\mathcal{G}_i)$ and $V(\mathcal{G}_{i'})$ are anticomplete in G. Since $s(\gamma+2l-2)^l \geq \gamma > 1$, it follows that for each $i \in I$, there is a cycle C_i in $G[V(\mathcal{G}_i) \cup \{x, z_i\}]$ of length at least $h + 2$ where $x \in C_i$. Also, the sets $\{V(\mathcal{G}_i) \cup \{z_i\} : i \in I\}$ are pairwise disjoint and anticomplete in G, which in turn implies that the sets $\{C_i \setminus \{x\} : i \in I\}$ are pairwise disjoint and anticomplete in G. This yields a contradiction with the assumption that G is (c, h)-pinched, and so proves (8).

Henceforth, let $i, i' \in I$ and $\mathcal{G}_i, \mathcal{G}_{i'}$ be as in (8). Since $|\mathcal{G}_i| = |\mathcal{G}_{i'}| = s(\gamma+2l-2)^l \geq \gamma+2l-2$, we may choose $\mathcal{G}' \subseteq \mathcal{G}_{i'}$ with $|\mathcal{G}'| = \gamma + 2l - 2$. Write $\mathcal{G} = \mathcal{G}_i$. It follows that both $\mathcal{G}$ and $\mathcal{G}'$ are plain and disentangled polypaths in G. For every path $L \in \mathcal{G}$, let us say L is *rigid* if there exists a vertex $x_L \in L$ as well as $\mathcal{G}'_L \subseteq \mathcal{G}'$ with $|\mathcal{G}'_L| = l$ such that x_L has a neighbor in every path in $\mathcal{G}'_L$. We claim that:

(9) *Not all paths in $\mathcal{G}$ are rigid, that is, there is a path $L_0 \in \mathcal{G}$ such that every vertex in L_0 has neighbors in more than $l - 1$ paths in $\mathcal{G}'$.*

Suppose not. For every $L \in \mathcal{G}$, let $x_L \in L$ and $\mathcal{G}'_L \subseteq \mathcal{G}'$ with $|\mathcal{G}'_L| = l$ be as in the definition of a rigid path. Since $|\mathcal{G}| = s(\gamma + 2l - 2)^l$ and $|\mathcal{G}'| = \gamma + 2l - 2$, it follows that there exists $\mathcal{S} \subseteq \mathcal{G}$ with $|\mathcal{S}| = s$ and $\mathcal{L} \subseteq \mathcal{G}'$ with $|\mathcal{L}| = l$ such that for every $L \in \mathcal{S}$, we have $\mathcal{G}'_L = \mathcal{L}$. Let $S = \{x_L : L \in \mathcal{S}\}$. Then every vertex in S has a neighbor in every path in $\mathcal{L}$. But now since $\mathcal{G}$ and $\mathcal{G}'$ are plain and disentangled, it follows that $(S, \mathcal{L})$ is a plain (s, l)-constellation in G, a contradiction with the second bullet of (6). This proves (9).

By (9), there exists a path $L_0 \in \mathcal{G}$ that is not rigid. Let $\mathcal{T}$ be the set of all paths $L' \in \mathcal{G}'$ for which some vertex in ∂L_0 has a neighbor in L' in G. Then by (9), we have $|\mathcal{T}| \leq 2l - 2$. Consequently, there are γ distinct paths $L'_1, \ldots, L'_\gamma \in \mathcal{G}'$ such that for every $i \in \mathbb{N}_\gamma$, the ends of L_0 are anticomplete to L'_i in G. For each $i \in \mathbb{N}_\gamma$, let x_i be the end of L'_i which is adjacent to x. By (8), traversing L'_i starting at x_i, we may choose $x'_i \in L'_i$ to be the first vertex in L'_i with a neighbor in L_0^*. Let $S_0 = \{x'_i : i \in \mathbb{N}_\gamma\}$. It follows that every vertex in S_0 has a neighbor in L_0 while S_0 is anticomplete to the ends of L_0, and so (S_0, L_0) is a $(\gamma, 1)$-constellation in G. Also, by (9), (S_0, L_0) is $(l-1)$-meager. This calls for an application of Lemma 4.3 to (S_0, L_0), which implies that there exists a $c(h + 2)$-alignment (S, L, π) in G with $S \subseteq S_0$ and $L \subseteq L_0$. In fact, since S_0 is anticomplete to ∂L_0, we may assume that $L \subseteq L_0^*$.

Let u be the end of L for which (S, L, π) satisfies (AL). For every $i \in \mathbb{N}_c$, choose $\varepsilon_i, \delta_i, \theta_i \in \mathbb{N}_\gamma$ such that

$$x'_{\varepsilon_i} = \pi(i(h + 2) - h - 1);$$
$$x'_{\delta_i} = \pi(i(h + 2) - 1);$$

$$x'_{\theta_i} = \pi(i(h+2));$$

Traversing L starting at u, let u_i be the last neighbor of x'_{ε_i} in L, let v_i be the first neighbor of x'_{δ_i} in L, and let w_i be the first neighbor of x'_{θ_i} in L. It follows that $u, u_1, v_1, w_1 \dots, u_c, v_c, w_c$ are pairwise distinct, appearing on L in this order. Also, for every $i \in \mathbb{N}_c$, $P_i = u_i$-L-v_i is a path of length at least h in G, and $P_1, \dots, P_c$ are contained in distinct components of $L \setminus \{w_i : i \in \mathbb{N}_c\}$. Thus, $P_1, \dots, P_c$ are pairwise disjoint and anticomplete in G. Now, for every $i \in \mathbb{N}_c$, consider the cycle

$$C_i = x\text{-}x_{\varepsilon_i}\text{-}L'_{\varepsilon_i}\text{-}x'_{\varepsilon_i}\text{-}u_i\text{-}P_i\text{-}v_i\text{-}x'_{\delta_i}\text{-}L'_{\delta_i}\text{-}x_{\delta_i}\text{-}x.$$

Then $C_1, \dots, C_c$ are c cycles in G each of length at least $h+4$, all going through x and otherwise pairwise disjoint and anticomplete. This contradicts the assumption that G is (c, h)-pinched, hence completing the proof of Theorem 5.1. ■

6. The end

We conclude the paper with the proof of our main result, which we restate:

Theorem 3.2. *For all $c, h, s, t \in \mathbb{N}$, there is a constant $\tau = \tau(c, h, s, t) \in \mathbb{N}$ such that for every t-clean (c, h)-pinched graph G of treewidth more than τ, there is an (s, h)-array in G.*

Proof. Let $\sigma = \sigma(c, h, s, t)$ and $\lambda = \lambda(c, h, s, t)$ be as in Theorem 4.1. Let $\Omega = \Omega(c, h, \lambda, \sigma, t)$ be as in Theorem 5.1. We define $\tau(c, h, s, t) = \xi(\Omega, t)$, where $\xi(\cdot, \cdot)$ comes from Theorem 2.1. Let G be a t-clean (c, h)-pinched graph of treewidth more than τ. By Theorem 2.1, G contains a strong Ω-block. Therefore, since G does not contain K_t and $K_{t,t}$, it follows from Theorem 5.1 that there exists a plain (σ, λ)-constellation in G. But now by Theorem 4.1, there is an (s, h)-array in G, as desired. ■

7. Acknowledgments

Our thanks to the anonymous referees for suggesting several improvements.

References

[1] P. Aboulker, I. Adler, E. J. Kim, N. L. D. Sintiari, and N. Trotignon. On the tree-width of even-hole-free graphs. *European J. Combin.*, 98:Paper No. 103394, 21, 2021.

[2] T. Abrishami, B. Alecu, M. Chudnovsky, S. Hajebi, and S. Spirkl. Induced subgraphs and tree decompositions VII. Basic obstructions in H-free graphs. *J. Combin. Theory Ser. B*, 164:443–472, 2024.

[3] T. Abrishami, M. Chudnovsky, C. Dibek, S. Hajebi, P. Rzążewski, S. Spirkl, and K. Vušković. Induced subgraphs and tree decompositions II. Toward walls and their line graphs in graphs of bounded degree. *J. Combin. Theory Ser. B*, 164:371–403, 2024.

[4] B. Alecu, M. Chudnovsky, S. Hajebi, and S. Spirkl. Induced subgraphs and tree decompositions IX. Grid theorem for perforated graphs. Manuscript available at https://arxiv.org/abs/2305.15615, 2023.

[5] H. L. Bodlaender. Dynamic programming on graphs with bounded treewidth. In *Automata, languages and programming (Tampere, 1988)*, volume 317 of *Lecture Notes in Comput. Sci.*, pages 105–118. Springer, Berlin, 1988.

[6] M. Bonamy, E. Bonnet, H. Déprés, L. Esperet, C. Geniet, C. Hilaire, S. Thomassé, and A. Wesolek. Sparse graphs with bounded induced cycle packing number have logarithmic treewidth. *J. Combin. Theory Ser. B*, 167:215–249, 2024.

[7] J. Davies. Appeared in an Oberwolfach report. https://ems.press/journals/owr/articles/9790352 (2022).

[8] Z. Dvořák. Induced subdivisions and bounded expansion. *European J. Combin.*, 69:143–148, 2018.

[9] S. Hajebi, Y. Li, and S. Spirkl. Hitting all maximum stable sets in P_5-free graphs. *J. Combin. Theory Ser. B*, 165:142–163, 2024.

[10] R. Hickingbotham, F. Illingworth, B. Mohar, and D. R. Wood. Treewidth, circle graphs, and circular drawings. *SIAM J. Discrete Math.*, 38(1):965–987, 2024.

[11] T. Korhonen. Grid induced minor theorem for graphs of small degree. *J. Combin. Theory Ser. B*, 160:206–214, 2023.

[12] V. Lozin and I. Razgon. Tree-width dichotomy. *European J. Combin.*, 103:Paper No. 103517, 8, 2022.

[13] C. Pohoata. Unavoidable induced subgraphs of large graphs. Senior thesis, Princeton University, 2014.

[14] F. P. Ramsey. On a Problem of Formal Logic. *Proc. London Math. Soc. (2)*, 30(4):264–286, 1929.

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

[16] N. L. D. Sintiari and N. Trotignon. (Theta, triangle)-free and (even hole, K_4)-free graphs—part 1: Layered wheels. *J. Graph Theory*, 97(4):475–509, 2021.