TREE INDEPENDENCE NUMBER
III. THETAS, PRISMS AND STARS

MARIA CHUDNOVSKY[§], SEPEHR HAJEBI[†,‡], AND NICOLAS TROTIGNON[⋆]

ABSTRACT. We prove that for every $t \in \mathbb{N}$ there exists $\tau = \tau(t) \in \mathbb{N}$ such that every (theta, prism, $K_{1,t}$)-free graph has tree independence number at most τ (where we allow "prisms" to have one path of length zero).

1. INTRODUCTION

Graphs in this paper have finite and non-empty vertex sets, no loops and no parallel edges. The set of all positive integers is denoted by $\mathbb{N}$, and for every $n \in \mathbb{N}$, we write $[n]$ for the set of all positive integers no greater than n.

Let $G = (V(G), E(G))$ be a graph. A *clique* in G is a set of pairwise adjacent vertices. A *stable* or *independent* set in G is a set of vertices no two of which are adjacent. The maximum cardinality of a stable set is denoted by $\alpha(G)$, and the maximum cardinality of a clique in G is denoted by $\omega(G)$. For a graph H we say that G *contains* H if H is isomorphic to an induced subgraph of G. We say that G is *H-free* if G does not contain H. For a set $\mathcal{H}$ of graphs, G is *$\mathcal{H}$-free* if G is H-free for every $H \in \mathcal{H}$. For a subset X of $V(G)$, we denote by $G[X]$ the induced subgraph of G with vertex set X, we often use "X" to denote both the set X of vertices and the graph $G[X]$.

Let $X \subseteq V(G)$. We write $N_G(X)$ for the set of all vertices in $G \setminus X$ with at least one neighbor in X, and we define $N_G[X] = N_G(X) \cup X$. When there is no danger of confusion, we omit the subscript "G". For $Y \subseteq V(G)$, we write $N_Y(X) = N_G(X) \cap Y$ and $N_Y[X] = N_Y(X) \cup X$. When $X = \{v\}$ is a singleton, we write $N_Y(v)$ for $N_Y(\{v\})$ and $N_Y[v]$ for $N_Y[\{v\}]$.

Let $x \in V(G)$ and let $Y \subseteq V(G)$. We say that *x is complete to Y in G* if $N_Y[x] = Y$, and we say that *x is anticomplete to Y in G* if $N_G[x] \cap Y = \emptyset$. In particular, if $x \in Y$, then x is neither complete nor anticomplete to Y in G. For subsets X, Y of $V(G)$, we say that *X and Y are complete in G* if every vertex in X is complete to Y in G, and we say that *X and Y are anticomplete in G* if every vertex in X is anticomplete to Y in G. In particular, if X and Y are either complete or anticomplete in G, then $X \cap Y = \emptyset$.

Date: 12^{th} December, 2025.

[§] Princeton University, Princeton, NJ, USA. Supported by NSF grant DMS-2348219 and by AFOSR grant FA9550-22-1-0083.

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

[‡] Corresponding Author (`shajebi@uwaterloo.ca`).

[⋆] CNRS, ENS de Lyon, Université Lyon 1, LIP UMR 5668, 69342 Lyon Cedex 07, France. Partially supported by the French National Research Agency under research grant ANR DIGRAPHS ANR-19-CE48-0013-01 and the LABEX MILYON (ANR-10-LABX-0070) of Université de Lyon, within the program Investissements d'Avenir (ANR-11-IDEX-0007) operated by the French National Research Agency (ANR), and H2020-MSCA-RISE project CoSP- GA No. 823748.

For a graph $G = (V(G), E(G))$, a *tree decomposition* (T, β) of G consists of a tree T and a map $\beta : V(T) \to 2^{V(G)}$ with the following properties:

- For every $v \in V(G)$, there exists $t \in V(T)$ with $v \in \beta(t)$.
- For every $v_1v_2 \in E(G)$, there exists $t \in V(T)$ with $v_1, v_2 \in \beta(t)$.
- $T[\{t \in V(T) \mid v \in \beta(t)\}]$ is connected for all $v \in V(G)$.

The *treewidth* of G, denoted $\mathrm{tw}(G)$, is the smallest integer $w \in \mathbb{N}$ such that G admits a tree decomposition (T, β) with $|\beta(t)| \leq w + 1$ for all $t \in V(T)$. The *tree independence number* of G, denoted $\mathsf{tree}\text{-}\alpha(G)$, is the smallest integer $s \in \mathbb{N}$ such that G admits a tree decomposition (T, β) with $\alpha(G[\beta(t)]) \leq s$ for all $t \in V(T)$.

Both the treewidth and the tree independence number are of great interest in structural and algorithmic graph theory (see [1, 3, 4, 6, 8] for detailed discussions). They are also related quantitatively because, by Ramsey's theorem [11], graphs of bounded clique number and bounded tree independence number have bounded treewidth (see also Lemma 3.2 in [8]). Dallard, Milanič, and Štorgel [8] conjectured that the converse is also true in *hereditary* classes of graphs (meaning classes which are closed under taking induced subgraphs). Let us say that a graph class $\mathcal{G}$ is (tw, ω)*-bounded* if there is a function $f : \mathbb{N} \to \mathbb{N}$ such that every graph $G \in \mathcal{G}$ satisfies $\mathrm{tw}(G) \leq f(\omega(G))$.

Conjecture 1.1 (Dallard, Milanič, and Štorgel [8]). *For every hereditary class $\mathcal{G}$ which is (tw, ω)-bounded, there exists $\tau = \tau(\mathcal{G}) \in \mathbb{N}$ such that $\mathsf{tree}\text{-}\alpha(G) \leq \tau$ for all $G \in \mathcal{G}$.*

Conjecture 1.1 was recently refuted [5] by two of the authors of this paper. It is still natural to ask: which (tw, ω)-bounded hereditary classes have bounded tree independence number? So far, the list of hereditary classes known to be of bounded tree independence number is not very long (see [1, 7, 8] for a few). More hereditary classes are known to be (tw, ω)-bounded. The reasons for the existence of the bound are often highly non-trivial, and it is not known whether the corresponding class has bounded tree independence number. A notable instance is the class of all (theta, prism)-free graphs excluding a fixed forest [2], which we will focus on in this paper.

Let us first give a few definitions. Let P be a graph which is a path. Then we write, for $k \in \mathbb{N}$, $P = p_1\text{-}\cdots\text{-}p_k$ to mean $V(P) = \{p_1, \ldots, p_k\}$, and for all $i, j \in [k]$, the vertices p_i and p_j are adjacent in P if and only if $|i - j| = 1$. We call the vertices p_1 and p_k the *ends of* P, and we say that P is a *path from* p_1 *to* p_k or a *path between* p_1 *and* p_k. We refer to $V(P) \setminus \{p_1, p_k\}$ as the *interior of* P and denote it by P^*. The *length* of a path is its number of edges. Given a graph G, by a *path in* G we mean an induced subgraph of G which is a path. Similarly, for $t \in \mathbb{N} \setminus \{1, 2\}$, given a t-vertex graph C which is a cycle, we write $C = c_1\text{-}\cdots\text{-}c_t\text{-}c_1$ to mean $V(C) = \{c_1, \ldots, c_t\}$, and for all $i, j \in [t]$, the vertices c_i and c_j are adjacent in C if and only if $|i - j| \in \{1, t - 1\}$. The *length* of a cycle is its number of edges (which is the same as its number of vertices). For a graph G, a *hole* in G is an induced subgraph of G which is a cycle of length at least four.

A *theta* is a graph Θ consisting of two non-adjacent vertices a, b, called the *ends of* Θ, and three pairwise internally disjoint paths P_1, P_2, P_3 of length at least two in Θ from a to b, called the *paths of* Θ, such that P_1^*, P_2^*, P_3^* are pairwise anticomplete in Θ (see Figure 1). A *prism* is a graph Π consisting of two triangles $\{a_1, a_2, a_3\}, \{b_1, b_2, b_3\}$ called the *triangles of* Π, and three pairwise disjoint paths P_1, P_2, P_3 in Π, called the *paths of* Π, such that for each $i \in \{1, 2, 3\}$, P_i has ends a_i, b_i, for all distinct $i, j \in \{1, 2, 3\}$, a_ia_j and b_ib_j are the only edges of Π with an end in P_i and an end in P_j, and for every

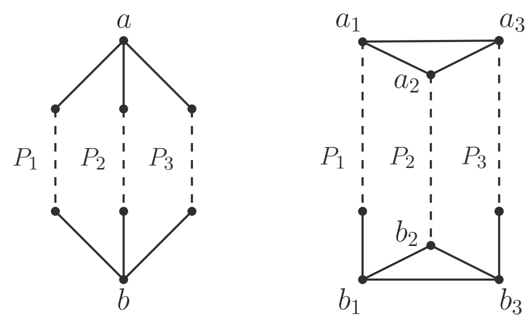

FIGURE 1. A theta (left) and a prism (right). Dashed lines represent paths of arbitrary (possibly zero) length.

$i \neq j \in \{1, 2, 3\}$ $P_i \cup P_j$ is a hole (see Figure 1). If follows that if P_2 has length zero, then each of P_1, P_3 has length at least two. We remark that the last condition is non-standard; the paths of a prism are usually of non-zero length, and a prism with a length-zero path is sometimes called a "line-wheel." For a graph G, a *theta in* G is an induced subgraph of G which is a theta and a *prism in* G is an induced subgraph of G which is a prism.

The following was proved in [2] to show that the local structure of the so-called "layered wheels" [12] is realized in all theta-free graphs of large treewidth. It also characterizes all forests, and remains true when only the usual "prisms" (with no length-zero path) are excluded:

Theorem 1.2 (Abrishami, Alecu, Chudnovsky, Hajebi, Spirkl [2])**.** *Let F be a graph. Then the class of all (theta, prism, F)-free graphs is* (tw, ω)*-bounded if and only if F is a forest.*

We propose the following strengthening (again, this may be true even with the usual "prisms" excluded):

Conjecture 1.3. *For every forest F, there is a constant $\tau = \tau(F) \in \mathbb{N}$ such that for every (theta, prism, F)-free graph G, we have* $\mathsf{tree}\text{-}\alpha(G) \leq \tau$.

As far as we know, Conjecture 1.3 remains open even for paths. But our main result settles the case of stars. For every $t \in \mathbb{N}$, let $\mathcal{C}_t$ be the class of all (theta, prism, $K_{1,t}$)-free graphs. We prove that:

Theorem 1.4. *For every $t \in \mathbb{N}$, there is a constant $f_{1.4} = f_{1.4}(t) \in \mathbb{N}$ such that every graph $G \in \mathcal{C}_t$ satisfies* $\mathsf{tree}\text{-}\alpha(G) \leq f_{1.4}$.

2. Outline of the main proof

Like several earlier results [1, 4, 3] coauthored by the first two authors of this work, the proof of Theorem 1.4 deals with "balanced separators." Let G be a graph and let $w : V(G) \to \mathbb{R}^{\geq 0}$. For every $X \subseteq V(G)$, we write $w(X) = \sum_{v \in X} w(v)$. We say that that w is a *normal weight function on* G if $w(V(G)) = 1$. Given a graph G and a weight function w on G, a subset X of $V(G)$ is called a *w-balanced separator* if for every component D of $G \setminus X$, we have $w(D) \leq 1/2$. The main step in the proof of Theorem 1.4 is the following:

Theorem 2.1. *For every $t \in \mathbb{N}$, there is a constant $f_{2.1} = f_{2.1}(t) \in \mathbb{N}$ with the following property. Let $G \in \mathcal{C}_t$ and let w be a normal weight function on G. Then there exists $Y \subseteq V(G)$ such that $|Y| \leq f_{2.1}$ and $N[Y]$ is a w-balanced separator in G.*

As shown below, Theorem 1.4 follows by combining Theorem 2.1 and the following (this is not a difficult result; see [4] for a proof):

Lemma 2.2 (Chudnovsky, Gartland, Hajebi, Lokshtanov and Spirkl; see Lemma 7.1 in [4])**.** *Let $s \in \mathbb{N}$ and let G be a graph. If for every normal weight function w on G, there is a w-balanced separator X_w in G with $\alpha(X_w) \leq s$, then we have* $\textsf{tree-}\alpha(G) \leq 5s$.

Proof of Theorem 1.4 assuming Theorem 2.1. Let $c = f_{2.1}(t)$. We prove that $f_{1.4}(t) = 5ct$ satisfies the theorem. Let w be a normal weight function on G. By Theorem 2.1, there exists $Y \subseteq V(G)$ such that $|Y| \leq c$ and $X_w = N[Y]$ is a w-balanced separator in G. Assume that there is a stable set S in X_w with $|S| > ct$. Since $S \subseteq N[Y]$, it follows that there is a vertex $y \in Y$ with $|N[y] \cap S| \geq t$. But now G contains $K_{1,t}$, a contradiction. We deduce that $\alpha(X_w) \leq ct$. Hence, by Lemma 2.2, we have $\textsf{tree-}\alpha(G) \leq 5ct = f_{1.4}(t)$. This completes the proof of Theorem 1.4. ■

It remains to prove Theorem 2.1. The idea of the proof is the following. In [3] a technique was developed to prove that separators satisfying the conclusion of Theorem 2.1 exist. It consists of showing that the graph class in question satisfies two properties: being "amiable" and being "amicable." Here we use the same technique. To prove that a graph class is amiable, one needs to analyze the structure of connected subgraphs containing neighbors of a given set of vertices. To prove that a graph is amicable, it is necessary to show that certain carefully chosen pairs of vertices can be separated by well-structured separators. Most of the remainder of the paper is devoted to these two tasks. Section 3 and Section 4 contain structural results asserting the existence of separators that will be used to establish amicability. Section 5 contains definitions and previously known results related to amiability. Section 6 contains the proof of the fact that the class $\mathcal{C}_t$ is amiable. Section 7 uses the results of Sections 3 and 4 to deduce that $\mathcal{C}_t$ is amicable, and to complete the proof of Theorem 2.1.

3. Breaking a wheel

A *wheel* in a graph G is a pair $W = (H, c)$ where H is a hole in G and $c \in G \setminus H$ has at least three neighbors in H. We also use W to denote the vertex set $H \cup \{c\} \subseteq V(G)$. A *sector* of the wheel (H, c) is a path of non-zero length in H whose ends are adjacent to c and whose internal vertices are not. A wheel is *special* if it has exactly three sectors, one sector has length one and the other two (called the *long* sectors) have length at least two (see Figure 2 – A special wheel is sometimes referred to as a "short pyramid.")

For a wheel $W = (H, c)$ in a graph G, we define the set $Z(W) \subseteq H \cup \{c\}$ as follows (see Figure 2). If W is non-special, then $Z(W) = N_H[c]$. Now assume that W is special. Let ab be the sector of length one of W and let d be the neighbor of c in $H \setminus \{a, b\}$. Then we define $Z(W) = \{a, b, c\} \cup N_H[d]$.

Let G be a graph. By a *separation* in G we mean a triple (L, M, R) of pairwise disjoint subsets of $V(G)$ with $L \cup M \cup R = V(G)$, such that neither L nor R is empty and L and R are anticomplete in G. Let $x, y \in V(G)$ be distinct. We say that a set $M \subseteq V(G) \setminus \{x, y\}$ *separates x and y in G* if there exists a separation (L, M, R) in G with $x \in L$ and $y \in R$. Also, for disjoint sets $X, Y \subseteq V(G)$, we say that a set $M \subseteq V(G) \setminus (X \cup Y)$ *separates X and Y* if there exists a separation (L, M, R) in G with $X \subseteq L$ and $Y \subseteq R$. If $X = \{x\}$, we say that M *separates x and Y* to mean M separates X and Y.

We have two results in this section; one for the non-special wheels and one for special wheels:

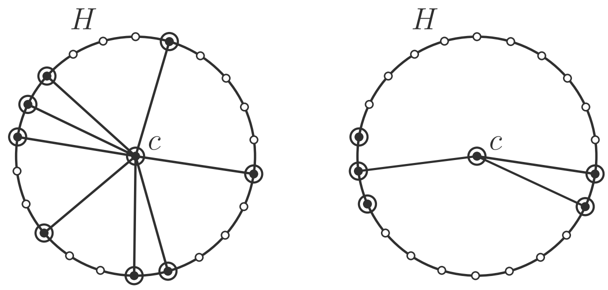

FIGURE 2. A non-special wheel W (left) and a special wheel W (right). Circled nodes represent the vertices in $Z(W)$.

Theorem 3.1. *Let G be a (theta, prism)-free graph, let $W = (H, c)$ be a non-special wheel in G such that H has length at least seven. Let $a, b \in G \setminus N[Z(W)]$ belong to (the interiors of) distinct sectors of W. Then $N[Z(W)]$ separates a and b in G.*

Proof. Let $S = N_H(c)$ and let $T = N[c] \cup (N[S] \setminus H)$. Then $T \subseteq N[Z(W)]$, and so it suffices to show that T separates a and b (note that $a, b \notin T$). We begin with the following:

(1) *Assume that some vertex $v \in G \setminus (W \cup T)$ has either a unique neighbor or two non-adjacent neighbors in some sector $P = p\text{-}\cdots\text{-}p'$ of W. Let b' be the neighbor of p in $W \setminus P$ and b'' be the neighbor of p' in $W \setminus P$. Then $N_H(v) \subseteq P \cup \{b', b''\}$.*

Otherwise, v has a neighbor $d \in H \setminus (P \cup \{b', b''\})$. Also, c has a neighbor $d' \in H \setminus (P \cup \{b', b''\})$, as otherwise W would be a prism or a special wheel. We choose d and d' such that the path Q in $H \setminus P$ from d to d' is minimal. If v has a unique neighbor a in P, then $P \cup Q \cup \{c, v\}$ is a theta in G with ends a and c, a contradiction. Also, if v has two non-adjacent neighbors in P, then $P \cup Q \cup \{v, c\}$ contains a theta with ends c and v. This proves (1).

(2) *For every $v \in G \setminus (W \cup T)$, there exists a sector P of W such that $N_H(v) \subseteq P$.*

Suppose there exists a sector $P = p\text{-}\cdots\text{-}p'$ such that v has two non-adjacent neighbors in P. Then, by (1), we may assume up to symmetry that v is adjacent to the neighbor b of p in $H \setminus P$. By (1), b is the unique neighbor of v in some sector Q of W. So the fact that v has at least two neighbors in P contradicts (1) applied to v and Q.

Suppose there exists a sector $P = p\text{-}\cdots\text{-}p'$ such that v has a unique neighbor a in P. By (1), we may assume that $N_H(v) = \{a, b', b''\}$ where b' is the neighbor of p in $W \setminus P$ and b'' is the neighbor of p' in $W \setminus P$ (because $N_H(v) = \{a, b'\}$ or $N_H(v) = \{a, b''\}$ would imply that v and H form a theta). Let $Q = p\text{-}\cdots\text{-}q$ be the sector of W that contains b'. By (1) applied to v and Q, we have $ap \in E(G)$ and $b''q \in E(G)$. So, b'' is the unique neighbor of v in the sector $R = p'\text{-}\cdots\text{-}q$ of W. By (1) applied to v and R, we have $ap' \in E(G)$ and $b'q \in E(G)$. So H has length six, a contradiction.

We proved that for every sector P of W, either v has no neighbors in P, or v has two neighbors in P, and those neighbors are adjacent. We may therefore assume that v has neighbors in at least three distinct sectors of W, because if v has neighbors in exactly two of them, then $H \cup \{v\}$ would be a prism. So, suppose that $P = p\text{-}\cdots\text{-}p'$, $Q = q\text{-}\cdots\text{-}q'$ and $R = r\text{-}\cdots\text{-}r'$ are three distinct sectors of W, and v is adjacent to $x, x' \in P$, to

$y, y' \in Q$ and to $z, z' \in R$. Suppose up to symmetry that p, x, x', p', q, y, y', q', r, z, z' and r' appear in this order along H. Then there is a theta in G with ends c, v and paths v-x-P-p-c, v-y-Q-q-c and v-z-R-r-c, a contradiction. This proves (2).

To conclude the proof, suppose for a contradiction that the interiors of two distinct sectors of W are contained in the same connected component of $G \setminus T$. Then there exists a path $Y = v$-$\cdots$-w in $G \setminus T$ and two sectors $P = p$-$\cdots$-p' and $Q = q$-$\cdots$-q' of W such that v has neighbors in P^* and w has neighbors in Q^*. By (2), v is anticomplete to $W \setminus P$ and w is anticomplete to $W \setminus Q$ (in particular, Y has length at least one). By choosing such a path Y to be minimal, we deduce that Y^* is anticomplete to H.

Suppose that v has a unique neighbor, or two distinct and non-adjacent neighbors in P. Next, assume that w has a neighbor d in H that is distinct from b' and b'' where b' is the neighbor of p in $W \setminus P$ and b'' is the neighbor of p' in $W \setminus P$, then let d' be a neighbor of c in $H \setminus (P \cup \{b', b''\})$ (d' exists for otherwise, W would be a prism or a special wheel). We choose d and d' such that the path R in $H \setminus P$ from d to d' is minimal. We now see that if v has a unique neighbor a in P, then $P \cup Y \cup R \cup \{c\}$ contains a theta with ends a and c, a contradiction. Also, if v has two distinct non-adjacent neighbors in P, then $P \cup Y \cup R \cup \{c\}$ contains a theta with ends c and v. So, w has only two possible neighbors in H, namely, b' and b''. Due to symmetry, we may assume that $b'w \in E(G)$ (so $b''w \notin E(G)$). It follows that b' is non-adjacent to c. If v has a unique neighbor in P, then $H \cup Y$ is a theta in G, so v has a neighbor in P that is non-adjacent to p. In particular, there exists a path R' from v to p' in $P \cup \{v\}$ that contains no neighbor of p. It follows that $R' \cup Q \cup Y \cup \{c\}$ is a theta in G with ends b' and c.

We deduce that v has exactly two neighbors in P, and those neighbors are adjacent. By the same argument, we can prove that w has exactly two neighbors in P that are adjacent. But now $H \cup Y$ is a prism in G, a contradiction. This completes the proof of Theorem 3.1. ■

Theorem 3.2. *Let G be a (theta, prism)-free graph and let $W = (H, c)$ be a special wheel in G whose long sectors have lengths at least three. Let $a'', b'' \in G \setminus N[Z(W)]$ belong to (the interiors of) distinct sectors of W. Then $N[Z(W)]$ separates a'' and b'' in G.*

Proof. Let ab be the sector of length one of W and let d be the neighbor of c in $H \setminus \{a, b\}$. Let a' be the neighbor of d in the long sector of W containing a and let b' be the neighbors of d in the long sector of W containing b. Then $Z(W) = \{a, a', b, b', c, d\}$. Let P be the path in $H \setminus d$ from a to a' and let Q be the path of $H \setminus d$ from b to b'. Assume, without loss of generality, that $a'' \in P^* \setminus N[Z(W)]$ and let $b'' \in Q^* \setminus N[Z(W)]$.

Let $T = N[c] \cup (N[\{a, b, a', b', d\}] \setminus H)$. Then $T \subseteq N[Z(W)]$, and so it suffices to show that T separates a'' and b'' (note that $a'', b'' \notin T$). Suppose not. Then there exists a path $Y = v$-$\cdots$-w in $G \setminus T$ such that v has neighbors in P^*, w has neighbors in Q^*, $Y \setminus v$ is anticomplete to $W \setminus P$ and $Y \setminus w$ is anticomplete to $W \setminus Q$ (note that possibly $v = w$).

Let x be the neighbor of v in P closest to a along P and let x' be the neighbor of v in P closest to a' along P. Let y be the neighbor of w in Q closest to b along Q and let y' be the neighbor of w in Q closest to b' along Q.

If $x = x'$, then there is a theta in G with ends x and d and paths x-P-a'-d, x-P-a-c-d and x-v-Y-w-y'-Q-b'-d. So, $x \neq x'$, and symmetrically we have $y \neq y'$. If $xx' \notin E(G)$, then there is a theta in G with ends v and d and paths v-x'-P-a'-d, v-x-P-a-c-d and v-Y-w-y'-Q-b'-d. So, $xx' \in E(G)$, and symmetrically we can prove that $yy' \in E(G)$. But now $H \cup Y$ is a prism in G, a contradiction. This completes the proof of Theorem 3.2. ■

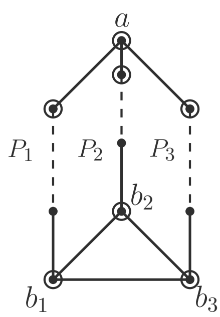

FIGURE 3. A pyramid Σ. Dashed lines represent paths of arbitrary (possibly zero) length, and circled nodes represent the vertices in $Z(\Sigma)$.

4. Breaking a pyramid

A *pyramid* is a graph Σ consisting of a vertex a, a triangle $\{b_1, b_2, b_3\}$ disjoint from a and three paths P_1, P_2, P_3 in Σ of length at least two, such that for each $i \in [3]$, the ends of P_i are a and b_i, and for all distinct $i, j \in [3]$, the sets $V(P_i) \setminus \{a\}$ and $V(P_j) \setminus \{a\}$ are disjoint, b_ib_j is the only edge of G with an end in $V(P_i) \setminus \{a\}$ and an end in $V(P_j) \setminus \{a\}$, and for every $i \neq j \in \{1, 2, 3\}$ $P_i \cup P_j$ is a hole (the assumption that P_1, P_2, P_3 have length at least two is non-standard; usually, one of the paths is allowed to have length 1, and our definition above would refer to a "long" pyramid.)

We say that a is the *apex* of Σ, the triangle $\{b_1, b_2, b_3\}$ is the *base* of Σ, and P_1, P_2, P_3 are the *paths* of Σ. We also define $Z(\Sigma) = N_\Sigma[a] \cup \{b_1, b_2, b_3\}$ (so we have $|Z(\Sigma)| = 7$). For a graph G, by a *pyramid in G* we mean an induced subgraph of G which is a pyramid (see Figure 3).

The main result of this section, Theorem 4.1 below, follows from much more general results of [2]. However, there is also a short and self-contained proof, which we include here:

Theorem 4.1. *Let G be a (theta, prism)-free graph and let Σ be a pyramid in G with apex a, base $\{b_1, b_2, b_3\}$ and paths P_1, P_2 and P_3 as in the definition. Let $u, v \in G \setminus N[Z(\Sigma)]$ belong to distinct paths of Σ. Then $N[Z(\Sigma)]$ separates u and v in G.*

Proof. Suppose not. Then there exist $u, v \in G \setminus N[Z(\Sigma)]$, belonging to distinct paths of Σ, such that $N[Z(\Sigma)]$ does not separate u and v in G. It follows that for distinct $i, j \in [3]$, there exists a path $Q = x\text{-}\cdots\text{-}y$ in $G \setminus (\Sigma \cup N[Z(\Sigma)])$ such that x has a neighbor in P_i^* and y has a neighbor in P_j^*. We choose $i, j \in [3]$ and Q subject to the minimality of Q. By symmetry, we may assume that $i = 1$ and $j = 2$.

From the minimality of Q and the fact that $Q \subseteq V(G) \setminus (\Sigma \cup N[Z(\Sigma)])$, it follows that:

- $N_{P_1}(x) \subseteq P_1 \setminus Z(\Sigma)$, and $Q \setminus x$ and P_1 are anticomplete in G.
- $N_{P_2}(y) \subseteq P_2 \setminus Z(\Sigma)$, and $Q \setminus y$ and P_2 are anticomplete in G.

Now, if some vertex of Q has a neighbor in P_3, then by the minimality of Q, we must have $x = y$. In particular, x has neighbors in P_1, P_2 and P_3. Since a and x are not adjacent in G (for otherwise there is a theta in G), it follows that the three paths in G from a to x with interiors in P_1, P_2 and P_3 form a theta in G with ends a and x, a contradiction. We deduce that Q and P_3 are anticomplete in G.

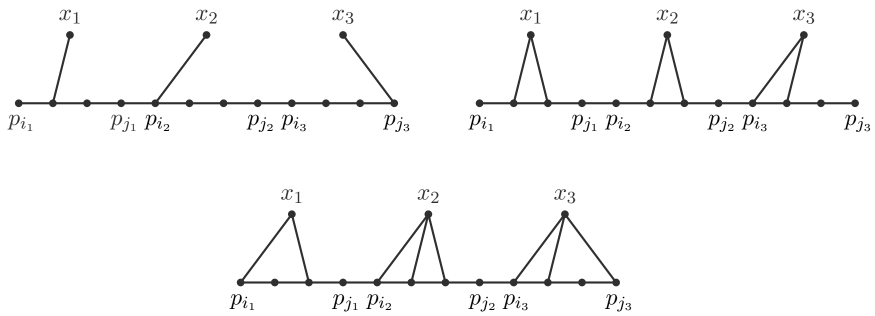

FIGURE 4. A consistent alignment which is spiky (top left), triangular (top right) and wide (bottom).

Let x' be the neighbor of x in P_1 closest to a along P_1 and let x'' be the neighbor of x in P_1 closest to b_1 along P_1. Similarly, let y' be the neighbor of y in P_2 closest to a along P_2 and let y'' be the neighbor of y in P_2 closest to b_2 along P_2. Recall that $x', x'' \in P_1 \setminus Z(\Sigma)$ and $y', y'' \in P_2 \setminus Z(\Sigma)$. If $x' = x''$, then there is a theta in G with ends a, x' and paths $a\text{-}P_1\text{-}x'$, $a\text{-}P_2\text{-}y'\text{-}y\text{-}Q\text{-}x\text{-}x'$ and $a\text{-}P_3\text{-}b_3\text{-}b_1\text{-}P_1\text{-}x'$. Also, if x' and x'' are distinct and adjacent in G, then there is a prism in G with triangles $x''xx'$ and $b_1b_2b_3$ and paths $x''\text{-}P_1\text{-}b_1$, $x\text{-}Q\text{-}y\text{-}y''\text{-}P_2\text{-}b_2$ and $x'\text{-}P_1\text{-}a\text{-}P_3\text{-}b_3$. Hence, we have $x' \neq x''$ and $x'x'' \notin E(G)$. But now there is a theta in G with ends a, x and paths $a\text{-}P_1\text{-}x'\text{-}x$, $a\text{-}P_2\text{-}y'\text{-}y\text{-}Q\text{-}x$ and $a\text{-}P_3\text{-}b_3\text{-}b_1\text{-}P_1\text{-}x''\text{-}x$, a contradiction. This completes the proof of Theorem 4.1. ■

5. ALIGNMENTS AND CONNECTIFIERS

This section covers a number of definitions and a result from [3], which we will use in the proof of Theorem 2.1.

Let G be a graph, let P be a path in G and let $X \subseteq V(G) \setminus P$. We say that (P, X) is an *alignment* if every vertex of X has at least one neighbor in P and one may write $P = p_1\text{-}\cdots\text{-}p_n$ and $X = \{x_1, \ldots, x_k\}$ for $k, n \in \mathbb{N}$ such that there exist $1 \leq i_1 \leq j_1 < i_2 \leq j_2 < \cdots < i_k \leq j_k \leq n$ where $N_P(x_l) \subseteq p_{i_l}\text{-}P\text{-}p_{j_l}$ for every $l \in [k]$. This is a little different from the definition in [3], but the difference is not substantial, and using this definition is more convenient for us here. In this case, we say that $x_1, \ldots, x_k$ is *the order on X given by the alignment* (P, X). An alignment (P, X) is *wide* if each of $x_1, \ldots, x_k$ has two non-adjacent neighbors in P, *spiky* if each of $x_1, \ldots, x_k$ has a unique neighbor in P and *triangular* if each of $x_1, \ldots, x_k$ has exactly two neighbors in P and those neighbors are adjacent. An alignment is *consistent* if it is wide, spiky or triangular. See Figure 4.

By a *caterpillar* we mean a tree C with maximum degree three such that no two branch vertices in C are adjacent, and such that there exists a path P in C containing all branch vertices of C. We call a minimal such path P the *spine* of C. (We note that our definition of a "caterpillar" is non-standard in multiple ways.) By a *subdivided star* we mean a graph isomorphic to a subdivision of the complete bipartite graph $K_{1,\delta}$ for some $\delta \geq 3$. In other words, a subdivided star is a tree with exactly one branch vertex, which we call its *root.* For a graph H, a vertex v of H is said to be *simplicial* if $N_H(v)$ is a clique. We denote by $\mathcal{Z}(H)$ the set of all simplicial vertices of H. Note that for every tree T, $\mathcal{Z}(T)$ is the set of all leaves of T. An edge e of a tree T is said to be a *leaf-edge* of T if e is incident with a leaf of T. It follows that if H is the line graph of a tree T, then $\mathcal{Z}(H)$ is the set of all vertices in H corresponding to the leaf-edges of T.

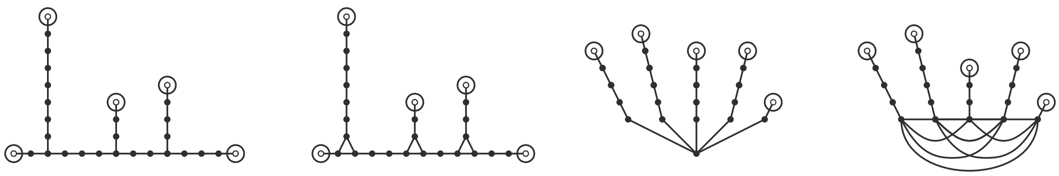

FIGURE 5. Examples of a connectifier. Circled nodes represent the vertices in X.

Let H be a graph that is either a caterpillar, or the line graph of a caterpillar, or a subdivided star with root r, or the line graph of a subdivided star with root r. We define an induced subgraph of H, denoted by $P(H)$, which we will use throughout the paper. If H is a path (possibly of length zero), then let $P(H) = H$. If H is a caterpillar, then let $P(H)$ be the spine of H. If H is the line graph of a caterpillar C, then let $P(H)$ be the path in H consisting of the vertices of H that correspond to the edges of the spine of C. If H is a subdivided star with root r, then let $P(H) = \{r\}$. If H is the line graph of a subdivided star S with root r, let $P(H)$ be the clique of H consisting of the vertices of H that correspond to the edges of S incident with r. The *legs* of H are the components of $H \setminus P(H)$. Let G be a graph and let H be an induced subgraph of G that is either a caterpillar, or the line graph of a caterpillar, or a subdivided star or the line graph of a subdivided star. Let $X \subseteq V(G) \setminus H$ such that every vertex of X has a unique neighbor in H and $N_H(X) = \mathcal{Z}(H)$ (see Figure 5). We call (H, X) a *connectifier*. Also, if H is a single vertex and $X \subseteq N(H)$, we call (H, X) a *connectifier* as well. We say that the connectifier (H, X) is *concentrated* if H is a subdivided star or the line graph of a subdivided star or a singleton.

Let (H, X) be a connectifier in G which is not concentrated. So H is a caterpillar or the line graph of a caterpillar. Let S be the set of vertices of $H \setminus P(H)$ that have neighbors in $P(H)$. Then $(P(H), S)$ is an alignment. Let $s_1, \dots, s_k$ be the corresponding order on S given by $(P(H), S)$. Now, order the vertices of X as $x_1, \dots, x_k$ where for every $i \in [k]$, the vertex x_i has a neighbor in the leg of H containing s_i. We say that $x_1, \dots, x_k$ is *the order on X given by* (H, X).

The following was proved in [3]:

Theorem 5.1 (Chudnovsky, Gartland, Hajebi, Lokshtanov and Spirkl; Theorem 5.2 in [3])**.** *For every integer $h \in \mathbb{N}$, there is a constant $f_{5.1} = f_{5.1}(h) \in \mathbb{N}$ with the following property. Let G be a connected graph. Let $S \subseteq V(G)$ such that $|S| \geq f_{5.1}$, the graph $G \setminus S$ is connected and every vertex of S has a neighbor in $G \setminus S$. Then there exists $S' \subseteq S$ with $|S'| = h$ as well as an induced subgraph H of $G \setminus S$ for which one of the following holds.*

- *(H, S') is a connectifier, or*
- *H is a path and every vertex in S' has a neighbor in H.*

6. Amiability

The two notions of "amiability" and "amicability," first introduced in [3], are at the heart of the proof of Theorem 2.1. We deal with the former in this section and leave the latter for the next one.

Let $s \in \mathbb{N}$ and let G be a graph. An *s-trisection in G* is a separation (D_1, Y, D_2) in G such that the following hold.

- Y is a stable set with $|Y| = s$.
- D_1 and D_2 are components of $G \setminus Y$ with $N(D_1) = N(D_2) = Y$.
- D_1 is a path and for every $y \in Y$ there exists $d_y \in D_1$ such that $N_Y(d_y) = \{y\}$.

(The reader may notice that we will never use the second condition in the third bullet point. It was however necessary in [3], so we keep it for easier cross-referencing.)

We say that a graph class $\mathcal{G}$ is *amiable* if there is a function $\sigma : \mathbb{N} \to \mathbb{N}$ with the following property. Let $x \in \mathbb{N}$, let $G \in \mathcal{G}$ and let (D_1, Y, D_2) be a $\sigma(x)$-trisection in G. Then there exist $H \subseteq D_2$ and $X \subseteq Y$ with $|X| = x$ such that the following hold.

- (D_1, X) is a consistent alignment.
- (H, X) is either a connectifier or a consistent alignment.
- If (H, X) is not a concentrated connectifier, then the orders given on X by (D_1, X) and by (H, X) are the same.

In this case, we say that *H and X are given by amiability*. The main result of this section is the following:

Theorem 6.1. *For every $t \in \mathbb{N}$, the class $\mathcal{C}_t$ is amiable. Moreover, with notation as in the definition of amiability, if (H, X) is a connectifier, then we have $|H| > 1$.*

In order to prove Theorem 6.1, first we prove the following lemma:

Lemma 6.2. *Let $d, s \in \mathbb{N}$, let G be a theta-free graph and let Y be a stable set in G of cardinality $3s(d+1)$. Let P be a path in $G \setminus Y$ such that every vertex in Y has a neighbor in P, and each vertex of P has fewer than d neighbors in Y. Assume that for every two vertices $y, y' \in Y$, there is a path R in G from y to y' such that P and R^* are disjoint and anticomplete in G. Then there is an s-subset S of Y such that (P, S) is a consistent alignment.*

Proof. For every vertex $y \in Y$, let P_y be the (unique) path in P with the property that y is complete to the ends of P_y and anticomplete to $P \setminus P_y$. Let I be the graph with $V(I) = Y$ such that two distinct vertices $y, y' \in Y$ are adjacent in I if and only if $P_y \cap P_{y'} \neq \emptyset$. Then I is an interval graph and so I is perfect [10]. Since $|V(I)| = 3s(d+1)$, it follows that I contains either a clique of cardinality $d+1$ or a stable set of cardinality $3s$.

Assume that I contains a clique of cardinality $d+1$. Then there exists $C \subseteq Y$ with $|C| = d+1$ and $p \in P$ such that $p \in P_y$ for every $y \in C$. Since $p \in P$ has fewer than d neighbors in $C \subseteq Y$, it follows that there are at least two vertices $y, y' \in C \setminus N(p)$. Since $p \in P_y \cap P_{y'}$, it follows that $P \setminus \{p\}$ has two components, and each of y and y' has a neighbor in each component of $P \setminus \{p\}$. It follows that there are two paths P_1 and P_2 from y to y' with disjoint and anticomplete interiors contained in P. On the other hand, there is a path R in G from y to y' such that P and R^* are disjoint and anticomplete in G. It follows that P_1, P_2 and R are pairwise internally disjoint and anticomplete. But now there is a theta in G with ends y, y' and paths P_1, P_2, R, a contradiction.

We deduce that I contains a stable set S' of cardinality $3s$. From the definition of I, it follows that (P, S') is an alignment. Hence, since every vertex in S' has one, two adjacent, or at least two non-adjacent neighbors in P, there exists $S \subseteq S' \subseteq Y$ with $|S| = s$ such that (P, S) is a consistent alignment. This completes the proof of Lemma 6.2. ■

Proof of Theorem 6.1. For every $x \in \mathbb{N}$, let

$$s = f_{5.1}(3x^2(t+1))$$

and let

$$\sigma(x) = 3s(t+1).$$

We will show that $\mathcal{C}_t$ is amiable with respect to $\sigma : \mathbb{N} \to \mathbb{N}$ as defined above. Let $x \in \mathbb{N}$, let $G \in \mathcal{C}_t$ and let (D_1, Y, D_2) be a $\sigma(x)$-trisection in G. Then Y is a stable set of cardinality $3s(t+1)$, D_1 is a path in $G \setminus Y$ and every vertex in Y has a neighbor in D_1. Moreover, since G is $K_{1,t}$-free, no vertex in D_1 has t or more neighbors in Y, and since $N(D_2) = Y$, it follows that for every two vertices $y, y' \in Y$, there is a path R in G from y to y' with $R^* \subseteq D_2$, and so D_1 and R^* are disjoint and anticomplete in G. By Lemma 6.2, there exists $S \subseteq Y$ with $|S| = s$ such that (D_1, S) is a consistent alignment.

Now, we show that there exists $H \subseteq D_2$ as well as an x-subset X of $S \subseteq Y$ such that H and X satisfy the definition of amiability. Since D_2 is connected and every vertex in $S \subseteq Y$ has a neighbor in D_2, it follows that $D_2 \cup S$ is connected too. Since $|S| = s = f_{5.1}(3x^2(t+1))$, it follows from Theorem 5.1 that there exists $S' \subseteq S$ with $|S'| = 3x^2(t+1)$ and an induced subgraph H_2 of D_2 for which one of the following holds:

- (H_2, S') is a connectifier.
- H_2 is a path and every vertex of S' has a neighbor in H_2.

First, assume that (H_2, S') is a concentrated connectifier. Then, since $|S'| \geq t$ and G is $K_{1,t}$-free, it follows that $|H_2| > 1$. Now, since $|S'| \geq x$, we may choose a concentrated connectifier (H, X) where X is an x-subset of $S' \subseteq S \subseteq Y$ and H is an induced subgraph $H_2 \subseteq D_2$ with $|H| > 1$. In particular, H and X satisfy the definition of amiability.

Next, assume that (H_2, S') is a connectifier which is not concentrated. Consider the orders on S' given by (D_1, S') and by (H_2, S'). Since $|S'| \geq x^2$, it follows from the Erdős-Szekeres theorem [9] that there is an x-subset X of $S' \subseteq S \subseteq Y$ as well as an induced subgraph H of $H_2 \subseteq D_2$ such that:

- (D_1, X) is a consistent alignment (because (D_1, S) is);
- (H, X) is a connectifier which is not concentrated; and
- The orders given on X by (D_1, X) and by (H, X) are the same.

It follows that H and X satisfy the definition of amiability.

Finally, assume that H_2 is a path and every vertex in S' has a neighbor in H_2. Let $H = H_2$. Recall that (D_1, S') is an alignment. In particular, S' is a stable set of cardinality $3x^2(t+1)$, and since G is $K_{1,t}$-free, no vertex in H_2 has t or more neighbors in S'. Also, for every two vertices $y, y' \in S$, there is a path R in G from y to y' such that $R^* \subseteq D_1$, and so H and R^* are disjoint and anticomplete in G. By Lemma 6.2, there exists $S'' \subseteq S' \subseteq S$ with $|S''| = x^2$ such that (H, S'') is a consistent alignment. Consider the order on S'' given by (D_1, S'') and by (H, S''). Since $|S''| = x^2$, it follows from the Erdős-Szekeres theorem [9] that there is an x-subset X of $S'' \subseteq S' \subseteq S \subseteq Y$ such that such that:

- (D_1, X) is a consistent alignment (because (D_1, S) is);
- (H, X) is a consistent alignment (because (H, S'') is); and
- The orders given on X by (D_1, X) and by (H, X) are the same.

So H and X satisfy the definition of amiability. This completes the proof of Theorem 6.1 ■

7. Amicability

Here we complete the proof of Theorem 2.1, beginning with the following definition.

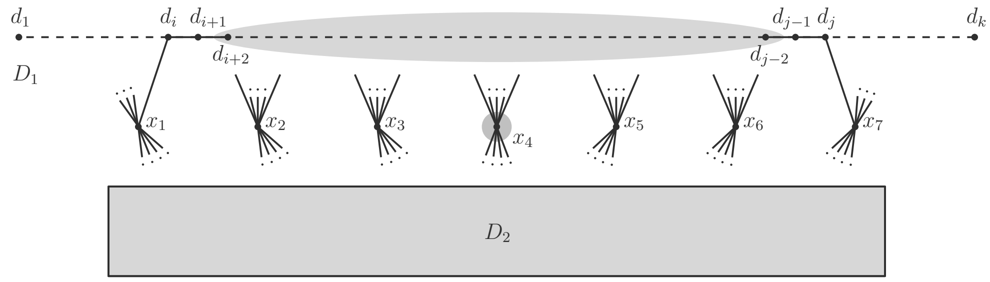

FIGURE 6. Amicability – Note that Z is contained in the highlighted set.

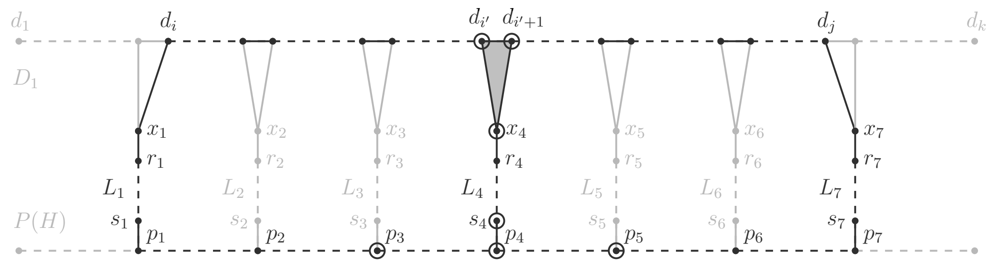

FIGURE 7. H is a caterpillar. Circled nodes depict the vertices in $Z(\Sigma)$.

Let $m \in \mathbb{N}$ and let $\mathcal{G}$ be a graph class. We say that $\mathcal{G}$ is *m-amicable* if $\mathcal{G}$ is amiable and the following holds. Let $\sigma : \mathbb{N} \to \mathbb{N}$ be as in the definition of amiability for $\mathcal{G}$. Let $G \in \mathcal{G}$ and let (D_1, Y, D_2) be a $\sigma(7)$-trisection in G. Let $X = \{x_1, \dots, x_7\} \subseteq Y$ be given by amiability such that $x_1, \dots, x_7$ is the order on X given by (D_1, X). Let $D_1 = d_1\text{-}\cdots\text{-}d_k$ such that traversing D_1 from d_1 to d_k, the first vertex in D_1 with a neighbor in X is a neighbor of x_1. Let $i \in [k]$ be maximum such that x_1 is adjacent to d_i and let $j \in [k]$ be minimum such that x_7 is adjacent to d_j. Then there exists a subset Z of $D_2 \cup \{d_k : i + 2 \leq k \leq j - 2\} \cup \{x_4\}$ with $|Z| \leq m$ such that $N[Z]$ separates d_i and d_j. It follows that $N[Z]$ separates $d_1\text{-}D_1\text{-}d_i$ and $d_j\text{-}D_1\text{-}d_k$ (see Figure 6).

We prove that:

Theorem 7.1. *For every $t \in \mathbb{N}$, the class $\mathcal{C}_t$ is $\max\{2t, 7\}$-amicable.*

Proof. By Theorem 6.1, $\mathcal{C}_t$ is amiable, and with notation as in the definition of amiability, if (H, X) is a connectifier, then we have $|H| > 1$. Let $\sigma : \mathbb{N} \to \mathbb{N}$ be as in the definition of amiability for $\mathcal{C}_t$. Let $G \in \mathcal{C}_t$ and let (D_1, Y, D_2) be a $\sigma(7)$-trisection in G. Let $X = \{x_1, \dots, x_7\} \subseteq Y$ be given by amiability such that $x_1, \dots, x_7$ is the order on X given by the consistent alignment (D_1, X). Let $D_1 = d_1\text{-}\cdots\text{-}d_k$ and $i, j \in [k]$ be as in the definition of amicability. Our goal is to show that there exists a subset Z of $D_2 \cup \{d_k : i + 2 \leq k \leq j - 2\} \cup \{x_4\}$ with $|Z| \leq \max\{2t, 7\}$ such that $N[Z]$ separates d_i and d_j.

Let $i' \in [k]$ be minimum such that x_4 is adjacent to $d_{i'}$, let $j' \in [k]$ be maximum such that x_4 is adjacent to $d_{j'}$, and let H be the induced subgraph of D_2 given by amiability. It follows that $i + 2 < i' \leq j' < j - 2$, (H, X) is either a connectifier with $|H| > 1$ or a consistent alignment, and if (H, X) is not a concentrated connectifier, then $x_1, \dots, x_7$ is the order on X given by (H, X). When (H, X) is a connectifier with $|H| > 1$, then for each $l \in [7]$, let r_l be the unique neighbor of x_l in H (so $r_l \in \mathcal{Z}(H)$) and let L_l

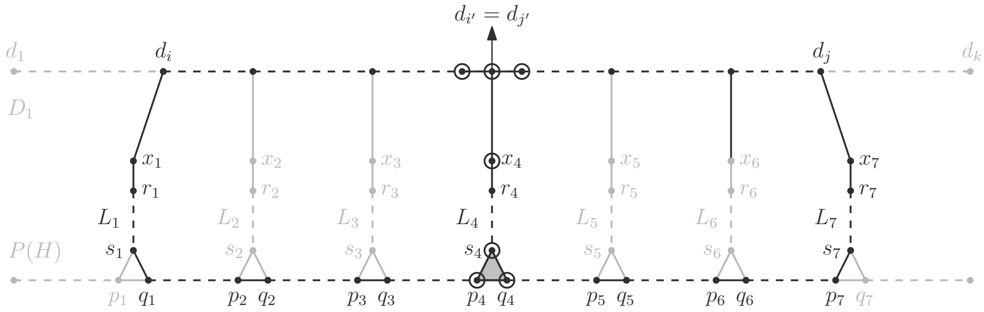

FIGURE 8. H is the line graph of a caterpillar and (D_1, X) is spiky. Circled nodes represent the vertices in $Z(\Sigma)$.

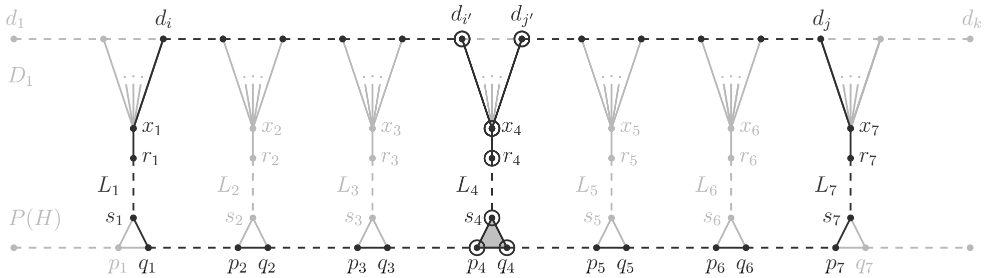

FIGURE 9. H is the line graph of a caterpillar and (D_1, X) is wide. Circled nodes represent the vertices in $Z(\Sigma)$.

be the (unique) shortest path in H from r_l to a vertex $s_l \in N_H[P(H)]$. It follows that $s_l \in H \setminus P(H)$ unless H is the line graph of a subdivided star where not all edges of the star are subdivided, in which case we have $r_l = s_l \in P(H) = \mathcal{Z}(H) = H$.

First, consider the case where H is a caterpillar. It follows that for each $l \in [7]$, we have $s_l \in H \setminus P(H)$ and s_l has a unique neighbor $p_l \in P(H)$. Since G is theta-free, it follows that (D_1, X) is triangular, and so $j' = i' + 1$ (see Figure 7). Let Σ be the pyramid with apex p_4, base $\{d_{i'}, x_4, d_{j'}\}$ and paths

$$P_1 = p_4\text{-}P(H)\text{-}p_1\text{-}s_1\text{-}L_1\text{-}r_1\text{-}x_1\text{-}d_i\text{-}D_1\text{-}d_{i'};$$

$$P_2 = p_4\text{-}s_4\text{-}L_4\text{-}r_4\text{-}x_4;$$

$$P_3 = p_4\text{-}P(H)\text{-}p_7\text{-}s_7\text{-}L_7\text{-}r_7\text{-}x_7\text{-}d_j\text{-}D_1\text{-}d_{j'}.$$

Then $Z(\Sigma)$ is a 7-subset of $D_2 \cup \{d_k : i + 2 \leq k \leq j - 2\} \cup \{x_4\}$. Moreover, we have $d_i \in P_1^* \setminus N[Z(\Sigma)]$ and $d_j \in P_3^* \setminus N[Z(\Sigma)]$. Therefore, by Theorem 4.1, $N[Z(\Sigma)]$ separates d_i and d_j, as desired.

Second, consider the case where H is the line graph of a caterpillar. It follows that for each $l \in [7]$, we have $s_l \in H \setminus P(H)$ and s_l has exactly two neighbors $p_l, q_l \in P(H)$, where p_l and q_l are adjacent, and the vertices $p_1, q_1, p_2, q_2, \ldots, p_7, q_7$ appear on $P(H)$ in this order. Since G is prism-free, it follows that (D_1, X) is either spiky or wide. Suppose that (D_1, X) is spiky (see Figure 8). Then $i' = j'$. Let Σ be the pyramid with apex $d_{i'} = d_{j'}$, base $\{p_4, s_4, q_4\}$ and paths

$$P_1 = d_{i'}\text{-}D_1\text{-}d_i\text{-}x_1\text{-}r_1\text{-}L_1\text{-}s_1\text{-}q_1\text{-}P(H)\text{-}p_4;$$

$$P_2 = d_{i'}\text{-}x_4\text{-}r_4\text{-}L_4\text{-}s_4;$$

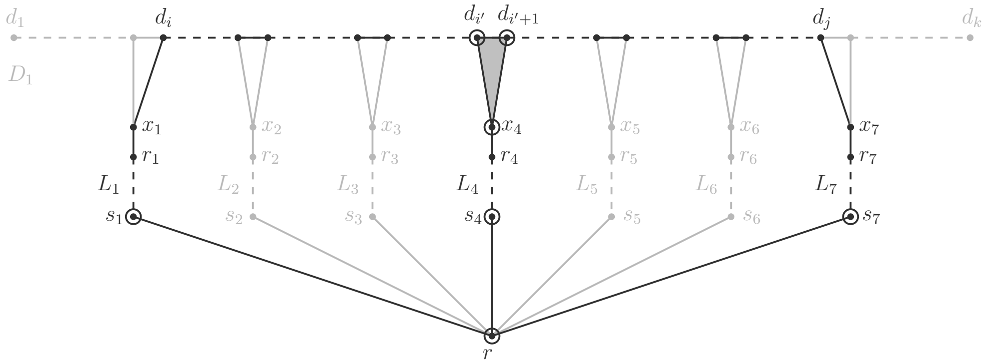

FIGURE 10. H is a subdivided star. Circled nodes represent the vertices in $Z(\Sigma)$.

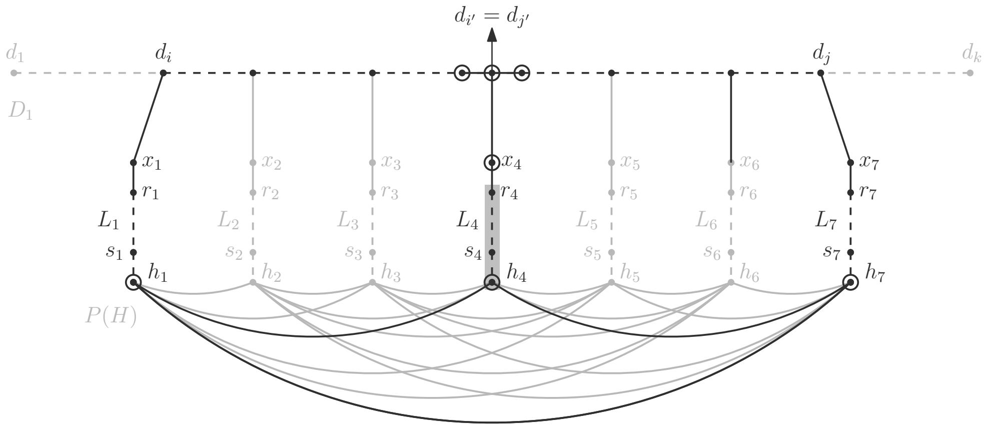

FIGURE 11. H is the line graph of a subdivided star and (D_1, X) is spiky. Circled nodes represent the vertices in $Z(\Sigma)$, and the highlighted path may be of length zero.

$$P_3 = d_{i'}\text{-}D_1\text{-}d_j\text{-}x_7\text{-}r_7\text{-}L_7\text{-}s_7\text{-}p_7\text{-}P(H)\text{-}q_4.$$

Then $Z(\Sigma)$ is a 7-subset of $D_2 \cup \{d_k : i+2 \leq k \leq j-2\} \cup \{x_4\}$. Moreover, we have $d_i \in P_1^* \setminus N[Z(\Sigma)]$ and $d_j \in P_3^* \setminus N[Z(\Sigma)]$. So by Theorem 4.1, $N[Z(\Sigma)]$ separates d_i and d_j. Now assume that (D_1, X) is wide (see Figure 9). Then $j' - i' > 1$. Let Σ be the pyramid with apex x_4, base $\{p_4, s_4, q_4\}$ and paths

$$P_1 = x_4\text{-}d_{i'}\text{-}D_1\text{-}d_i\text{-}x_1\text{-}r_1\text{-}L_1\text{-}s_1\text{-}q_1\text{-}P(H)\text{-}p_4;$$

$$P_2 = x_4\text{-}r_4\text{-}L_4\text{-}s_4;$$

$$P_3 = x_4\text{-}d_{j'}\text{-}D_1\text{-}d_j\text{-}x_7\text{-}r_7\text{-}L_7\text{-}s_7\text{-}p_7\text{-}P(H)\text{-}q_4.$$

Let $Z = (N(x_4) \cap \Sigma) \cup \{p_4, s_4, q_4\}$. Then $Z(\Sigma)$ is a 7-subset of $D_2 \cup \{d_k : i+2 \leq k \leq j-2\} \cup \{x_4\}$. Also, we have $d_i \in P_1^* \setminus N[Z(\Sigma)]$ and $d_j \in P_3^* \setminus N[Z(\Sigma)]$. So by Theorem 4.1, $N[Z(\Sigma)]$ separates d_i and d_j, as required.

Third, consider the case where H is a subdivided star with root r. It follows that $P(H) = \{r\}$ and $H \neq \{r\}$ (because $|H| > 1$). For each $l \in [7]$, we have $r_l, s_l \in H \setminus P(H)$ and r_l is a leaf of H. Since G is theta-free, it follows that (D_1, X) is triangular and so

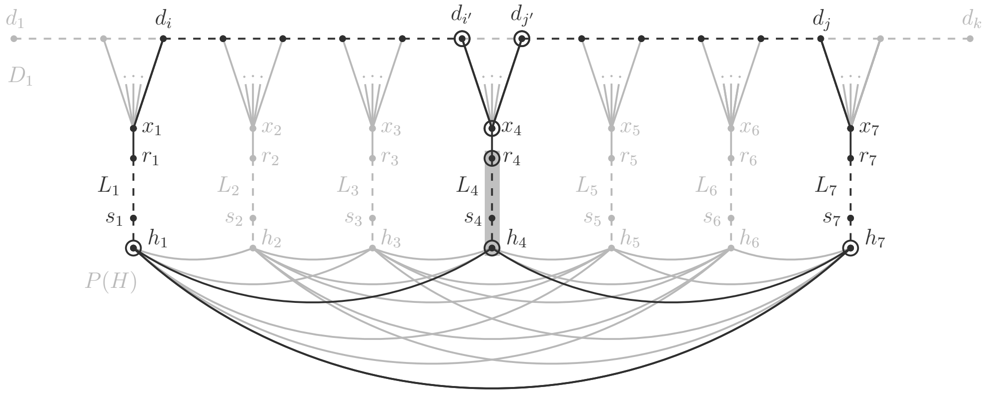

FIGURE 12. H is the line graph of a subdivided star, (D_1, X) is wide and the vertices r_4, s_4, h_4 are not all the same. Circled nodes represent the vertices in $Z(\Sigma)$, and the highlighted path has length at least one.

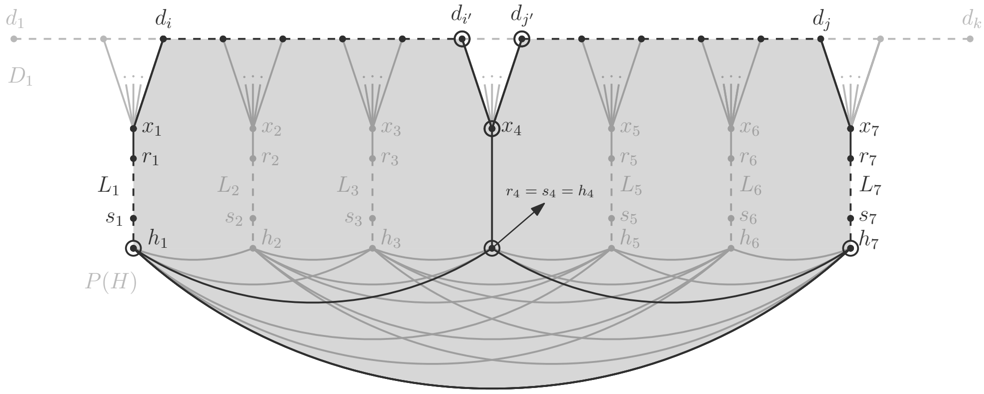

FIGURE 13. H is the line graph of a subdivided star, (D_1, X) is wide and $r_4 = s_4 = h_4$. The hole C is highlighted, and circled nodes represent the vertices in $Z(W)$.

$j' - i' = 1$ (see Figure 10). Let Σ be the pyramid with apex r, base $\{d_{i'}, x_4, d_{j'}\}$ and paths

$$P_1 = r\text{-}s_1\text{-}L_1\text{-}r_1\text{-}x_1\text{-}d_i\text{-}D_1\text{-}d_{i'};$$

$$P_2 = r\text{-}s_4\text{-}L_4\text{-}r_4\text{-}x_4;$$

$$P_3 = r\text{-}s_7\text{-}L_7\text{-}r_7\text{-}x_7\text{-}d_j\text{-}D_1\text{-}d_{j'}.$$

Then $Z(\Sigma)$ is a 7-subset of $D_2 \cup \{d_k : i+2 \leq k \leq j-2\} \cup \{x_4\}$. Also, we have $d_i \in P_1^* \setminus N[Z(\Sigma)]$ and $d_j \in P_3^* \setminus N[Z(\Sigma)]$. So it follows from Theorem 4.1 that $N[Z(\Sigma)]$ separates d_i and d_j, as desired.

Fourth, consider the case where H is the line graph of a subdivided star. It follows that for each $l \in [7]$, either we have $s_l \in P(H)$, in which case we set $h_l = s_l$, or we have $s_l \in H \setminus P(H)$, in which case we choose h_l to be the unique neighbor of s_l in $P(H)$. Since G is prism-free, it follows that (D_1, X) is either spiky or wide. There are now three cases to analyze:

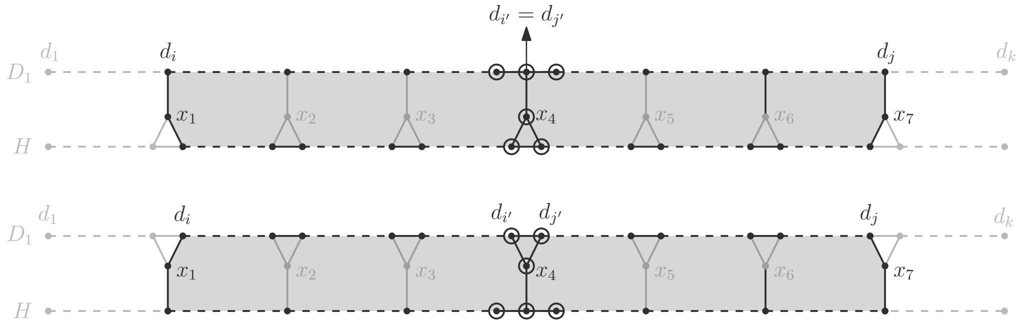

FIGURE 14. One of (D_1, X) and (H, X) is spiky and the other is triangular. The hole C is highlighted, and circled nodes represent the vertices in $Z(W)$.

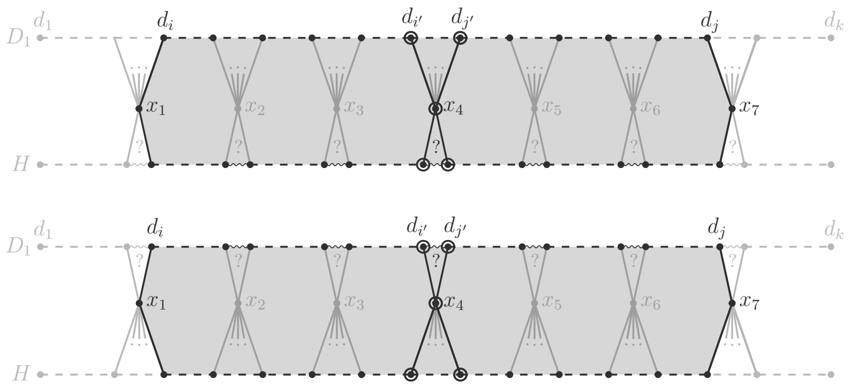

FIGURE 15. One of (D_1, X) and (H, X) is wide. The hole C is highlighted, and circled nodes represent the vertices in $Z(W)$.

Case 1. Suppose that (D_1, X) is spiky (see Figure 11). Then we have $i' = j'$. Consider the pyramid Σ in G with apex $d_{i'} = d_{j'}$, base $\{h_1, h_4, h_7\}$ and paths

$$P_1 = d_{i'}\text{-}D_1\text{-}d_i\text{-}x_1\text{-}r_1\text{-}L_1\text{-}s_1\text{-}h_1;$$

$$P_2 = d_{i'}\text{-}x_4\text{-}r_4\text{-}L_4\text{-}s_4\text{-}h_4;$$

$$P_3 = d_{i'}\text{-}D_1\text{-}d_j\text{-}x_7\text{-}r_7\text{-}L_7\text{-}s_7\text{-}h_7.$$

Then $Z(\Sigma)$ is a 7-subset of $D_2 \cup \{d_k : i+2 \leq k \leq j-2\} \cup \{x_4\}$. Moreover, we have $d_i \in P_1^* \setminus N[Z(\Sigma)]$ and $d_j \in P_3^* \setminus N[Z(\Sigma)]$. Thus, by Theorem 4.1, $N[Z(\Sigma)]$ separates d_i and d_j.

Case 2. Suppose that (D_1, X) is wide and the vertices r_4, s_4, h_4 are not all the same (see Figure 12). Then $j' - i' > 1$. Let Σ be the pyramid with apex x_4, base $\{h_1, h_4, h_7\}$ and paths

$$P_1 = x_4\text{-}d_{i'}\text{-}D_1\text{-}d_i\text{-}x_1\text{-}r_1\text{-}L_1\text{-}s_1\text{-}h_1;$$

$$P_2 = x_4\text{-}r_4\text{-}L_4\text{-}s_4\text{-}h_4;$$

$$P_3 = x_4\text{-}d_{j'}\text{-}D_1\text{-}d_j\text{-}x_7\text{-}r_7\text{-}L_7\text{-}s_7\text{-}h_7.$$

Then $Z(\Sigma)$ is a 7-subset of $D_2 \cup \{d_k : i+2 \leq k \leq j-2\} \cup \{x_4\}$, and we have $d_i \in P_1^* \setminus N[Z(\Sigma)]$ and $d_j \in P_3^* \setminus N[Z(\Sigma)]$. It follows from Theorem 4.1 that $N[Z(\Sigma)]$ separates d_i and d_j.

Case 3. Suppose that (D_1, X) is wide and $r_4 = s_4 = h_4$ (see Figure 13). Then $j' - i' > 1$. Let $C = x_4\text{-}d_{i'}\text{-}D_1\text{-}d_i\text{-}x_1\text{-}r_1\text{-}L_1\text{-}s_1\text{-}h_1\text{-}h_7\text{-}s_7\text{-}L_7\text{-}r_7\text{-}x_7\text{-}d_j\text{-}D_1\text{-}d_{j'}\text{-}x_4$. Then C is a hole on more than seven vertices and $W = (C, h_4)$ is a special wheel in G where $Z(W) = \{d_{i'}, d_{j'}, h_1, h_4, h_7, x_4\}$; in particular, $Z(W)$ is a 6-subset of $D_2 \cup \{d_k : i+2 \leq k \leq j-2\} \cup \{x_4\}$. By Theorem 3.2, $N[Z(W)]$ separates d_i and d_j.

Finally, assume that (H, X) is a consistent alignment. Recall that (D_1, X) is also a consistent alignment, and that (D_1, X) and (H, X) give the same order $x_1, \dots, x_7$ on X. Let R be the unique path in G from x_1 to x_7 with $R^* \subseteq H$. Then $C = d_i\text{-}x_1\text{-}R\text{-}x_7\text{-}d_j\text{-}D_1\text{-}d_i$ is a hole on more than seven vertices in G. Also, since G is (theta, prism)-free, it follows that either one of (D_1, X) and (H, X) is spiky and the other is triangular, or at least one of (D_1, X) and (H, X) is wide. In the former case, $W = (C, x_4)$ is a special wheel (see Figure 14). It follows from Theorem 3.2 that $Z(W)$ is a 6-subset of $D_2 \cup \{d_k : i+2 \leq k \leq j-2\} \cup \{x_4\}$ such that $N[Z(W)]$ separates d_i and d_j. In the latter case, $W = (C, x_4)$ is a non-special wheel (see Figure 15). Since G is $K_{1,t}$-free, it follows that $Z(W) = N_C[x_4] \subseteq D_2 \cup \{d_k : i+2 \leq k \leq j-2\} \cup \{x_4\}$ has cardinality at most $2t$. Moreover, by Theorem 3.1, $N[Z(W)]$ separates d_i and d_j. This completes the proof of Theorem 7.1. ■

We also need the following result from [3]:

Theorem 7.2 (Chudnovsky, Gartland, Hajebi, Lokshtanov, Spirkl [3])**.** *For every $m \in \mathbb{N}$ and every m-amicable graph class $\mathcal{G}$, there is a constant $f_{7.2} = f_{7.2}(\mathcal{G}, m) \in \mathbb{N}$ with the following property. Let $\mathcal{G}$ be a graph class which is m-amicable. Let $G \in \mathcal{G}$ and let w be a normal weight function on G. Then there exists $Y \subseteq V(G)$ such that*

- *$|Y| \leq f_{7.2}$, and*
- *$N[Y]$ is a w-balanced separator in G.*

Now, defining $f_{2.1}(t) = f_{7.2}(\mathcal{C}_t, \max\{2t, 7\})$ for every $t \in \mathbb{N}$, Theorem 2.1 is immediate from Theorems 7.1 and 7.2.

8. Acknowledgement

Part of this work was done when Nicolas Trotignon visited Maria Chudnovsky at Princeton University with the generous support of the H2020-MSCA-RISE project CoSP-GA No. 823748.

References

[1] T. Abrishami, B. Alecu, M. Chudnovsky, S. Hajebi, and S. Spirkl. Tree independence number I. (Even hole, diamond, pyramid)-free graphs. *J. Graph Theory* (2024), 1–21.

[2] T. Abrishami, B. Alecu, M. Chudnovsky, S. Hajebi, and S. Spirkl. Induced subgraphs and tree decompositions VIII. Excluding a forest in (theta, prism)-free graphs. *Combinatorica*, 44(5):921–948, 2024.

[3] M. Chudnovsky, P. Gartland, S. Hajebi, D. Lokshtanov, and S. Spirkl. Induced subgraphs and tree decompositions XV. Even-hole-free graphs have logarithmic treewidth. Preprint available at `https://arxiv.org/abs/2402.14211`.

[4] M. Chudnovsky, S. Hajebi, D. Lokshtanov, and S. Spirkl. Tree independence number II. Three-path-configurations. *J. Combin. Theory Ser. B*, 176:74–96, 2026.

[5] M. Chudnovsky and N. Trotignon. On treewidth and maximum cliques. *Innov. Graph Theory*, 2:223–243, 2025.

[6] C. Dallard, F. V. Fomin, P. A. Golovach, T. Korhonen, and M. Milanič. Computing Tree Decompositions with Small Independence Number. *ACM Trans. Algorithms*, 22(1), Nov. 2025.

[7] C. Dallard, M. Krnc, O. Kwon, M. Milanič, A. Munaro, K. Štorgel, and S. Wiederrecht. Treewidth versus clique number. IV. Tree-independence number of graphs excluding an induced star. Preprint available at `https://arxiv.org/pdf/2402.11222`.

[8] C. Dallard, M. Milaniˇ c, and K. ˇ Storgel. Treewidth versus clique number. II. Tree-independence number. *J. Combin. Theory Ser. B*, 164:404–442, 2024.

[9] P. Erdös and G. Szekeres. A combinatorial problem in geometry. *Compositio Math.*, 2:463–470, 1935.

[10] M. C. Golumbic. Algorithmic aspects of perfect graphs. In *Topics on perfect graphs*, volume 88 of *North-Holland Math. Stud.*, pages 301–323. North-Holland, Amsterdam, 1984.

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

[12] 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.