TREE INDEPENDENCE NUMBER
II. THREE-PATH-CONFIGURATIONS.

MARIA CHUDNOVSKY[†], SEPEHR HAJEBI [§], DANIEL LOKSHTANOV[‡], AND SOPHIE SPIRKL[§‖]

Abstract. A *three-paths-configuration* is a graph consisting of three pairwise internally-disjoint paths the union of every two of which is an induced cycle of length at least four. A graph is *3PC-free* if no induced subgraph of it is a three-paths-configuration. We prove that 3PC-free graphs have poly-logarithmic tree-independence number. More explicitly, we show that there exists a constant c such that every n-vertex 3PC-free graph graph has a tree decomposition in which every bag has stability number at most $c(\log n)^2$. This implies that the Maximum Weight Independent Set problem, as well as several other natural algorithmic problems, that are known to be NP-hard in general, can be solved in quasi-polynomial time if the input graph is 3PC-free.

1. Introduction

All graphs in this paper are finite and simple and all logarithms are base 2. Let $G = (V(G), E(G))$ be a graph. For a set $X \subseteq V(G)$ we denote by $G[X]$ the subgraph of G induced by X, and by $G \setminus X$ the subgraph of G induced by $V(G) \setminus X$. In this paper, we use induced subgraphs and their vertex sets interchangeably. For a graph G, H we say that G *contains* H if H is isomorphic to $G[X]$ for some $X \subseteq V(G)$. We say that G is H*-free* if G does not contain H. For a family $\mathcal{H}$ of graphs, we say that G is $\mathcal{H}$*-free* if G is H-free for every $H \in \mathcal{H}$.

Let $v \in V(G)$. The *open neighborhood of* v, denoted by $N_G(v)$, is the set of all vertices in $V(G)$ adjacent to v. The *closed neighborhood of* v, denoted by $N_G[v]$, is $N(v) \cup \{v\}$. Let $X \subseteq V(G)$. The *open neighborhood of* X, denoted by $N_G(X)$, is the set of all vertices in $V(G) \setminus X$ with at least one neighbor in X. The *closed neighborhood of* X, denoted by $N_G[X]$, is $N_G(X) \cup X$. When there is no danger of confusion, we omit the subscript G. Let $Y \subseteq V(G)$ be disjoint from X. We say X is *complete* to Y if all edges with an end in X and an end in Y are present in G, and X is *anticomplete* to Y if there are no edges between X and Y.

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

(i) For every $v \in V(G)$, there exists $t \in V(T)$ such that $v \in \chi(t)$.
(ii) For every $v_1v_2 \in E(G)$, there exists $t \in V(T)$ such that $v_1, v_2 \in \chi(t)$.
(iii) For every $v \in V(G)$, the subgraph of T induced by $\{t \in V(T) \mid v \in \chi(t)\}$ is connected.

Date: June 14, 2024.
[§]Department of Combinatorics and Otimization, University of Waterloo, Waterloo, Ontario, Canada.
[†] Princeton University, Princeton, NJ, USA. Supported by NSF-EPSRC Grant DMS-2120644 and by AFOSR grant FA9550-22-1-0083.
[‡] Department of Computer Science, University of California Santa Barbara, Santa Barbara, CA, USA. Supported by NSF grant CCF-2008838.
[‖] 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].

For each $t \in V(T)$, we refer to $\chi(t)$ as a *bag of* (T, χ). The *width* of a tree decomposition (T, χ), denoted by width(T, χ), is $\max_{t \in V(T)} |\chi(t)| - 1$. The *treewidth* of G, denoted by tw(G), is the minimum width of a tree decomposition of G. Treewidth, first introduced by Robertson and Seymour in the context of graph minors, is an extensively studied graph parameter, mostly due to the fact that graphs of bounded treewidth exhibit interesting structural [12] and algorithmic [3] properties.

A *stable (or independent) set* in a graph G is a set of pairwise non-adjacent vertices of G. The *stability (or independence) number* $\alpha(G)$ of G is the size of a maximum stable set in G. Given a graph G with weights on its vertices, the MAXIMUM WEIGHT INDEPENDENT SET (MWIS) problem is the problem of finding a stable set in G of maximum total weight. This problem is known to be NP-hard [8], but it can be solved efficiently (in polynomial time) in graphs of bounded treewidth. Closer examination of the algorithm motivated Dallard, Milanič and Štorgel [6] to define a related graph width parameter, specifically targeting the complexity of the MWIS problem. The *independence number* of a tree decomposition (T, χ) of G is $\max_{t \in V(T)} \alpha(G[\chi(t)])$. The *tree independence number* of G, denoted tree-$\alpha(G)$, is the minimum independence number of a tree decomposition of G. Graphs with large treewidth and small tree-α are graphs whose large treewidth can be explained by the presence of a large clique. It is shown in [6] that if a graph is given together with a tree decomposition with bounded independence number, then the MWIS problem can be solved in polynomial time. [5] presents an algorithm that constructs such tree decompositions efficiently in graphs of bounded tree-α, yielding an efficient algorithm for the MWIS problem for graphs of bounded tree-α.

A *hole* in a graph is an induced cycle of length at least four. A *path* in a graph is an induced subgraph that is a path. The *length* of a path or a hole is the number of edges in it. Given a path P with ends a, b, the *interior* of P, denoted by P^*, is the set $P \setminus \{a, b\}$.

A *theta* is a graph consisting of two distinct vertices a, b and three paths P_1, P_2, P_3 from a to b, such that $P_i \cup P_j$ is a hole for every $i, j \in \{1, 2, 3\}$. It follows that a is non-adjacent to b and the sets P_1^*, P_2^*, P_3^* are pairwise disjoint and anticomplete to each other. If a graph G contains an induced subgraph H that is a theta, and a, b are the two vertices of degree three in H, then we say that G contains a theta *with ends a and b*.

A *pyramid* is a graph consisting of a vertex a and a triangle $\{b_1, b_2, b_3\}$, and three paths P_i from a to b_i for $1 \leq i \leq 3$, such that $P_i \cup P_j$ is a hole for every $i, j \in \{1, 2, 3\}$. It follows that $P_1 \setminus a, P_2 \setminus a, P_3 \setminus a$ are pairwise disjoint, and the only edges between them are of the form $b_i b_j$. It also follows that at most one of P_1, P_2, P_3 has length exactly one. We say that a is the *apex* of the pyramid and that $b_1 b_2 b_3$ is the *base* of the pyramid.

A *generalized prism* is a graph consisting of two triangles $\{a_1, a_2, a_3\}$ and $\{b_1, b_2, b_3\}$, and three paths P_i from a_i to b_i for $1 \leq i \leq 3$, and such that $P_i \cup P_j$ is a hole for every $i, j \in \{1, 2, 3\}$. It follows that P_1^*, P_2^*, P_3^* are pairwise disjoint and anticomplete to each other, $|\{a_1, a_2, a_3\} \cap \{b_1, b_2, b_3\}| \leq 1$, and if $a_1 = b_1$, then $P_2^* \neq \emptyset$ and $P_3^* \neq \emptyset$. Moreover, the only edges between P_i and P_j are $a_i a_j$ and $b_i b_j$. A *prism* is a generalized prism whose triangles are disjoint. A *pinched prism* is a generalized prism whose triangles meet.

A *three-paths-configuration (3PC)* is a graph that is either a theta, or pyramid, or a generalized prism (see Figure 1). It is easy to check that this definition is equivalent to the one in the abstract. Let $\mathcal{C}$ be the class of (theta, pyramid, generalized prism)-free graphs; $\mathcal{C}$ is also known as the class of *3PC-free* graphs.

The following is the main result of [2]:

Theorem 1.1 ([2]). *For every integer $t > 0$ there exists a constant $c(t)$ such that for every n-vertex graph $G \in \mathcal{C}$ that contains no clique of size t,* $\text{tw}(G) \leq c(t) \log n$.

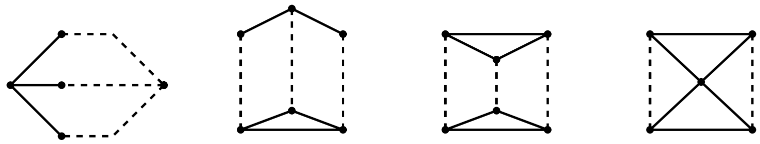

FIGURE 1. The three-path-configurations. From left to right: A theta, a pyramid, a prism and a pinched prism (dashed lines depict paths of non-zero length).

This is a strengthening of a conjecture of [14] that theta-free graphs with no 3-vertex clique have logarithmic treewidth It was also shown in [14] that there exist triangle-free graphs in $\mathcal{C}$ with arbitrarily large treewidth (in fact, treewidth logarithmic in the number of vertices), and so the bound of Theorem 1.1 is asymptotically best possible. A consequence of Theorem 1.1 is that the MWIS problem (as well as many others) can be solved in polynomial time on 3PC-free graphs with bounded clique number.

It is now natural to ask about 3PC-free graphs with no bound on the clique number. Since the complete bipartite graph $K_{2,3}$ is a theta, and therefore is forbidden in graphs in $\mathcal{C}$, one would expect these graphs to behave well with respect to tree-α. Our main result here confirms this. We prove:

Theorem 1.2. *There exists a constant c such that for every integer $n > 1$ every n-vertex graph $G \in \mathcal{C}$ has tree independence number at most $c(\log n)^2$.*

Note that since the class of theta-free graphs is "χ-bounded" (see [13] for details), Theorem 1.2 yields a weakening of Theorem 1.1, that for every integer $t > 0$, there exists a constant $c(t)$ such that for every n-vertex graph $G \in \mathcal{C}$ that contains no clique of size t, $\text{tw}(G) \leq c(t)(\log n)^2$. On the other hand, since the only construction of 3PC-graphs with large treewidth known to date is the construction of [14] where all graphs have clique number at most four, we do not know if the bound of Theorem 1.2 is asymptotically tight, or whether it can be made linear in $\log n$ (in which case, it would imply Theorem 1.1).

Another result in this paper that may be of independent interest is the following:

Theorem 1.3. *Let $G \in \mathcal{C}$ with $|V(G)| = n$, and let $a, b \in V(G)$ be non-adjacent. Then there is a set $X \subseteq V(G) \setminus \{a, b\}$ with $\alpha(X) \leq 32 \log n$ and such that every component of $G \setminus X$ contains at most one of a, b.*

1.1. **Proof outline and organization.** The proof of Theorem 1.2 follows an outline similar to [4], but requires several new techniques and ideas. We sketch it in this subsection, postponing the precise definitions for later. We begin by exploring the effect that the presence of "useful wheels" has on 3PC-free graphs, and show that every useful wheel can be broken by a cutset that is contained in the union of the neighborhoods of three vertices. This is done in Section 2.

For a graph G a function $w : V(G) \to [0, 1]$ is a *normal weight function* on G if $w(V(G)) = 1$. Let $c \in [0, 1]$ and let w be a normal weight function on G. A set $X \subseteq V(G)$ is a *(w, c)-balanced separator* if $w(D) \leq c$ for every component D of $G \setminus X$. The set X is a *w-balanced separator* if X is a $(w, \frac{1}{2})$-balanced separator. We show:

Theorem 1.4. *There is an integer d with the following property. Let $G \in \mathcal{C}$, and let w be a normal weight function on G. Then there exists $Y \subseteq V(G)$ such that*

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

This is done in Section 3; the proof is similar to the proof of an analogous statement in [4].

In Section 4 we prove Theorem 1.3. The key insight here is that a stronger result can (and should) be proved, showing that every two "cooperative subgraphs", disjoint and anticomplete to each other, can be separated by removing a set with logarithmic stability number. The proof of this strengthening follows by relatively standard structural analysis.

In Section 6 we develop a technique that uses results of Section 3 and Section 4 and produces a balanced separator of small stability number in a graph. This technique does not depend on the particular graph-class in question, but only on the validity of statements similar to Theorems 3.1 and 4.1. We also rely on a lemma from Section 5, which is proved here for theta-free graphs, but can be generalized in several ways. Section 6 is completely different from [4], and requires several new ideas.

In Section 7 we deduce Theorem 1.2 from the building blocks developed so far. We finish with Section 8 discussing the algorithmic implications of Theorem 1.2.

2. Structural results

In this section we prove a theorem asserting the existence of certain cutsets in graphs in $\mathcal{C}$.

Let G be a graph. Let $X, Y, Z \subseteq V(G)$. We say that X *separates* Y from Z if no component of $G \setminus X$ meets both Y and Z. Let W be a hole in G and $v \in G \setminus W$. A *sector* of (W, v) is a path P of W of length at least one, such that both ends of P are adjacent to v, and v is anticomplete to P^*. A sector P is *long* if $P^* \neq \emptyset$. A *useful wheel* in G is a pair (W, v) where W is a hole of length at least seven and (W, v) has at least two long sectors. We prove:

Theorem 2.1. *Let $G \in \mathcal{C}$ and let (W, v) be a useful wheel in G. Let S be a long sector of W with ends s_1, s_2. Then $((N(s_1) \cup N(s_2)) \setminus W) \cup N(v)$ separates S^* from $W \setminus S$.*

Proof. Let $X = ((N(s_1) \cup N(s_2)) \setminus W) \cup N(v)$ Suppose for a contradiction that there is a component of $G \setminus X$ intersecting both S^* and $W \setminus S$. It follows that there is a path $P = p_1\text{-}\dots\text{-}p_k$ in $G \setminus X$, possibly with $k = 1$, such that p_1 has a neighbor in S^* and p_k has a neighbor in $W \setminus S$. In particular, P is disjoint from and anticomplete to $\{s_1, s_2, v\}$.

Choose P with $|P| = k$ as small as possible. It follows that

- we have $P \subseteq G \setminus (W \cup X)$;
- P^* is anticomplete to $W \cup \{v\}$;
- if $k > 1$, then p_1 is anticomplete to $W \setminus S^*$ and p_k is anticomplete to S.

Let s_1, s_2 be the ends of S, and let t_1 and t_2 be the (unique) neighbors of s_1 and s_2 in $W \setminus S^*$, respectively. Since (W, v) has at least two long sectors, it follows that s_1, s_2, t_1, t_2 are all distinct, and that $W \setminus N[S] \neq \emptyset$. In particular, since $W \cup \{v\}$ is not a pyramid and (W, v) is not a pinched prism, it follows that v has a neighbor in $w \in W \setminus N[S]$.

Traversing S from s_1 to s_2, let u_1 and u_2 be the first and the last neighbor of p_1 in S, respectively. It follows that $u_1, u_2 \in S^*$. We deduce:

(1) $N(p_k) \cap (W \setminus S) \subseteq \{t_1, t_2\}$.

Suppose not. Then there is a path Q in G from p_k to v such that $Q^* \subseteq W \setminus N[S]$. Assume that $u_1 = u_2$. Then there is a theta in G with u_1, v and paths $u_1\text{-}S\text{-}s_1\text{-}v$, $u_1\text{-}S\text{-}s_2\text{-}v$, $u_1\text{-}p_1\text{-}P\text{-}p_k\text{-}Q\text{-}v$. Next, assume that u_1 and u_2 are distinct and non-adjacent. Then there is a theta in G with ends p_1, v and paths $p_1\text{-}u_1\text{-}S\text{-}s_1\text{-}v$, $p_1\text{-}u_2\text{-}S\text{-}s_2\text{-}v$, $p_1\text{-}P\text{-}p_k\text{-}Q\text{-}v$. Since G is theta-free, it follows that u_1, u_2

are distinct and adjacent. But now there is a pyramid in G with apex v, base $p_1u_1u_2$ and paths p_1-P-p_k-Q-v, u_1-S-s_1-v, u_2-S-s_2-v, a contradiction. This proves (1).

(2) *We have $u_1 = u_2$.*

Suppose not. By (1) and without loss of generality, we may assume that p_k is adjacent to t_1. Assume first that t_1 and v are not adjacent. If u_1 and u_2 are not adjacent either, then there is a theta in G with ends p_1, s_1 and paths p_1-u_1-S-s_1, p_1-u_2-S-s_2-v-s_1, p_1-P-p_k-t_1-s_1, and if u_1, u_2 are adjacent, then there is a pyramid in G with apex s_1, base $p_1u_1u_2$ and paths p_1-P-p_k-t_1-s_1, u_1-S-s_1, u_2-S-s_2-v-s_1. Since G is (theta, pyramid)-free, it follows that t_1 and v are adjacent. Assume that u_1 and u_2 are not adjacent. Then there is a pyramid in G with apex p_1, base s_1t_1v and paths paths s_1-S-u_1-p_1, t_1-p_k-P-p_1, v-s_2-S-u_2-p_1. Again, since G is pyramid-free, it follows that u_1, u_2 are adjacent. But now there is a prism in G with triangles $u_1p_1u_2$, s_1t_1v and paths u_1-S-s_1, p_1-P-p_k-t_1, u_2-S-s_2-v, a contradiction. This proves (2).

Henceforth, let $u = u_1 = u_2$. It follows that:

(3) *We have $k = 1$.*

Suppose that $k > 1$. Since $W \cup \{p_k\}$ is not a theta with ends t_1, t_2, we may assume by (1) and without loss of generality, that $N(p_k) \cap (W \setminus S) = \{t_1\}$. But then $W \cup P$ is a theta in G with ends u, t_1, a contradiction. This proves (3).

Henceforth, let $p = p_1 = p_k$. Since $W \cup \{p\}$ is not a theta with one end u and the other end in $\{t_1, t_2\}$, (1) implies that $N(p) \cap W = \{t_1, t_2, u\}$. Recall that v has a neighbor $w \in W \setminus N[S]$.

(4) *We have $t_1w \in E(G)$ and $t_2v \notin E(G)$. Similarly, we have $t_2w \in E(G)$ and $t_1v \notin E(G)$.*

Suppose not. Then we may assume, without loss of generality, that either t_1 and w are not adjacent, or t_2 and v are adjacent. In either case, it follows that there is a path Q in G from p to v such that $t_2 \in P^* \subseteq W \setminus (S \cup N(t_1))$. Now there is a theta in G with ends p, s_1 and paths p-u-S-s_1, p-t_1-s_1, p-Q-v-s_1, a contradiction. This proves (4).

We will now finish the proof. From (4), it follows that $W \setminus S = t_1$-w-t_2 and $N(v) \cap W = \{s_1, s_2, w\}$. Recall also that $N(p) \cap W = \{t_1, t_2, u\}$. Since $|W| > 6$, it follows either s_1-S-u or s_2-S-u, say the former, has non-empty interior. But then there is a theta in G with ends s_1, u and paths s_1-S-u, s_1-t_1-p-u and s_1-v-s_2-S-u, a contradiction. ■

3. Dominated balanced separators

The goal of this section is to prove the following:

Theorem 3.1. *There is an integer d with the following property. Let $G \in \mathcal{C}$ and let w be a normal weight function on G. Then there exists $Y \subseteq V(G)$ such that*

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

We follow the outline of the proof of Theorem 8.1 in [4]. First we repeat several definitions from [4]. Let G be a graph, let $P = p_1$-$\dots$-p_n be a path in G and let $X = \{x_1, \dots, x_k\} \subseteq G \setminus P$. We say that (P, X) is an *alignment* if $N_P(x_1) = \{p_1\}$, $N_P(x_k) = \{p_n\}$, every vertex of X has a neighbor in P, and there exist $1 < j_2 < \dots < j_{k-1} < j_k = n$ such that $N_P(x_i) \subseteq p_{j_i}$-$P$-$p_{j_{i+1}-1}$ for $i \in \{2, \dots, k-1\}$. We also say that $x_1, \dots, x_k$ is *the order on X given by the alignment* (P, X). An alignment (P, X) is *wide* if each of $x_2, \dots, x_{k-1}$ has two non-adjacent neighbors in P, *spiky* if each of $x_2, \dots, x_{k-1}$ has a unique neighbor in P and *triangular* if each of $x_2, \dots, x_{k-1}$ has exactly two neighbors in P and they are adjacent. An alignment is *consistent* if it is wide, spiky or triangular.

The first step in the proof of Theorem 3.1 is the following:

Theorem 3.2. *For every integer $x \geq 6$, there exists an integer $\sigma = \sigma(x) \geq 1$ with the following property. Let $G \in \mathcal{C}$ and assume that $V(G) = D_1 \cup D_2 \cup Y$ where*

- *Y is a stable set with $|Y| = \sigma$,*
- *D_1 and D_2 are components of $G \setminus Y$,*
- *$N(D_1) = N(D_2) = Y$,*
- *$D_1 = d_1\text{-}\cdots\text{-}d_k$ is a path, and*
- *for every $y \in Y$ there exists $i(y) \in \{1, \dots, k\}$ such that $N(d_{i(y)}) \cap Y = \{y\}$.*

Then there exist $X \subseteq Y$ with $|X| = x$ and $H_2 \subseteq D_2$ such that

1. *(D_1, X) is a consistent alignment.*
2. *One of the following holds.*
   - *We have $|H_2| = 1$ (so $H_2 \cup X$ is a star), and (D_1, X) is wide.*
   - *(H_2, X) is a consistent alignment, the orders given on X by (D_1, X) and by (H_2, X) are the same, and at least one of (D_1, X) and (H_2, X) is wide.*

The proof of Theorem 3.2 requires two preliminary results. The first one is Theorem 3.3 below from [4]. Following [4], by a *caterpillar* we mean a tree C with maximum degree three such that there exists a path P of C where all branch vertices of C belong to P. (Our definition of a caterpillar is non-standard for two reasons: a caterpillar is often allowed to be of arbitrary maximum degree, and a spine often contains all vertices of degree more than one.) A *claw* is the graph $K_{1,3}$. For a graph H, a vertex v of H is said to be *simplicial* if $N_H(v)$ is a clique.

Theorem 3.3 (Chudnovsky, Gartland, Hajebi, Lokshtanov, Spirkl; Theorem 5.2 in [4])**.** *For every integer $h \geq 1$, there exists an integer $\mu = \mu(h) \geq 1$ with the following property. Let G be a connected graph. Let $Y \subseteq G$ such that $|Y| \geq \mu$, $G \setminus Y$ is connected and every vertex of Y has a neighbor in $G \setminus Y$. Then there is a set $Y' \subseteq Y$ with $|Y'| = h$ and an induced subgraph H of $G \setminus Y$ for which one of the following holds.*

- *H is a path and every vertex of Y' has a neighbor in H.*
- *H is a caterpillar, or the line graph of a caterpillar, or a subdivided star or the line graph of a subdivided star. Moreover, every vertex of Y' has a unique neighbor in H and every vertex of $H \cap N(Y')$ is simplicial in H.*

The second one is:

Lemma 3.4. *Let $c, x \geq 1$ be integers. Let G be a theta-free graph and assume that $V(G) = D_1 \cup D_2 \cup Y$ where*

- *Y is a stable set with $|Y| = 3x(c+2)$;*
- *D_1 and D_2 are components of $G \setminus Y$;*
- *$N(D_1) = N(D_2) = Y$;*
- *D_1 is a path; and*
- *for every $d \in D_1$, we have $|N(d) \cap Y| \leq c$.*

Then there exist $X \subseteq Y$ with $|X| = x$ such that (D_1, X) is a consistent alignment.

Proof. For every vertex $y \in Y$, let P_y be the path in D_1 such that y is complete to the ends of P_y and anticomplete to $D_1 \setminus P_y$. 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 by [9] I is perfect. Since $|V(I)| = 3x(c+2)$, we deduce that I contains either a clique of cardinality $c+2$ of a stable set of cardinality $3x$.

Assume that I contains a clique of cardinality $c+2$. Then there exists $C \subseteq Y$ with $|C| = c+2$ and $d \in D_1$ such that $d \in P_y$ for every $y \in C$. It follows that for every $y \in C$, either y is adjacent to d, or $D_1 \setminus d$ has two components and y has a neighbor in each of them. Since $|N(d) \cap Y| \leq c$, we deduce that there are two vertices $y, y' \in C \subseteq Y$ as well as two paths P_1 and P_2 from y to y' with disjoint and anticomplete interiors contained in D_1. On the other hand, since D_2 is connected and $N(D_2) = Y$, it follows that there is a path P_3 in G from y to y' whose interior is contained in D_2. But now there is a theta in G with ends y, y' and paths P_1, P_2, P_3, a contradiction.

We deduce that I contains a stable set S of cardinality $3x$. From the definition of I, it follows that (D_1, S) is an alignment. Hence, there exists $X \subseteq S \subseteq Y$ with $|X| = x$ such that (D_1, S) is a consistent alignment. This completes the proof of Lemma 3.4. ■

We are now ready to prove Theorem 3.2:

Proof of Theorem 3.2. Let $\sigma(x) = 18\mu(3(x^2+1)(x+1))$, where $\mu(\cdot)$ comes from Theorem 3.3. We begin with the following:

(5) *Every vertex in D_1 has at most four neighbors in Y.*

Suppose for a contradiction that for some $i \in \{1, \dots, k\}$, there is a subset $Z \subseteq Y$ of cardinality five such that d_i is complete to Z. It follows that for every $y \in Z$, we have $i(y) \neq i$, and so there is 3-subset T of Z such that either $i(y) < i$ for all $y \in T$ or $i < i(y)$ for all $y \in T$. Consequently, there are two distinct vertices $y, y' \in T \subseteq Z \subseteq Y$ for which d_i is disjoint from and anticomplete to $d_{i(y)}$-D_1-$d_{i(y')}$. On the other hand, since D_2 is connected and $N(D_2) = Y$, it follows that there is a path Q in G from y to y' whose interior is contained in D_2. But now there is a theta in G with ends y, y' and paths y-$d_{i(y)}$-D_1-$d_{i(y')}$-y', y-d_i-y', Q, a contradiction. This proves (5).

From (5), Lemma 3.4 and the choice of $\sigma(x)$, it follows that:

(6) *There exists $Y_1 \subseteq Y$ with $|Y_1| = \mu(3(x^2+1)(x+1))$ such that (D_1, Y_1) is consistent alignment.*

Henceforth, let Y_1 be as in (6). Since $G_1 = G[Y_1 \cup D_2]$ and $G_1 \setminus Y_1 = D_2$ are both connected, we can apply Theorem 3.3 to G_1 and Y_1. It follows that there is a set $Y' \subseteq Y_1$ with $|Y'| = 3(x^2+1)(x+1)$ and an induced subgraph H of $G_1 \setminus Y_1 = D_2$ for which one of the following holds.

- H is a path and every vertex of Y' has a neighbor in H.
- H is a caterpillar, or the line graph of a caterpillar, or a subdivided star or the line graph of a subdivided star. Moreover, every vertex of Y' has a unique neighbor in H and every vertex of $H \cap N(Y')$ is simplicial in H.

Assume that the second bullet above holds. By (6), (D_1, Y') is a consistent alignment. But then it is straightforward to observe that G contains either a theta, a prism or a pyramid, a contradiction. It follows that H is indeed a path and every vertex of Y' has a neighbor in H.

Now, assume that some vertex in $z \in H$ has at least x neighbors in Y'. Choose $X \subseteq N(z) \cap Y' \subseteq Y$ with $|X| = x$. Let $H_2 = \{x\}$. By (6), (D_1, X) is a consistent alignment. Note that if (D_1, X) is spiky, then $D_1 \cup X \cup \{z\}$ contains a theta, and if (D_1, X) is triangular, then $D_1 \cup X \cup \{z\}$ contains a pyramid. Therefore, (D_1, X) is wide. But now X and H_2 satisfy Theorem 3.2.

Therefore, we may assume that every vertex in H has fewer than x neighbors in Y'. Let $H_2 = H$. Since $|Y'| = 3(x^2+1)(x+1)$, it follows from Lemma 3.4 that there exists $X' \subseteq Y'$ with $|X'| = x^2+1$ such that (H_2, X') is a consistent alignment. Also, by (6), (D_1, X') is a consistent alignment. This, along with the Erdős-Szekeres Theorem [7], implies that there exists $X \subseteq X' \subseteq Y' \subseteq Y$ with $|X| = x$ such that both (D_1, X) and (H_2, X) are consistent alignments, and the orders given on X by (D_1, X) and (H_2, X) are the same. Moreover, since G is (theta, pyramid, pinched prism)-free, it follows that at least one of (D_1, X) and (H_2, X) is wide. Hence, X and H_2 satisfy Theorem 3.2. This completes the proof. ■

Now, as in [4], we will show that the class $\mathcal{C}$ is "amiable" and "amicable", and then use Theorem 8.5 of [4] to complete to the proof. The details are below. In [4], a graph class $\mathcal{G}$ is said to be *amiable* if, under the same assumptions as that of Theorem 3.2 for a graph $G \in \mathcal{G}$, there exists $X \subseteq Y$ with $|X| = x$ and $H_2 \subseteq D_2$ satisfying one of several possible outcomes. In particular, the outcome of Theorem 3.2 is one of the possible outcomes in the definition of an amiable class. Therefore, from Theorem 3.2, we deduce that:

Corollary 3.5. *The class $\mathcal{C}$ is amiable.*

Following [4], for an integer $m > 0$, a graph class $\mathcal{G}$ is said to be *m-amicable* if $\mathcal{G}$ is amiable, and the following holds for every $G \in \mathcal{G}$. Let σ, $Y, D_1 = d_1\text{-}\cdots\text{-}d_k$ and D_2 be as in the definition of an amiable class (and so as in Theorem 3.2) with $|Y| = \sigma(7)$. Let X and $H_2 \subseteq D_2$ be as the conclusion of Theorem 3.2. Let $\{x_1, \dots, x_7\}$ be the order given on X by (D_1, X). Let i be maximum such that x_1 is adjacent to d_i, and let j be minimum such that x_7 is adjacent to d_j. Then there exists a subset $Z \subseteq D_2 \cup \{d_{i+2}, \dots, d_{j-2}\} \cup \{x_4\}$ with $|Z| \leq m$ such that $N[Z]$ separates d_i from d_j. Consequently, $N[Z]$ separates $d_1\text{-}D_1\text{-}d_i$ from $d_j\text{-}D_1\text{-}d_k$. We prove:

Theorem 3.6. *The class $\mathcal{C}$ is 3-amicable.*

Proof. By Corollary 3.5, $\mathcal{C}$ is an amiable class. With same notation as in the definition of a 3-amicable class, our goal is to show that there exists a subset $Z \subseteq D_2 \cup \{d_{i+2}, \dots, d_{j-2}\} \cup \{x_4\}$ with $|Z| \leq 3$ such that $N[Z]$ separates d_i from d_j. Consequently, $N[Z]$ separates $d_1\text{-}D_1\text{-}d_i$ from $d_j\text{-}D_1\text{-}d_k$.

Let $l \in \{1, \dots, k\}$ be minimum such that x_4 is adjacent to d_l, and let $m \in \{1, \dots, k\}$ be maximum such that x_4 is adjacent to d_m. It follows that $i + 2 < l \leq m < j - 2$.

Let R be the (unique) path in H_2 with ends r_1, r_2 (possibly $r_1 = r_2$) such that x_1 is adjacent to r_1 and anticomplete to $R \setminus \{r_1\}$ and x_7 is adjacent to r_2 and anticomplete to $R \setminus \{r_2\}$. Then x_4 has a neighbor in R. Traversing R from r_1 to r_2, let z_1 and z_2 be the first and the last neighbor of x_4 in R.

Let $W = d_i\text{-}D_1\text{-}d_j\text{-}x_7\text{-}r_2\text{-}R\text{-}r_1\text{-}x_1\text{-}d_i$. Then W is a hole in G and $|W| \geq 7$. Since X and H_2 are outcomes of Theorem 3.2, it follows that (W, x_4) is a useful wheel in G. In particular, $S = d_l\text{-}D_1\text{-}d_i\text{-}x_1\text{-}r_1\text{-}R\text{-}z_1$ and $S' = d_m\text{-}D_1\text{-}d_j\text{-}x_7\text{-}r_2\text{-}R\text{-}z_2$ are two long sector of (W, x_4)

Note that d_1 and z_1 are the ends of the sector S from (W, x_4). Let $Z = \{x_4, d_l, z_1\}$. Then we have $Z \subseteq D_2 \cup \{d_{i+2}, \dots, d_{j-2}\} \cup \{x_4\}$, $d_i \in S^* \setminus N[Z]$ and $d_j \in W \setminus (S \cup N[Z])$. Hence, by Theorem 2.1, $N[Z]$ separates d_i from d_j, as desired. ■

The following is a restatement of Theorem 8.5 of [4]:

Theorem 3.7 (Chudnovsky, Gartland, Hajebi, Lokshtanov, Spirkl [4])**.** *For every integer $m > 0$ and every m-amicable graph class $\mathcal{G}$, there is an integer $d > 0$ with the following property. Let $\mathcal{G}$ be a graph class which is m-amicable. Let $G \in \mathcal{C}$ and let w be a normal weight function on G. Then there exists $Y \subseteq V(G)$ such that*

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

Now Theorem 3.1 is immediate from Theorems 3.6 and 3.7.

4. Separating a pair of vertices

For $X \subseteq V(G)$, a component D of $G \setminus X$ is *full for* X if $N(D) = X$. $X \subseteq V(G)$ is a *minimal separator* in G if there exist two distinct full components for X. The goal of this section is to prove the following:

Theorem 4.1. *Let $G \in \mathcal{C}$ with $|V(G)| = n$, and let $a, b \in V(G)$ be non-adjacent. Then there is a set $X \subseteq V(G) \setminus \{a, b\}$ with $\alpha(X) \leq 16 \times 2 \log n$ and such that every component of $G \setminus X$ contains at most one of a, b.*

We need the following result from [1].

Lemma 4.2. *Let x_1, x_2, x_3 be three distinct vertices of a graph G. Assume that H is a connected induced subgraph of $G \setminus \{x_1, x_2, x_3\}$ such that $V(H)$ contains at least one neighbor of each of x_1, x_2, x_3, and that $V(H)$ is minimal subject to inclusion. Then, one of the following holds:*

(i) For some distinct $i, j, k \in \{1, 2, 3\}$, there exists P that is either a path from x_i to x_j or a hole containing the edge $x_i x_j$ such that
  - *$V(H) = V(P) \setminus \{x_i, x_j\}$; and*
  - *either x_k has two non-adjacent neighbors in H or x_k has exactly two neighbors in H and its neighbors in H are adjacent.*

(ii) There exists a vertex $a \in V(H)$ and three paths P_1, P_2, P_3, where P_i is from a to x_i, such that
  - *$V(H) = (V(P_1) \cup V(P_2) \cup V(P_3)) \setminus \{x_1, x_2, x_3\}$;*
  - *the sets $V(P_1) \setminus \{a\}$, $V(P_2) \setminus \{a\}$ and $V(P_3) \setminus \{a\}$ are pairwise disjoint; and*
  - *for distinct $i, j \in \{1, 2, 3\}$, there are no edges between $V(P_i) \setminus \{a\}$ and $V(P_j) \setminus \{a\}$, except possibly $x_i x_j$.*

(iii) There exists a triangle $a_1 a_2 a_3$ in H and three paths P_1, P_2, P_3, where P_i is from a_i to x_i, such that
  - *$V(H) = (V(P_1) \cup V(P_2) \cup V(P_3)) \setminus \{x_1, x_2, x_3\}$;*
  - *the sets $V(P_1)$, $V(P_2)$ and $V(P_3)$ are pairwise disjoint; and*
  - *for distinct $i, j \in \{1, 2, 3\}$, there are no edges between $V(P_i)$ and $V(P_j)$, except $a_i a_j$ and possibly $x_i x_j$.*

For a graph G and two subsets $X, Y \subseteq V(G)$ we define *the distance in G between X and Y* as the length (number of edges) of the shortest path of G with one end in X and the other in Y. We denote the distance between X and Y by $\text{dist}_G(X, Y)$. Thus X and Y are disjoint if and only if $\text{dist}_G(X, Y) > 0$, and X, Y are anticomplete to each other if and only if $\text{dist}_G(X, Y) > 1$. In order to prove Theorem 4.1 we will prove a stronger statement. Let $H \subseteq G$. We denote by $\delta_G(H)$ the set of vertices of H that have a neighbor in $G \setminus H$ (so $\delta(H) = N(G \setminus H)$). We say that H is *cooperative* if one of the following holds:

- H is a clique, or
- $N_H(H \setminus \delta(H)) = \delta(H)$ and $H \setminus \delta(H)$ is connected.

The following lemma summarizes the property of cooperative subgraphs that is of interest to us.

Lemma 4.3. *Let $H \subseteq G$ be cooperative and let $\{n_1, n_2, n_3\}$ be a stable set in $N(H)$. Assume that there exist distinct $h_1, h_2, h_3 \in \delta(H)$ such that $n_i h_j$ is an edge if and only if $i = j$. Then there is $K \subseteq H \cup \{n_1, n_2, n_3\}$ such that*

(1) K is a subdivided claw or the line graph of a subdivided claw
(2) $\{n_1, n_2, n_3\}$ is the set of simplicial vertices of K

Proof. If $\{h_1, h_2, h_3\}$ is a triangle, then $\{n_1, n_2, n_3, h_1, h_2, h_3\}$ is the line graph of a subdivided claw and the lemma holds. This we may assume that at least one pair h_ih_j is non-adjacent, and in particular H is not a clique. It follows that $N_H(H \setminus \delta(H)) = \delta(H)$ and $H \setminus \delta(H)$ is connected.

Next suppose that h_1h_2 and h_2h_3 are edges. Then h_1h_3 is not an edge, and $\{n_1, n_2, n_3, h_1, h_2, h_3\}$ is a subdivided claw and the lemma holds. Thus we may assume that at most one of the pairs h_ih_j is an edge.

Suppose h_1h_3 is an edge. Let $P = p_1\text{-}\dots\text{-}p_k$ be path such that $p_1 = h_2$, p_k has a neighbor in $\{h_1, h_3\}$, and $P \setminus p_1 \subseteq H \setminus \delta(H)$; choose P with k as small as possible. If p_k is adjacent to exactly one of h_1, h_3, then $P \cup \{h_1, h_2, h_3, n_1, n_2, n_3\}$ is a subdivided claw; and if p_k is adjacent to both h_1 and h_3, then $P \cup \{h_1, h_2, h_3, n_1, n_2, n_3\}$ is the line graph of a subdivided claw; in both cases the lemma holds. Thus we may assume that $\{h_1, h_2, h_3\}$ is a stable set.

Let R be a minimal connected subgraph of $H \setminus \delta(H)$ such that each of h_1, h_2, h_3 has a neighbor in R. We apply Lemma 4.2. Suppose that the first outcome holds; we may assume that R is path from h_1 to h_2. If h_3 has two non-adjacent neighbors in R, then $R \cup \{h_1, h_2, h_3, n_1, n_2, n_3\}$ contains a subdivided claw; and if h_3 has exactly two neighbors in R and they are adjacent, then $R \cup \{h_1, h_2, h_3, n_1, n_2, n_3\}$ is the line graph of a subdivided claw; in both cases the theorem holds. If the second outcome of Lemma 4.2 holds, then $R \cup \{h_1, h_2, h_3, n_1, n_2, n_3\}$ is a subdivided claw; and if the third outcome of Lemma 4.2 holds, then $R \cup \{h_1, h_2, h_3, n_1, n_2, n_3\}$ is the line graph of a subdivided claw. Thus in all cases the lemma holds. ■

We also need the following:

Lemma 4.4. *Let $G \in \mathcal{C}$ and let H_1, H_2 be cooperative subgraphs of G, disjoint and anticomplete to each other. Then $\alpha(N(H_1) \cap N(H_2)) < 17$.*

Proof. Suppose there is a stable set $N \subseteq N(H_1) \cap N(H_2)$ with $|N| = 17$. Suppose first that some vertex $h_1 \in H_1$ has at least five neighbors in N; let $n_1, \dots, n_5 \in N \cap N(h_1)$. If some $h_2 \in H_2$ has three three neighbors in $\{n_1, \dots, n_5\}$, say $\{n_1, n_2, n_3\}$, then $\{h_1, h_2, n_1, n_2, n_3\}$ is a theta with ends h_1, h_2, a contradiction. So no such h_2 exists. It follows that there exist $h'_1, h'_2, h'_3 \in H_2$ and $n'_1, n'_2, n'_2 \in \{n_1, \dots, n_5\}$ such that $h'_in'_j$ is an edge if and only if $i = j$. By Lemma 4.3 there exists $K \subseteq H_2 \cup \{n'_1, n'_2, n'_3\}$ such that K is a subdivided claw or the line graph of a subdivided claw, and $\{n'_1, n'_2, n'_3\}$ is the set of simplicial vertices of K. But now $K \cup h_1$ is a theta or a pyramid in G, a contradiction.

It follows that for every $h_1 \in H_1$, $|N(h_1) \cap N| \leq 4$. Since $|N| = 17$ and $N \subseteq N(H_1)$, there exist $h_1, \dots, h_5 \in H$, and $n_1, \dots, n_5 \in N$ such that h_in_j is an edge if and only if $i = j$.

By renumbering $n_1, \dots, n_5$ if necessary we may assume that one of the following holds:

- there exists $k \in H_2$ such that $n_1, n_2, n_3 \in N(k)$; in this case set $K_2 = \{k, n_1, n_2, n_3\}$, or
- there exist $k_1, k_2, k_3 \in H_2$ such that k_in_j is an edge if and only if $i = j$. In this case let $K_2 \subseteq H_2 \cup \{n_1, n_2, n_3\}$ be such that K_2 is a subdivided claw or the line graph of a subdivided claw and $\{n_1, n_2, n_3\}$ is the set of simplicial vertices of K_2 (such K_2 exists by Lemma 4.3).

By Lemma 4.3 there exists $K_1 \subseteq H_1 \cup \{n_1, n_2, n_3\}$ such that K_1 is a subdivided claw or the line graph of a subdivided claw, and $\{n_1, n_2, n_3\}$ is the set of simplicial vertices of K_1. But now $K_1 \cup K_2$ is a theta, a pyramid or a prism in G, a contradiction. ■

We will now prove the following strengthening of Theorem 4.1

Theorem 4.5. *Let $G \in \mathcal{C}$ with $|V(G)| = n$, and let H_1, H_2 be cooperative subgraphs of G, disjoint and anticomplete to each other. Then there is a set $X \subseteq V(G) \setminus (H_1 \cup H_2)$ with $\alpha(X) \leq 16 \times 2\log(n + 1 - |H_1| - |H_2|)$ and such that X separates H_1 from H_2.*

Proof. Write $G_1 = G$, $H_2^1 = H_2$ and $N_1 = N_{G_1}(H_1) \cap N_{G_1}(H_2)$. Define $G_2 = G_1 \setminus N_1$, $H_2^2 = N_{G_2}[H_2^1]$ and $N_2 = N_{G_2}(H_1) \cap N_{G_2}(H_2^2)$. Let $G_3 = G_2 \setminus N_2$.

(7) $\text{dist}_{G_3}(H_1, H_2) > 3$.

Let $P = p_1\text{-}\dots\text{-}p_k$ be a shortest path in G_3 from H_1 to H_2. Then P is a path in G, $p_1 \in H_1$, $p_k \in H_2$, $P \setminus p_1$ is anticomplete to H_1, and $P \setminus p_k$ is anticomplete to H_2. Since H_1 is anticomplete to H_2, it follows that $k \geq 3$. If $k = 3$, then $p_1 \in N_1$, a contradiction. Suppose $k = 4$. Then $p_3, p_4 \notin N_1$. It follows that $p_3 \in N_{G_2}(H_2) = H_2^2$, and therefore $p_4 \in N_2$; again a contradiction. This proves that $k > 4$, and (7) follows.

It follows immediately from the definition of a cooperative subgraph that:

(8) *For $i \in \{1, 2\}$, H_1 and H_2^i are both cooperative in G_i.*

Now Lemma 4.3 implies that

(9) $\alpha(N_i) < 17$ *for every $i \in \{1, 2\}$.*

If H_1 and H_2 belong to different components of G_3, then $N_1 \cup N_2$ separates H_1 from H_2 in G. Since by (9) $\alpha(N_1 \cup N_2) \leq 16 \times 2$, the theorem holds. Thus we may assume that there is a component F of G_3 such that $H_1 \cup H_2 \subseteq F$.

(10) *There is a minimal separator Z in F such that $\text{dist}_F(Z, H_i) \geq 2$ for $i \in \{1, 2\}$, and there exist distinct full components F_1, F_2 for Z such that $H_i \subseteq F_i$.*

Let $X = N_F^2(H_1)$. Since $\text{dist}_F(H_1, H_2) \geq 4$, it follows that X separates H_1 from H_2 in F, and that $\text{dist}_F(H_2, X) \geq 2$. For $i \in \{1, 2\}$ let D_i be the component of $F \setminus X$ such that $H_i \subseteq D_i$. Let $Y = N(D_1)$, let D_2' be the component of $F \setminus Y$ such that $D_2 \subseteq D_2'$, and let $Z = N(D_2)$. Then $Z \subseteq Y$, and D_1, D_2' are full components for Z. Setting $F_1 = D_1$ and $F_2 = D_2'$, (10) follows.

Let Z be as in (10). We are now ready to complete the proof of the theorem. The proof is by induction on $n - |H_1| - |H_2|$. If $n - |H_1| - |H_2| = 0$, then $X = \emptyset$ works. Since $F_1 \cap F_2 = \emptyset$, we may assume that $|F_1 \setminus H_1| \leq \frac{n - |H_1| - |H_2|}{2}$. Let $F' = F_1 \cup F_2 \cup Z$. Let $H_2' = F_2 \cup Z$. Then $\delta_{F'}(H_2') = Z$ and H_2' is cooperative in F'. Since $\text{dist}_F(H_1, Z) \geq 2$, we have that H_1 is anticomplete to H_2'. Also

$$|F'| - |H_1| - |H_2'| \leq |F_1| - |H_1| - 1 \leq \frac{n - |H_1| - |H_2|}{2} - 1.$$

Inductively, there exists $X' \subseteq F' \setminus (H_1 \cup H_2')$ with

$$\alpha(X') \leq 16 \times 2\log(\frac{n - |H_1| - |H_2|}{2}) \leq 16 \times 2\log(n - |H_1| - |H_2| + 1) - 16 \times 2.$$

Let $X = X' \cup N_1 \cup N_2$. Then X separates H_1 from H_2 in G. By (9) $\alpha(X) \leq 16 \times 2\log(n - |H_1| - |H_2| + 1)$ as required. ∎

5. Large stable subsets in neighborhoods

In this section we prove a statement about theta-free graphs which we expect to use in future papers.

Lemma 5.1. *Let G be a theta-free graph. Let $c \geq 2$ be integer, and let Y be a set with $\alpha(Y) > 24c^2$. Let Z be the set of all vertices such that $\alpha(N(z) \cap Y) > \frac{\alpha(Y)}{c}$. Then $\alpha(Z) < 2c$.*

Proof. Suppose not, and let $I \subseteq Z$ be a stable set of size $2c$. For every $z \in I$, let $J'(z)$ be a stable set in $N(z) \cap Y$ with $|J'(z)| = \frac{\alpha(Y)}{c}$.

(11) *For all distinct $z, z' \in I$, $\alpha(N[z] \cap N[z']) \leq 2$.*

Suppose that $\alpha(N[z] \cap N[z']) \geq 3$. Since z is non-adjacent to z', there exists a stable set of size three in $N(z) \cap N(z')$, and we get a theta with ends z, z', a contradiction. This proves (11).

(12) *For all distinct $z, z' \in I$, $|J'(z) \cap N(J'(z') \setminus N(z))| \leq 4$.*

Suppose not, and let $\{n_1, .., n_5\} \subseteq J'(z) \cap N(J'(z'))$. Then $\{n_1, \dots, n_5\}$ is a stable set. If some $h \in J'(z')$ has three neighbors in $\{n_1, .., n_5\}$, then we get a theta with ends z, h; so no such h exists. It follows that there exist $h_1, h_2, h_3 \in J'(z')$ such that (permuting $n_1, \dots, n_5$ if necessary) $n_i h_j$ is an edge if and only if $i = j$. But now $\{z, n_1, n_2, n_3, h_1, h_2, h_3, z'\}$ is a theta with ends z, z', again a contradiction. This proves (12).

Let $J(z) = J'(z) \setminus \bigcup_{z' \in I \setminus z}(N[z'] \cup (N(J'(z') \setminus N(z))))$. By (11) and (12) it follows that

$$|J(z)| \geq |J'(z)| - 6|I| \geq \frac{\alpha(Y)}{c} - 12c.$$

But for all distinct $z, z' \in I$ the sets $J(z), J(z')$ are disjoint and anticomplete to each other; it follows that $\bigcup_{z \in I} J(z)$ is a stable set of size $2c(\frac{\alpha(Y)}{c} - 12c)$. Consequently,

$$2c(\frac{\alpha(Y)}{c} - 12c) \leq \alpha(Y)$$

and so $\alpha(Y) \leq 24c^2$, a contradiction. ∎

6. From domination to stability

The last step in the proof of Theorem 1.2 is to transform balanced separators with small domination number into balanced separators with small stability number.

The results in this section are more general than what we need in this paper; again they are to be used in future papers in the series. Let L, d, r integers. We say that an n-vertex graph G is (L, d, r)-*breakable* if

(1) for every two disjoint and anticomplete cliques H_1, H_2 of G with $|H_1| \leq r$ and $|H_2| \leq r$, there is a set $X \subseteq G \setminus (H_1 \cup H_2)$ with $\alpha(X) \leq L$ separating H_1 from H_2, and

(2) for every normal weight function w on G and for every induced subgraph G' of G there exists a set $Y \subseteq V(G')$ with $|Y| \leq d$ such that for every component D of $G' \setminus N[Y]$, $w(D) \leq \frac{1}{2}$.

We prove:

Theorem 6.1. *Let $d > 0$ be an integer and let $C(d) = 100d^2$. Let $L, d, n, r > 0$ be integers such that $r \leq d(2 + \log n)$ and let G be an n-vertex (L, d, r)-breakable theta-free graph. Then there exists a w-balanced separator Y in G such that $\alpha(Y) \leq C(d)\lceil\frac{d(2+\log n)}{r}\rceil(2 + \log n)L$.*

We start by proving a variant of Theorem 3.1 for $(L, d, 1)$-breakable graphs.

Theorem 6.2. *Let L, d be integers, and let G be an $(L, d, 1)$-breakable graph. Let w be a normal weight function on G. Then there exist a clique K in G and a set $Y(K) \subseteq V(G) \setminus K$ such that*

- *$|K| \leq d$,*
- *$\alpha(Y(K)) \leq d^2L$ and*
- *$N[K] \cup Y(K)$ is a w-balanced separator in G.*

Proof. Since G is $(L, d, 1)$-breakable there exists $X \subseteq V(G)$ with $|X| \leq d$ such that for every component D of $G \setminus N[X]$, $w(D) \leq \frac{1}{2}$. For every pair x, x' of non-adjacent vertices of X, let $Y(x, x')$ be a set with $\alpha(Y(x, x')) \leq L$ and $Y(x, x') \cap \{x, x'\} = \emptyset$ separating x form x' in G (such a set exists since G $(L, d, 1)$-breakable). Now let

$$Y = X \cup \bigcup_{x,x' \in X \text{ nonadjacent}} Y(x, x').$$

Then $\alpha(Y) \leq \binom{d}{2}L + d \leq d^2L$. If Y is a w-balanced separator of G, set $K = \emptyset$ and $Y(K) = Y$, and the theorem holds. Thus we may assume that there is a component D of $G \setminus Y$ with $w(D) > \frac{1}{2}$. Let $K \subseteq X$ be the set of vertices of X with a neighbor in D. Since every two non-adjacent vertices of X are separated by Y, it follows that K is a clique.

We claim that $N[K] \cup Y$ is a w-balanced separator in G. Suppose not, and let D' be the component of $G \setminus (N[K] \cup Y)$ with $w(D') > \frac{1}{2}$. Then $D' \subseteq D$. But $D \cap N(X) \subseteq D \cap N(K) \subseteq Y$, and consequently $D' \cap N[X] = \emptyset$, contrary to the fact that $N[X]$ is a w-balanced separator in G. Thus setting $Y(K) = Y$ the theorem holds. ■

We are now ready to prove Theorem 6.1.

Proof. Let G be an (L, d, r)-breakable graph. Let $C(d) = 100d^2$. We define several sequences of subgraphs of G and subsets of $V(G)$. Let $G_0 = G$; let $K_0 = Y(K_0) = L_0 = Z_0 = \emptyset$ and let $Bad_0 = V(G)$. Now suppose that we have defined $G_{i-1}, K_{i-1}, L_{i-1}, Z_{i-1}, Bad_{i-1}$ with the following properties:

I. $\alpha(Bad_{i-1}) \leq \frac{n}{2^{i-1}}$.
II. $\alpha(Z_{i-1}) \leq C(d)L \cdot \lceil\frac{d(i-1)}{r}\rceil(i - 1)$.
III. $K_{i-1} \subseteq Z_{i-1}$.
IV. If $v \in G_{i-1}$ has a neighbor in K_j for every $j \in \{1, \ldots, i - 1\}$, then $v \in Bad_{i-1}$.
V. $G_{i-1} = G \setminus Z_{i-1}$.
VI. If $i > 1$ and D is a connected subset of $G \setminus Z_{i-1}$ with $w(D) > \frac{1}{2}$, then for every $1 \leq j \leq i - 1$, $K_j \setminus L_{i-1}$ is anticomplete to D, and $K_j \cap L_{i-1} \cap N(D) \neq \emptyset$.
VII. If there is a connected subset D of $G \setminus Z_{i-1}$ with $w(D) > \frac{1}{2}$, then L_{i-1} is a clique and $|L_{i-1}| \leq d(i - 1)$.

We proceed as follows. If Z_{i-1} is a balanced separator of G, we stop the construction. Otherwise, we construct $G_i, K_i, L_i, Z_i, Bad_i$ and show that the properties above continue to hold. If $\alpha(Bad_{i-1}) > 96d^2$, let Z be obtained by applying Lemma 5.1 with $Y = Bad_{i-1}$ and $c = 2d$; then $\alpha(Z) \leq 4d$. If $\alpha(Bad_{i-1}) \leq 96d^2$, let $Z = Bad_{i-1}$.

In both cases $\alpha(Z) \leq 96d^2$. Let $G' = G_{i-1} \setminus Z$. Let $w'(v) = \frac{w(v)}{\Sigma_{v \in V(G')} w(v)}$. Then w' is a normal weight function on G', and for every $v \in G'$, $w'(v) \geq w(v)$. Let $K_i, Y(K_i)$ be as in Theorem 6.2 applied to G' and w'.

Let $v \in K_i$. By (VII) (for $i-1$), $L_{i-1} \setminus N(v)$ can be partitioned into at most $\lceil \frac{d(i-1)}{r} \rceil$ cliques each of size at most r. Since G is (L, d, r)-breakable, this implies that there exists a set $Z(v)$ with $\alpha(Z(v)) \leq \lceil \frac{d(i-1)}{r} \rceil L$, such that $Z(v)$ separates $\{v\}$ from $L_{i-1} \setminus N(v)$ in G. Let $Z' = \bigcup_{v \in K_i} Z(v)$; then $\alpha(Z') \leq d \lceil \frac{d(i-1)}{r} \rceil L$.

Let $Z_i = Z_{i-1} \cup Z' \cup Z \cup K_i \cup Y(K_i)$. Now (III) holds. Next we have:

$$\alpha(Z_i \setminus Z_{i-1}) \leq d \lceil \frac{d(i-1)}{r} \rceil L + 96d^2 + 1 + d^2 L \leq C(d) \lceil \frac{d(i-1)}{r} \rceil L.$$

It follows that $\alpha(Z_i) \leq C(d) \lceil \frac{di}{r} \rceil L \times i$, and (II) holds.

Let $G_i = G_{i-1} \setminus Z_i$; now (V) holds. Let $Bad_i = Bad_{i-1} \cap N_{G'}(K_i) \cap V(G_i)$; now (IV) holds. Since either $Bad_{i-1} \subseteq Z$, or $Z \cap K_i = \emptyset$, we have that for every $v \in K_i$, $\alpha(N_{G'}(v) \cap Bad_{i-1}) \leq \frac{n}{2^i d}$. It follows that $\alpha(Bad_i) \leq \frac{n}{2^i}$, and (I) holds. If Z_i is a w'-balanced separator in G_{i-1}, let $L_i = L_{i-1}$; now (VI) and (VII) hold vacuously. Thus we may assume not, and let D be the maximal connected subset of G_i with $w(D) > \frac{1}{2}$. We now define L_i and check that (VI) and (VII) hold. Since $D \subseteq G_{i-1}$, we know (by (VI) for $i-1$) that $K_j \setminus L_{i-1}$ is anticomplete to D for every $j \leq i-1$. Since $N[K_i] \cup Y(K_i)$ is a w-balanced separator in G', it follows that $N(D) \cap K_i \neq \emptyset$. Since $Z' \cap G_i = \emptyset$, it follows that $N(D) \cap K_i$ is complete to $N(D) \cap L_{i-1}$; let

$$L_i = (N(D) \cap L_{i-1}) \cup (N(D) \cap K_i).$$

Then L_i is a clique, and no vertex of $K_i \setminus L_i$ has a neighbor in D. Since $L_{i-1} \cap N(D) \subseteq L_i$, and since $K_j \setminus L_{i-1}$ is anticomplete to D for every $j \leq i-1$, it follows that for every $j \leq i$, $K_j \setminus L_i$ is anticomplete to D. Let $1 \leq j \leq i$; then $D \subseteq G_{j-1}$. It follows from (VI) (for $i-1$) that $K_j \cap L_{i-1} \cap N(D) \neq \emptyset$. Consequently $L_i \cap K_j \neq \emptyset$; thus (VI) holds. Since $L_i \setminus L_{i-1} \subseteq K_i$, (VII) holds. We have shown that properties (I)–(VII) are maintained at each step of the construction.

We can now complete the proof of the theorem. It follows immediately from (I) that there exists $k \leq 1 + \log n$ such that $Bad_k = \emptyset$. We claim that Z_k is a balanced separator in G (and in particular the construction stops). Suppose not. Then the construction continues and the sets $G_{k+1}, K_{k+1}, L_{k+1}, Z_{k+1}, Bad_{k+1}$ are defined. Also there exists a component D of $G \setminus Z_k = G_k$ such that $w(D) > \frac{1}{2}$. We apply (VI) with $i - 1 = k + 1$. It follows that there is a verex $v \in L_{k+1} \cap K_{k+1}$. Since $L_{k+1} \cap K_j \neq \emptyset$ for every $1 \leq j \leq k$, we deduce from (III), (IV), (V) and (VII) that $v \in Bad_k$, a contradiction.

Now by (II) we have

$$\alpha(Z_{k+1}) \leq C(d) \lceil \frac{d(k+1)}{r} \rceil L \times (k+1) \leq C(d) \frac{d(\log n + 2)}{r} (2 + \log n) L,$$

as required.

■

7. The proof of Theorem 1.2

We start with a lemma. (There are many versions of this lemma; we chose one with a simple proof, without optimizing the constants.)

Lemma 7.1. *Let G be a graph, let $c \in [\frac{1}{2}, 1)$, and let d be a positive integer. If for every normal weight function w on G, there is a (c, w)-balanced separator X_w with $\alpha(X_w) \leq d$, then the tree-independence number of G is at most $\frac{3-c}{1-c}d$.*

Proof. We will prove that for every set $Z \subseteq V(G)$ with $\alpha(Z) \leq \frac{2}{1-c}d$ there is a tree-decomposition (T, χ) of G such that $\alpha(\chi(t)) \leq \frac{3-c}{1-c}d$ for every $v \in T$, and there exists $t \in T$ such that $Z \subseteq \chi(t)$.

The proof is by induction on $|V(G)|$. Observe that every induced subgraph of G satisfies the assumption of the theorem.

Let $Z \subseteq V(G)$ with $\alpha(Z) \leq \frac{2}{1-c}d$. Let I be a stable set of Z with $|I| = \alpha(Z)$. Define a function w where $w(v) = \frac{1}{|I|}$ if $v \in I$, and $w(v) = 0$ if $v \notin I$. Then w is a normal weight function on G. Let X be a (c, w)-balanced separator with $\alpha(X) \leq d$. Then $V(G) \setminus X = V_1 \cup V_2$, where V_1 is anticomplete to V_2, and $|I \cap V_i| \leq c|I|$ for $i = 1, 2$. Define $G_i = V_i \cup X$ and $Z_i = (Z \cap V_i) \cup X$. Since $|I \cap V_1| \leq c|I|$, it follows that $|I \cap G_2| \geq (1-c)|I|$. Since $\alpha(X) \leq d$, we have that $|I \cap V_2| \geq (1-c)|I| - d$. It follows that $\alpha(Z \cap V_1) \leq c|I| + d$. Consequently,

$$\alpha(Z_1) \leq \alpha(Z \cap V_1) + \alpha(X) \leq c|I| + 2d \leq \frac{2}{1-c}d.$$

Similarly, $\alpha(Z_2) \leq \frac{2}{1-c}d$.

Inductively, for $i = 1, 2$ there is a tree-decomposition (T_i, χ_i) of G_i such that $\alpha(\chi_i(t)) \leq \frac{3-c}{1-c}d$ for every $v \in T_i$, and there exists $t_i \in T_i$ such that $Z_i \subseteq \chi(t)$. Now let T be obtained from the disjoint union of T_1 and T_2 by adding a new vertex t_0 adjacent to t_1 and t_2 (and with no other neighbors). Define $\chi(t) = \chi_i(t)$ for every $t \in T_i$, and let $\chi(t_0) = Z \cup X$. Since

$$\alpha(Z \cup X) \leq \alpha(Z) + \alpha(X) \leq \frac{2}{1-c}d + d = \frac{3-c}{1-c}d,$$

(T, χ) satisfies the conclusion of the theorem. ■

Next we restate and prove Theorem 1.2:

Theorem 7.2. *There exists a constant c such that for every integer $n > 1$ every n-vertex graph $G \in \mathcal{C}$ has tree independence number at most at most $c(\log n)^2$.*

Proof. The proof is by induction on n. Let $G \in \mathcal{C}$. Let d be as in Theorem 3.1. Let $C(d)$ be as in Theorem 6.1, and let $c = 165C(d)$. We may assume that $n \geq c$. Let $r = d(2 + \log n)$, and let $L = 32 \log n$. By Theorems 3.1 and 4.5, and since every clique is cooperative, it follows that G is (L, d, r)-breakable. Now by Theorem 6.1, for every normal weight function w on G, there exists a w-balanced separator Y in G such that $\alpha(Y) \leq C(d)\frac{d(2+\log n)}{r}(2 + \log n)L \leq 33C(d)(\log n)^2$. Now Theorem 7.2 follows from Lemma 7.1. ■

8. Algorithmic consequences

Theorem 7.2 implies quasi-polynomial time (namely $2^{(\log n)^{O(1)}}$ time) algorithms for a number of problems. In particular Dallard et al. [6] gave $n^{O(k)}$ time algorithms for Maximum Weight Independent Set and Maximum Weight Induced Matching assuming that a tree decomposition with independence number at most k is given as input. Subsequently Dallard et al. [5]

gave an algorithm that takes as input a graph G and integer k, runs in time $2^{O(k^2)}n^{O(k)}$ and either outputs a tree decomposition of G with independence number at most $8k$, or determines that the tree independence number of G is larger than k. Theorem 7.2, together with these results (setting $k = c\log^2 n$), immediately imply the following theorem.

Theorem 8.1. *Maximum Weight Independent Set and Maximum Weight Induced Matching admit algorithms with running time $n^{O((\log n)^3)}$ on graphs in $\mathcal{C}$.*

It is worth mentioning that the $n^{O(k)}$ time algorithm of Dallard et al [6] works for a slightly more general packing problem (see their Theorem 7.2 for a precise statement) that simultaneously generalizes Maximum Weight Independent Set and Maximum Weight Induced Matching. Thus we could have stated Theorem 8.1 for this even more general problem.

Lima et al. [11] observed that the algorithm of Dallard et al [6] can be generalized to a much more general class of problems. In particular they show that for every integer ℓ and CMSO_2 formula ϕ, there exists an algorithm that takes as input a graph G of tree independence at most k, and a weight function $w : V(G) \to \mathbb{N}$, runs in time $f(k, \phi, \ell)n^{O(\ell k)}$ and outputs a maximum weight vertex subset S such that $G[S]$ has treewidth at most ℓ and $G[S] \models \phi$. This formalism captures Maximum Weight Independent Set, Maximum Weight Induced Matching as well as Maximum Weight Induced Forest, recognition of many well-studied graph classes (including $\mathcal{C}$) and a host of other problems. We remark that their result (Theorem 6.2 of [11]) is stated in terms of clique number rather than treewidth, however at the very beginning in the proof they show that in this context bounded clique number implies treewidth at most ℓ and then proceed to prove the theorem as stated here.

Unfortunately the algorithm of [11] does not give any meaningful results when combined with Theorem 7.2. The reason is that the function $f(k, \phi, \ell)$ bounding the running time of the algorithm is a tower of exponentials, which leads to super-exponential running time bounds even when $k = c\log^2 n$. However it turns out that the algorithm of [11] can be modified to run in time $(f(\ell, \phi)n)^{O(\ell k)}$ [10], which is quasi-polynomial for every fixed ℓ, ϕ when $k = O(\log^2 n)$. This improvement immediately leads to the following theorem.

Theorem 8.2. *For every integer ℓ and CMSO_2 formula ϕ, there exists an algorithm that takes as input a graph $G \in \mathcal{C}$ and a weight function $w : V(G) \to \mathbb{N}$, runs in time $(f(\phi, \ell)n)^{O(\ell \log^2 n)}$ and outputs a maximum weight vertex subset S such that $G[S]$ has treewidth at most ℓ and $G[S] \models \phi$.*

We refer to [11] for a discussion of the set of problems that are captured by Theorem 8.2.

9. Acknowledgments

Peter Gartland proved Theorem 3.1 independently using a different method. We thank him for sharing his proof with us, and for many useful discussions.

References

[1] T. Abrishami, M. Chudnovsky, C. Dibek, and K. Vušković. Submodular functions and perfect graphs. *to appear in Mathematics of Operations Research.*

[2] T. Abrishami, M. Chudnovsky, S. Hajebi, and S. Spirkl. Induced subgraphs and tree decompositions III. Three-path-configurations and logarithmic treewidth. *Adv. Comb.*, pages Paper No. 6, 29, 2022.

[3] 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.

[4] 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.

[5] C. Dallard, F. Fomin, T. Korhonen, P. Golovach, and M. Milanič. Computing Tree Decompositions with Small Independence Number. Preprint available at https://arxiv.org/abs/2207.09993.
[6] C. Dallard, M. Milanič, and K. Štorgel. Treewidth versus clique number. II. Tree independence number. Preprint available at https://arxiv.org/abs/2111.04543.
[7] P. Erdös and G. Szekeres. A combinatorial problem in geometry. *Compositio Math.*, 2:463–470, 1935.
[8] M. R. Garey and D. S. Johnson. "Strong" NP-completeness results: motivation, examples, and implications. *J. Assoc. Comput. Mach.*, 25(3):499–508, 1978.
[9] 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.
[10] P. T. Limma, M. Milanič, P. Muršič, K. Okrasa, P. Rzążewski, and K. Štorgel. Private communication.
[11] P. T. Limma, M. Milanič, P. Muršič, K. Okrasa, P. Rzążewski, and K. Štorgel. Tree decompositions meet induced matchings: beyond max weight independent set. Preprint available at https://arxiv.org/abs/2402.15834.
[12] N. Robertson and P. D. Seymour. Graph minors. XVI. Excluding a non-planar graph. *J. Combin. Theory Ser. B*, 89(1):43–76, 2003.
[13] A. Scott and P. Seymour. A survey of χ-boundedness. *J. Graph Theory*, 95(3):473–504, 2020.
[14] 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.