

Independence and Efficient Domination on P_6 -free Graphs*

DANIEL LOKSHTANOV, University of Bergen

MARCIN PILIPCZUK, University of Warsaw

ERIK JAN VAN LEEUWEN, Utrecht University

In the MAXIMUM WEIGHT INDEPENDENT SET problem, the input is a graph G , every vertex has a non-negative integer weight, and the task is to find a set S of pairwise non-adjacent vertices, maximizing the total weight of the vertices in S . We give an $n^{O(\log^2 n)}$ time algorithm for this problem on graphs excluding the path P_6 on 6 vertices as an induced subgraph. Currently, there is no constant k known for which MAXIMUM WEIGHT INDEPENDENT SET on P_k -free graphs becomes NP-hard, and our result implies that if such a k exists, then $k > 6$ unless all problems in NP can be decided in quasi-polynomial time.

Using the combinatorial tools that we develop for the above algorithm, we also give a polynomial-time algorithm for MAXIMUM WEIGHT EFFICIENT DOMINATING SET on P_6 -free graphs. In this problem, the input is a graph G , every vertex has an integer weight, and the objective is to find a set S of maximum weight such that every vertex in G has exactly one vertex in S in its closed neighborhood, or to determine that no such set exists. Prior to our work, the class of P_6 -free graphs was the only class of graphs defined by a single forbidden induced subgraph on which the computational complexity of MAXIMUM WEIGHT EFFICIENT DOMINATING SET was unknown.

CCS Concepts: •Mathematics of computing → Graph theory; Graph algorithms; Combinatorics; •Theory of computation → Graph algorithms analysis; Branch-and-bound; Complexity classes; Problems, reductions and completeness;

Additional Key Words and Phrases: independence, efficient domination, P_6 -free graphs, (quasi)polynomial-time algorithms

ACM Reference format:

Daniel Lokshtanov, Marcin Pilipczuk, and Erik Jan van Leeuwen. 2016. Independence and Efficient Domination on P_6 -free Graphs. *ACM Trans. Algor.* 1, 1, Article 1 (January 2016), 25 pages.

DOI: 10.1145/nnnnnnn.nnnnnnn

1 INTRODUCTION

An *independent set* in a graph is a set of pairwise non-adjacent vertices. In the INDEPENDENT SET problem the input is a graph G on n vertices and an integer t , and the task is to determine whether G contains an independent set of size at least t . INDEPENDENT SET is a fundamental and extremely well-studied graph problem. It was one of the very first problems to be shown NP-complete [27, 38], and a significant amount of research [2, 22, 24, 33, 41, 42, 46, 55]¹ has gone into identifying classes of graphs on which the problem becomes polynomial-time solvable.

A complete classification of the complexity status of INDEPENDENT SET on all classes of graphs seems out of reach. However, obtaining such a classification for all classes of graphs defined by excluding a single connected graph H as an induced subgraph (we call such graphs H -free) looks like an attainable, yet very challenging, goal. In particular, Alekseev [1] showed in 1982 that INDEPENDENT SET remains NP-complete on H -free graphs whenever H is connected, but neither a path nor a subdivision of the claw.

The complexity of INDEPENDENT SET on classes of P_k -free graphs (we denote by P_k the path on k vertices) has been subject to intense scrutiny, but yielding rather modest progress. For P_4 -free graphs a polynomial-time algorithm was given by Corneil et

*An extended abstract of this paper appeared as: D. Lokshtanov, M. Pilipczuk, and E.J. van Leeuwen. Independence and Efficient Domination on P_6 -free Graphs. In R. Krauthgamer, editor, *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016, Arlington, VA, USA, January 10-12, 2016*, SIAM, pages 1784–1803, 2016.

¹This list is far from exhaustive, see the Information System on Graph Classes and their Inclusions (ISGCI) [25].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2016 ACM. Manuscript submitted to ACM

al. [24] in 1981, and it took more than 30 years until a polynomial-time algorithm for the problem on P_5 -free graphs was discovered by Lokshtanov et al. [41] in 2014. Meanwhile, a substantial amount of work was devoted to INDEPENDENT SET on subclasses of P_5 -free graphs [6, 7, 16, 31, 47, 59], and some progress has been reported on subclasses of P_6 -free graphs [4, 11, 12, 14, 30, 39, 40, 43, 45, 48, 50–53] and even on subclasses of P_7 -free graphs [17, 30, 43, 44, 49]. Recently, Bacsó et al. [3] and Brause [21] independently proved subexponential-time exact algorithms for INDEPENDENT SET on P_k -free graphs for any fixed k , generalizing an earlier such result for P_5 -free graphs [54].

In this paper we push the boundary of knowledge on the complexity of INDEPENDENT SET on P_k -free graphs a step forward by giving a $n^{O(\log^2 n)}$ -time algorithm for INDEPENDENT SET on P_6 -free graphs. Our algorithm also works for the weighted version of the problem. In the MAXIMUM WEIGHT INDEPENDENT SET problem the input is a graph G where every vertex has a non-negative integer weight, and the task is to compute an independent set in G that maximizes the sum of the weights of the vertices in it.

THEOREM 1.1. *There is an $n^{O(\log^2 n)}$ -time, polynomial-space algorithm for MAXIMUM WEIGHT INDEPENDENT SET on n -vertex P_6 -free graphs.*

The algorithm of Theorem 1.1 does not completely resolve the complexity status of INDEPENDENT SET on P_6 -free graphs, as it runs in quasipolynomial-time rather than polynomial time. However, Theorem 1.1 does imply that MAXIMUM WEIGHT INDEPENDENT SET on P_6 -free graphs is not NP-hard, unless all problems in NP can be solved in quasipolynomial-time.

This result hints at the existence of a polynomial-time algorithm for the problem also on P_6 -free graphs. Very recently, Grzesik et al. [34] actually proved that MAXIMUM WEIGHT INDEPENDENT SET can be solved in polynomial time on P_6 -free graphs.

On the way to developing our algorithm for INDEPENDENT SET, we prove several new combinatorial properties of P_6 - and P_7 -free graphs. We leverage these new combinatorial insights to develop a polynomial-time algorithm for the EFFICIENT DOMINATING SET problem on P_6 -free graphs. We say that a vertex *dominates* itself and all of its neighbors. An *efficient dominating set* in a graph G is a vertex set S such that every vertex v in the graph is dominated by exactly one vertex in S . Not all graphs have an efficient dominating set, and in the EFFICIENT DOMINATING SET problem the input is a graph G and the task is to determine whether G has an efficient dominating set. We remark that the problem also goes by the name PERFECT CODE [5]. Observe that we do not ask for the smallest or largest efficient dominating set, only whether there exists one. This is because whenever a graph G has an efficient dominating set, all such sets have the same cardinality [36]. In the weighted variant, called MAXIMUM WEIGHT EFFICIENT DOMINATING SET, every vertex has an integer weight and the task is to find a maximum weight efficient dominating set, if one exists. Since the weights may be negative, there is no real difference between maximizing and minimizing the weight of the solution. Our second main theorem is the following.

THEOREM 1.2. *There is a polynomial-time algorithm for MAXIMUM WEIGHT EFFICIENT DOMINATING SET on P_6 -free graphs.*

Prior to our work, the P_6 was the only graph H , connected or not, for which the complexity of EFFICIENT DOMINATING SET on H -free graphs was unknown [10]. Thus our work completes the complexity classification of EFFICIENT DOMINATING SET (and MAXIMUM WEIGHT EFFICIENT DOMINATING SET) on classes of graphs defined by a single forbidden induced subgraph and resolves the main open problem of [9, 10, 13, 15]. We remark that an alternative polynomial-time algorithm for MAXIMUM WEIGHT EFFICIENT DOMINATING SET on P_6 -free graphs has been independently obtained by Brandstädt and Mosca [18–20] using different methods.²

Methodology. The polynomial-time algorithm for MAXIMUM WEIGHT INDEPENDENT SET on P_5 -free graphs of Lokshtanov et al. [41] demonstrated that investigating *potential maximal cliques* and *minimal separators* (see Section 2 for definitions) yields valuable insights on the structure of P_5 -free graphs. Our algorithm for P_6 -free graphs is also based on studying potential maximal cliques and minimal separators. However, this is where the similarity between the two algorithms ends, as essentially all of the arguments used in the algorithm for P_5 -free graphs quickly break down for P_6 -free graphs.

At heart our algorithm is very simple: the algorithm picks a node v and proceeds recursively in two branches. In the first v is included in the independent set, and the algorithm needs to solve $G - N(v)$ recursively. In the second v is excluded from the

²Although [18] appeared on arXiv a few weeks after this paper, the authors of [18–20] contacted us and shared with us a preliminary version of [18–20] immediately after we posted our work.

independent set, and the algorithm is called recursively on $G - v$. If in any recursive call the graph becomes disconnected the algorithm solves the connected components independently. The crux of the analysis is to show that one can always cleverly chose the vertex v , such that after only a few branches either the size of the graph decreases by at least $0.1n$, or the graph breaks into connected components of size at most $0.9n$.

Roughly speaking, the vertex v to branch on is chosen as follows. The algorithm identifies a *nuke* in G : a relatively small vertex set S such that every connected component of $G - S$ has size at most $0.9n$ (for a formal definition of a nuke, see Definition 4.1). The algorithm then picks a vertex v with a large neighborhood in S to branch on. In order to guarantee the existence of a nuke S and a vertex v with a large neighborhood in S we prove the following theorem about minimal separators in P_7 -free graphs.

THEOREM 1.3. *There exists a positive constant $0 < \alpha \leq 1$ such that for every P_7 -free graph G , for every minimal separator S in G , and for every probability measure μ on S , there exists a vertex $v \in V(G)$ satisfying $\mu(N(v) \cap S) \geq \alpha$.*

The reason that Theorem 1.3 is not already sufficient to yield a quasipolynomial-time algorithm for INDEPENDENT SET on P_7 -free graphs is that, despite the similarity between the definitions of nukes and minimal separators, not all nukes are minimal separators. For P_6 -free graphs we are able to prove an analogue of Theorem 1.3 for nukes rather than minimal separators, and this is sufficient to give a $n^{O(\log^2 n)}$ -time algorithm for MAXIMUM WEIGHT INDEPENDENT SET. As a first step to lift Theorem 1.3 to work for nukes we generalize it to potential maximal cliques in P_7 -free graphs.

THEOREM 1.4. *There exists a positive constant $0 < \beta \leq 1$ such that for every connected P_7 -free graph G on at least two vertices, for every potential maximal clique Ω in G , and every probability measure μ on Ω , there exists a vertex $v \in V(G)$ satisfying $\mu(N(v) \cap \Omega) \geq \beta$.*

Theorem 1.4 turns out to be very useful not only in our quasipolynomial-time algorithm for MAXIMUM WEIGHT INDEPENDENT SET, but for the polynomial-time algorithm for MAXIMUM WEIGHT EFFICIENT DOMINATING SET as well. Indeed, an almost immediate consequence of Theorem 1.4 is that for any P_7 -free graph G , any efficient dominating set X in G and any potential maximal clique Ω in G , $|X \cap \Omega| \leq 1/\beta$ (see Lemma 6.2 for a simple proof).

The observation above strongly suggests that one can solve MAXIMUM WEIGHT EFFICIENT DOMINATING SET on P_7 -free graphs in polynomial time by doing dynamic programming over the tree decomposition of an arbitrarily chosen minimal triangulation of G . For P_7 -free graphs this approach fails, as is expected from the NP-completeness of EFFICIENT DOMINATING SET on P_7 -free graphs [15, 58]. On the other hand, for P_6 -free graphs, we are able to carry this approach through.

We mention here that this approach follows a completely disjoint direction from the one followed in recent papers [9, 13] that gave polynomial-time algorithms for subclasses of P_6 -free graphs. In particular, those papers show that one can reduce to MAXIMUM WEIGHT INDEPENDENT SET on the square of the graph by proving special properties of the square when the graph is from such a subclass and has an efficient dominating set.

Furthermore, the drawback of our approach is that it inherently leads to a huge degree of the polynomial in the running time bound in Theorem 1.2. A careful inspection of our proofs shows that we prove Theorem 1.4 with $\beta = 1/576$. As the dynamic programming algorithm enumerates all subsets of a potential maximal clique up to size $1/\beta$, this enumeration adds a factor of n^{576} to the running time bound. Since the running time bound of Brandstädt and Mosca [18–20] is much better, we refrain from a precise estimation of the running time bound of the algorithm of Theorem 1.2.

Outline of the paper. In § 2 we set up the definitions and necessary notations. In § 3 we prove Theorems 1.3 and 1.4, while § 4 contains the generalization of Theorem 1.3 to nukes. We then proceed to the main algorithmic results, § 5 contains the quasipolynomial-time algorithm for MAXIMUM WEIGHT INDEPENDENT SET, while § 6 contains the polynomial-time algorithm for MAXIMUM WEIGHT EFFICIENT DOMINATING SET, both on P_6 -free graphs. In § 7 we conclude with some open problems and counterexamples to the most natural generalizations of the structural results underlying our algorithms.

2 PRELIMINARIES

For all graph terminology not defined here, we refer to the monograph by Diestel [26]. For a graph G and sets $A, B \subseteq V(G)$, we denote $N_B(A) := N(A) \cap B$.

Let G be a graph; throughout, we assume that all graphs are finite, simple, and undirected. Given distinct, non-adjacent $s, t \in V(G)$, a set $S \subseteq V(G)$ is an s - t separator if s and t are in distinct components of $G \setminus S$. We say that $S \subseteq V(G)$ is a *minimal* s - t

separator if no $S' \subsetneq S$ is also an s - t separator. Then $S \subseteq V(G)$ is a (*minimal*) *separator* of G if S is a (minimal) s - t separator for some $s, t \in V(G)$. Given a separator $S \subseteq V(G)$, a component C of $G \setminus S$ is said to be *full* if every vertex of S has a neighbor in C . It can be shown that S is a minimal separator if and only if $G \setminus S$ has at least two full components (cf. Golumbic [32, Ch. 4 Ex. 10]).

A set $M \subseteq V(G)$ is a *module* of G if every vertex $v \in V(G) \setminus M$ is either fully adjacent or fully anti-adjacent to M ; that is, either $vu \in E(G)$ for each $u \in M$ or $vu \notin E(G)$ for each $u \in M$. A module M of G is *trivial* if $M = V(G)$, $M = \emptyset$, or $|M| = 1$. A graph is *prime* if it only has trivial modules. A *modular partition* \mathcal{M} of G is a set of disjoint modules of G with union $V(G)$. The *quotient graph* G/\mathcal{M} induced by \mathcal{M} has a vertex for each module of \mathcal{M} and has an edge between two vertices if and only if the corresponding modules are fully adjacent to each other. Observe that, by definition, a non-edge between two vertices in the quotient graph implies that the corresponding modules are fully anti-adjacent to each other. A module M of G is *proper* if $M \neq V(G)$. A module M of G is *strong* if for every other module M' of G , either $M \subseteq M'$, $M' \subseteq M$, or $M \cap M' = \emptyset$.

THEOREM 2.1 ([35, THEOREM 2]). *Let G be a connected graph on at least two vertices and let \mathcal{M} denote the set of maximal proper strong modules of G . Then \mathcal{M} is a modular partition of G , and the quotient graph G/\mathcal{M} is either a clique or a prime graph.*

A graph G is *chordal* if every induced cycle of G has length three. A *clique tree* for a graph G is a pair (T, Φ) where T is a tree and Φ is a bijection between $V(T)$ and the set of maximal cliques of G , and for every $v \in V(G)$, the vertices $t \in V(T)$ for which $v \in \Phi(t)$ induce a subtree of T . Equivalently, for every pair $t, t' \in V(T)$ and every $t'' \in V(T)$ on the path between t and t' , $\Phi(t) \cap \Phi(t') \subseteq \Phi(t'')$. It is known that a graph is chordal if and only if it admits a clique tree [23, 29, 56, 57]. Moreover, since the number of maximal cliques of a chordal graph G is at most $|V(G)|$ [28], a clique tree of a chordal graph can be constructed in polynomial time.

A *tree decomposition* of a graph G is a pair (T, Π) where T is a tree and $\Pi : V(T) \rightarrow 2^{V(G)}$ such that the following three properties hold: $V(G) = \bigcup_{t \in V(T)} \Pi(t)$; for each $uv \in E(G)$ there is a $t \in V(T)$ such that $u, v \in \Pi(t)$; for each $v \in V(G)$, the vertices $t \in V(T)$ for which $v \in \Pi(t)$ induce a subtree of T .

A *triangulation* of a graph G is a set $F \subseteq ((V(G) \times V(G)) \setminus \{vv \mid v \in V(G)\}) \setminus E(G)$ such that the graph $G + F := (V(G), E(G) \cup F)$ is chordal. We say that F is a *minimal triangulation* of G if no $F' \subsetneq F$ is a triangulation of G .

A *potential maximal clique* of G is a set $\Omega \subseteq V(G)$ such that Ω induces a maximal clique in some minimal triangulation of G . We need the following properties of potential maximal cliques due to Bouchitté and Todinca [8].

THEOREM 2.2 ([8, LEMMA 3.14]). *Let G be a graph. If $\Omega \subseteq V(G)$ is a potential maximal clique of G , then for every connected component C of $G \setminus \Omega$, the set $N_G(C) \subseteq \Omega$ is a minimal separator of G .*

THEOREM 2.3 ([8, THEOREM 3.15]). *Let G be a graph. A set $\Omega \subseteq V(G)$ is a potential maximal clique of G if and only if the following two conditions hold:*

- (1) *for every connected component C of $G \setminus \Omega$, we have $N_G(C) \subseteq \Omega$;*
- (2) *for every two distinct vertices $x, y \in \Omega$, either $xy \in E(G)$ or there exists a component C of $G \setminus \Omega$ such that $x, y \in N_G(C)$ (in this case we say that the non-edge xy is covered by the component C).*

Throughout the paper, we use the following notation for probability measures. Let U be a universe and $X \subseteq U$. If we define μ as a probability measure on X , then we implicitly assume that $\mu(v) = 0$ for all $v \in U \setminus X$. This implies that for any $Y \subseteq U$, it holds that $\mu(Y) = \mu(Y \cap X)$. Moreover, if μ is a probability measure on X and $Z \subseteq X$ with $\mu(Z) > 0$, then the *restriction* of μ to Z is the probability measure μ' on Z such that $\mu'(v) = \mu(v)/\mu(Z)$ for all $v \in Z$.

Finally, throughout the paper, we often identify graphs and connected components with their vertex sets.

3 HITTING SEPARATORS AND POTENTIAL MAXIMAL CLIQUES

3.1 Proof of Theorem 1.3

For the sake of convenience, we first restate Theorem 1.3.

THEOREM 1.3. *There exists a positive constant $0 < \alpha \leq 1$ such that for every P_7 -free graph G , for every minimal separator S in G , and for every probability measure μ on S , there exists a vertex $v \in V(G)$ satisfying $\mu(N(v) \cap S) \geq \alpha$.*




Fig. 1. Two possibilities for (x, y) being lucky.

We now prove the theorem. Let G be a graph, let S be a minimal separator of G , let μ be any probability measure on S , and let $0 < \alpha \leq 1$ be some constant chosen later. For the sake of contradiction, assume that for every $v \in V(G)$ we have $\mu(N(v)) < \alpha$. This implies that $\mu(x) < \alpha$ for every $x \in S$, because by the minimality of S , x has a neighbour v in some (full) component of $G \setminus S$, and thus $\mu(x) \leq \mu(N(v)) < \alpha$.

Let A and B be two full components of $G \setminus S$. We say that $x \in S$ is *lucky* (with respect to A) if there exists an induced P_4 in G with one endpoint in x and the remaining three vertices in A . We say that an ordered pair $(x, y) \in S \times S$ is *lucky* (with respect to A) if x is lucky or there exists an induced P_4 in G with endpoints x and y and its middle two vertices in A . We emphasize here that this definition is not symmetric with respect to x and y . The following lemma is the crucial step in the argumentation.

LEMMA 3.2. *Let G be a graph, and let S, μ, α, A , and B as above; in particular, $\mu(N(v)) < \alpha$ for every $v \in V(G)$. If we choose two vertices $x, y \in S$ independently at random according to distribution μ , then the probability that (x, y) is not lucky with respect to A (or B) is less than 6α .*

PROOF. If $|A| = 1$, then the single vertex a of A is adjacent to all vertices of S , as A is a full component of $G \setminus S$. Hence, $\mu(N(a)) = 1 > \alpha$, a contradiction. Therefore, $|A| > 1$.

Consider the graph $G[A]$, and let \mathcal{M} be the family of maximal proper strong modules of $G[A]$. Note that \mathcal{M} is a modular partition and that the quotient graph of this partition is a clique or a prime graph (Theorem 2.1), since $G[A]$ is connected and $|A| > 1$. Now pick two arbitrary vertices $p, q \in A$ in two distinct modules of \mathcal{M} that are adjacent in the quotient graph. We can indeed pick such p, q , because the quotient graph is connected, as $G[A]$ is connected. Moreover, $|\mathcal{M}| > 1$, because for all modules $M \in \mathcal{M}$ it holds that $M \neq V(G)$ by the fact that M is proper; note that $\mathcal{M} \neq \emptyset$, because $|A| > 1$ and thus the set of singleton modules (one for each vertex) is a family of proper strong modules.

Consider some $(x, y) \in S \times S$ that are chosen independently at random according to distribution μ . In the following, we continuously use that $\mu(N(v)) < \alpha$ for every $v \in V(G)$ and $\mu(u) < \alpha$ for every $u \in S$. With probability less than 2α we have $x \in N(p) \cup N(q)$, and with probability less than 2α we have $x = y$ or $xy \in E(G)$. Furthermore, with probability less than α we have $N(y) \cap A \subseteq N(x) \cap A$, since for a fixed choice of y and $v \in N(y) \cap A$, the probability that $x \in N(v)$ is at most α . Now assume that none of the aforementioned events happen, and pick arbitrary $r \in (N(y) \setminus N(x)) \cap A$.

Let \mathcal{F} be the family of connected components of $G[A \setminus N(x)]$. Consider any $C \in \mathcal{F}$ and any vertex $v \in N(x) \cap A$. If v is neither fully adjacent nor fully anti-adjacent to C , then since C is connected, there exist two neighbouring vertices $u, w \in C$ such that $u \in N(v)$ and $w \notin N(v)$. Since $u, w \notin N(x)$ by the definition of C and \mathcal{F} , the vertices x, v, u, w form a P_4 in G with one endpoint in x ; then x is lucky, and by extension (x, y) is lucky. Hence, we may assume that for every $C \in \mathcal{F}$ and every $v \in N(x) \cap A$, the vertex v is either fully adjacent or fully anti-adjacent to C . In particular, every $C \in \mathcal{F}$ is a module of $G[A]$.




Fig. 2. The left panel shows part of the reasoning of Lemma 3.2 with the choice of C being the connected component of $G[A \setminus N(x)]$ that contains p . The right panel shows an archetypical P_7 constructed in the proof of Theorem 1.3.

Consider the component $C \in \mathcal{F}$ that contains the vertex p ; note that C exists, because $x \notin N(p)$ by assumption (see Fig. 2). Since C is a module of $G[A]$ and \mathcal{M} is the family of maximal proper strong modules of $G[A]$, either there exists a module $M \in \mathcal{M}$ that contains C , or C is a union of several modules of \mathcal{M} and the quotient graph $G[A]/\mathcal{M}$ is a clique.

If $C \subseteq M$ for some $M \in \mathcal{M}$, then consider the module $M' \in \mathcal{M}$ that contains q . By the choice of M and M' , M' is fully adjacent to M , and in particular, M' is fully adjacent to C . Since $C \in \mathcal{F}$ and $C \subseteq M$, we have that M' cannot contain any vertices of any other component in \mathcal{F} . Hence, $M' \subseteq N(x)$. However, $q \notin N(x)$, a contradiction.

Therefore, C is a union of several modules of \mathcal{M} and the quotient graph $G[A]/\mathcal{M}$ is a clique. Then C is fully adjacent to $A \setminus C$, and in particular to every $C' \in \mathcal{F} \setminus \{C\}$, which implies that $\mathcal{F} = \{C\}$. Therefore, there exists a vertex $v \in N(x) \cap A$ with $A \setminus N(x) \subseteq N(v)$, because $G[A]$ is connected and $C = A \setminus N(x)$ is a module. Observe that $y \in N(v)$ with probability less than α , since $\mu(N(v)) < \alpha$. If this does not happen (i.e., $y \notin N(v)$), then x, v, r, y form a P_4 , as $r \in (N(y) \setminus N(x)) \cap A$. Hence, (x, y) is lucky.

By the union bound, the total probability that any of the aforementioned events happen is less than 6α . The lemma follows. \square

We are now ready to conclude the proof of Theorem 1.3. Pick three vertices $x, y_1, y_2 \in S$ independently at random, each with distribution μ . The goal will be to find a P_4 in A from x (possibly to y_1) and a P_4 in B from x (possibly to y_2) that jointly form a P_7 in G . Consider the following set of “bad” events. In the below, we repeatedly rely on Lemma 3.2 and the assumptions that $\mu(N(v)) < \alpha$ for every $v \in V(G)$ and $\mu(u) < \alpha$ for every $u \in S$.

- (x, y_1) is not lucky with respect to A ; this happens with probability less than 6α . Otherwise, let P^1 be the witnessing P_4 .
- (x, y_2) is not lucky with respect to B ; this happens with probability less than 6α . Otherwise, let P^2 be the witnessing P_4 .
- Some vertices from the set $\{x, y_1, y_2\}$ are equal or adjacent; this happens with probability less than 6α .
- One of the (two or three) vertices from $V(P^1) \cap A$ is adjacent to y_2 ; since the choice of x and y_1 is independent of the choice of y_2 , and the path P^1 is a function of the pair (x, y_1) only, this happens with probability less than 3α (y_2 needs to land outside the neighbourhoods of $V(P^1) \cap A$).
- One of the (two or three) vertices from $V(P^2) \cap B$ is adjacent to y_1 ; since the choice of x and y_2 is independent of the choice of y_1 , and the path P^2 is a function of the pair (x, y_2) only, this happens with probability less than 3α (y_1 needs to land outside the neighbourhoods of $V(P^2) \cap B$).

By the union bound, the probability that none of the aforementioned “bad” events happen is greater than $1 - 24\alpha$. Hence, for sure when $\alpha = \frac{1}{24}$, there is a choice of $x, y_1, y_2 \in S$ for which the paths P^1 and P^2 exist and jointly form a P_7 in G . Hence, if G is P_7 -free, then there is a vertex $v \in V(G)$ satisfying $\mu(N(v)) \geq \alpha$ for some constant $\alpha > 0$ (in fact even $\alpha \geq \frac{1}{24}$).




Fig. 3. Illustration of the proof of Lemma 3.3 for $H = P_4$.

3.2 Proof of Theorem 1.4

The main tool in the proof is the following general lemma.

LEMMA 3.3. *Let H be a graph on n_H vertices and m_H edges, let G be a graph, let Ω be a potential maximal clique in G , and let μ be a probability measure on Ω . Then there exists either:*

- (1) a vertex $v \in V(G)$ with $\mu(v) > \frac{1}{2n_H^2}$ or with $\mu(N(v)) > \frac{1}{2n_H^2}$;
- (2) a minimal separator $S \subseteq V(G)$ of G with $\mu(S) > \frac{1}{2n_H m_H}$; or
- (3) an induced subgraph of G isomorphic to a graph obtained from H by replacing every edge by a path of length at least two (i.e., subdividing at least once).

PROOF. Let H, G, Ω , and μ be as in the statement, and assume for the sake of contradiction that neither of the first two outcomes happen. Consider the following random experiment: independently, for every $p \in V(H)$, choose a vertex $x_p \in \Omega$ according to the distribution μ .

For two distinct vertices $p, q \in V(H)$, we have $x_p = x_q$ with probability at most $\frac{1}{2n_H^2}$, and $x_p x_q \in E(G)$ (i.e., $x_q \in N(x_p)$) with probability at most $\frac{1}{2n_H^2}$. Consequently, all vertices $X := \{x_p : p \in V(H)\}$ are pairwise distinct and nonadjacent with probability at least

$$1 - \binom{n_H}{2} \cdot 2 \cdot \frac{1}{2n_H^2} > \frac{1}{2}.$$

Assume that the aforementioned event happens. For every two distinct and nonadjacent vertices $x, y \in \Omega$, fix a component $C(x, y)$ of $G \setminus \Omega$ that covers the non-edge xy (i.e., $x, y \in N(C)$). For every edge $pq \in E(H)$, consider the component $C_{pq} := C(x_p, x_q)$. As the choices of x_r for distinct vertices $r \in V(H)$ are independent, the probability that $x_r \in N(C_{pq})$ for a fixed $r \in V(H) \setminus \{p, q\}$ is at most $\frac{1}{2n_H m_H}$, since $N(C_{pq})$ is a minimal separator of G by Theorem 2.2 and thus $\mu(N(C_{pq})) \leq \frac{1}{2n_H m_H}$ by assumption. Consequently, the probability that X is an independent set of size n_H and for every $pq \in E(H)$ we have $N(C_{pq}) \cap X = \{x_p, x_q\}$ is strictly greater than

$$\frac{1}{2} - n_H m_H \cdot \frac{1}{2n_H m_H} = 0.$$

If this event happens, then for every $pq \in E(H)$ choose a shortest path between x_p and x_q with internal vertices in C_{pq} . The union of all aforementioned paths forms an induced subgraph of G isomorphic to a graph obtained from H by replacing each edge with a path of length at least two, obtaining the last outcome. \square

For the sake of convenience, we restate Theorem 1.4, and then prove it using Lemma 3.3.

THEOREM 1.4. *There exists a positive constant $0 < \beta \leq 1$ such that for every connected P_7 -free graph G on at least two vertices, for every potential maximal clique Ω in G , and every probability measure μ on Ω , there exists a vertex $v \in V(G)$ satisfying $\mu(N(v) \cap \Omega) \geq \beta$.*

PROOF. Let G be a connected P_7 -free graph on at least two vertices, let Ω be a potential maximal clique of G , let μ be any probability measure on Ω , and let $0 < \beta \leq 1$ be some constant chosen later. Let α denote the constant of Theorem 1.3. For sure, if $\beta = \min\{\frac{\alpha}{24}, \frac{1}{32}\}$, then the following happens. Apply Lemma 3.3 with $H = P_4$ and consider its outcomes:

- (1) If $\mu(N(v)) > \frac{1}{32} \geq \beta$, then we are done. Otherwise, if $\mu(v) > \frac{1}{32} \geq \beta$, then by connectivity of G there is a vertex $u \in N(v)$ with $\mu(N(u)) \geq \beta$.
- (2) Note that $\mu(S) > \frac{1}{24}$. Apply Theorem 1.3 to S and the restriction μ' of μ to S ; recall that $\mu(S) > 0.1$, so this restriction is well defined. It follows that there is a $v \in V(G)$ with $\mu'(N(v)) \geq \alpha$ and thus $\mu(N(v)) > \frac{\alpha}{24} \geq \beta$.
- (3) By the choice of H , this implies the existence of an induced P_7 in G , a contradiction.

Therefore, there is a vertex $v \in V(G)$ satisfying $\mu(N(v)) \geq \beta$ for some constant $\beta > 0$ (in fact even $\beta \geq \frac{1}{576}$). \square

4 NUKING A GRAPH

In this section we study the following notion.

Definition 4.1 (nuke, shelter). For a constant $0 < \eta \leq 0.1$ and a threshold $\tau \geq 0$, a set of vertices X is a (η, τ) -nuke in a graph G if the following holds:

- (i) $(1 - 2\eta)|V(G)| \leq \tau \leq (1 - \eta)|V(G)|$
- (ii) $|X| \leq \eta|V(G)|$;
- (iii) for every connected component C of $G - X$ we have $|C| + |X| \leq \tau$.

Given a (η, τ) -nuke X in G , any connected component of $G - X$ is called a *shelter*.³ If the parameters η and τ are clear from the context, we will simply call the set X a *nuke* in G .

Intuitively, a nuke is a small set of vertices in G whose removal breaks G into connected components of multiplicatively smaller size. Our algorithm keeps track of a nuke X in the given input P_6 -free graph G and tries to branch on vertices of G so that X will be removed from G as quickly as possible. This motivation introduces two delicate issues that result in a slightly technical definition of a nuke. First, during branching we need to keep the threshold τ constant, while the size of G drops a bit – if we define the nuke so that, say, $|C| + |X| \leq (1 - \eta)|V(G)|$, a set X may stop to be a nuke due to a removal of a vertex from G and consequent decrease of the bound $(1 - \eta)|V(G)|$. Second, we would like to argue about *inclusion-wise minimal nukes*, which makes the measure $|C| + |X|$ (as opposed to simply $|C|$) more natural, as we can then assume that every element of an inclusion-wise minimal nuke is adjacent to at least two shelters.

The rest of this section is devoted to a proof of the following structural statement.

THEOREM 4.2. *There exists a constant $0 < \gamma \leq 1$ such that for every constant $0 < \eta \leq 0.1$, for every connected P_6 -free graph G on at least two vertices, for every threshold $(1 - 2\eta)|V(G)| \leq \tau \leq (1 - \eta)|V(G)|$, for every inclusion-wise minimal (η, τ) -nuke X in G , and for every probability measure μ on X , there exists a vertex $v \in V(G)$ with $\mu(N(v)) \geq \gamma$.*

Let η, τ, G, X , and μ be as in the statement of Theorem 4.2. We set $\gamma = 0.1\beta \leq 0.1$, where the constant β comes from Theorem 1.4. We will prove Theorem 4.2 by contradiction: assume that for every $v \in V(G)$ we have $\mu(N(v)) < \gamma$. We will unravel subsequent observations about the structure of G , leading to a final contradiction.

We start with the following observation.

CLAIM 4.3. *There exists a minimal triangulation \widehat{G} of G , such that X is a (η, τ) -nuke of \widehat{G} as well and, moreover, the shelters of $\widehat{G} - X$ are exactly the same as of $G - X$.*

PROOF. Consider the following completion G_0 of G : we first turn X into a clique and then, for every shelter C of $G - X$ we turn C into a clique and make it completely adjacent to X . Clearly, G_0 is a chordal graph, and the set of connected components of $G_0 - X$ and $G - X$ are the same. Consequently, any minimal triangulation \widehat{G} of G that is a subgraph of G_0 has the required properties. \square

We fix a minimal triangulation \widehat{G} of G satisfying the statement of Claim 4.3. Observe the following.

³The main motivation for introducing the notion of a shelter is to explicitly distinguish connected components of $G - X$ from connected components of $G - \Omega$ for some potential maximal clique Ω ; we will call the former *shelters*, while the latter will be simply *connected components*.

CLAIM 4.4. *If two vertices $v, u \in V(G) \setminus X$ appear in the same maximal clique of \widehat{G} , then they are contained in the same shelter of $G - X$.*

PROOF. Recall that X is a nuke of \widehat{G} as well, with $\widehat{G} - X$ having the same set of shelters as $G - X$. Furthermore, $uv \in E(\widehat{G})$. \lrcorner

CLAIM 4.5. *For any maximal clique Ω in \widehat{G} we have $\mu(\Omega) \leq 0.1$.*

PROOF. Suppose for the sake of contradiction that $\mu(\Omega) > 0.1$ for some maximal clique Ω in \widehat{G} . Since \widehat{G} is already triangulated, Ω is a potential maximal clique in \widehat{G} too. Moreover, by the conditions on G , \widehat{G} is P_7 -free (even P_6 -free) and has at least two vertices. Now apply Theorem 1.4 to \widehat{G} , Ω , and μ restricted to Ω ; recall that $\mu(\Omega) > 0.1$ by assumption, so this restriction is well defined. Let v be the resulting vertex. Then $\mu(N(v)) \geq 0.1\beta = \gamma$, a contradiction (recall that we assumed that $\mu(N(v)) < \gamma$ at the beginning of the proof of Theorem 4.2). \lrcorner

CLAIM 4.6. *There exists a maximal clique Ω in \widehat{G} such that for every connected component C of $G - \Omega$ we have $\mu(C) \leq 0.5$.*

PROOF. Construct the clique tree (T, Φ) of \widehat{G} . Now, starting from an arbitrary vertex v of T , iteratively advance to the neighbor of v in a subtree T' of $T - v$ for which $\mu(\Phi(T') \setminus \Phi(v)) > 0.5$, until no longer possible. If we can show this process is finite, then $\Phi(v)$ is a maximal clique in \widehat{G} with the desired property.

To show that this process is finite, note that in each iteration, there is at most one subtree into which we can advance by the properties of a clique tree. Now suppose that the process advances from a vertex v to its neighbor u and then back to v . Let T_v denote the subtree of $T - uv$ containing v and T_u the subtree of $T - uv$ containing u . First, note that $\mu(\Phi(T_u) \setminus \Phi(v)) > 0.5$ by assumption. However, $\mu(\Phi(T_v) \setminus \Phi(u)) > 0.5$ by assumption as well. By the properties of a clique tree, the latter implies that $\mu(\Phi(T_v) \setminus \Phi(v)) + \mu(\Phi(v) \setminus \Phi(u)) = \mu(\Phi(T_v) \setminus \Phi(u)) > 0.5$ and thus $\mu(\Phi(T_v) \setminus \Phi(v)) + \mu(\Phi(v)) > 0.5$. However, $\mu(\Phi(T_v) \setminus \Phi(v)) + \mu(\Phi(T_u) \setminus \Phi(v)) + \mu(\Phi(v)) \leq 1$, because $\Phi(T_v) \cap \Phi(T_u) \subseteq \Phi(v)$ by the properties of a clique tree. Combined, this means that $\mu(\Phi(T_u) \setminus \Phi(v)) < 0.5$, a contradiction. Hence, the process is finite by the finiteness of T . \lrcorner

Fix one maximal clique Ω as promised by Claim 4.6. We say that a component C of $G - \Omega$ is *nuked* if $C \cap X \neq \emptyset$.

CLAIM 4.7. *There are at least two nuked components.*

PROOF. By the choice of Ω , every nuked component contains at most half of the measure of X . Furthermore, by Claim 4.5, $\mu(\Omega) \leq 0.1$. Thus, there are at least two nuked components. \lrcorner

By Claim 4.4, all vertices of $\Omega \setminus X$ are contained in one shelter of $G - X$. Let D be this shelter; we set $D = \emptyset$ if $\Omega \subseteq X$.

CLAIM 4.8. $|D| \geq (0.5 - 3\eta)|V(G)| \geq 0.2|V(G)|$.

PROOF. By Claim 4.7, there exists a nuked component C with $|C| \leq |V(G)|/2$. Consider the set $X' = X \setminus C$. By the minimality of X , X' is not a (η, τ) -nuke in G . However, as $|X'| < |X| \leq \eta|V(G)|$, both condition (i) and (ii) of Definition 4.1 hold. Hence, the only reason for X' to not be a nuke is that condition (iii) does not hold, and thus there exists a shelter C' of $G - X'$ that is too large, that is, $|C'| + |X'| > \tau$. By the construction of X' , the shelters of $G - X'$ and $G - X$ are the same, except for $C \cup D$, which is a shelter of $G - X'$, but may contain multiple shelters of $G - X$. Therefore $C' = C \cup D$. Hence, using condition (i) of Definition 4.1,

$$(1 - 2\eta)|V(G)| \leq \tau < |C'| + |X'| \leq |C| + |D| + |X'| \leq |V(G)|/2 + |D| + |X| \leq (0.5 + \eta)|V(G)| + |D|.$$

Note that Claim 4.8 in particular implies that $D \neq \emptyset$, that is, Ω is not completely contained in X .

CLAIM 4.9. $X = N(D)$.

PROOF. Clearly, $N(D) \subseteq X$. By the minimality of X , it suffices to show that $N(D)$ is a nuke in G . Condition (i) of Definition 4.1 holds trivially, and condition (ii) holds by the fact that $|N(D)| \leq |X|$. Hence, it remains to verify that condition (iii) holds. Consider




Fig. 4. Illustration of the proof of Claim 4.12. In this and all subsequent figures in this section the nuke is depicted gray. Furthermore, in this and the next two figures the component D is marked with gray dots.

a shelter D' of $G - N(D)$. If $D' = D$, then $|D'| + |N(D)| \leq |D| + |X| \leq \tau$ by the assumption that X is a nuke. Otherwise, by Claim 4.8 and the assumption $\eta \leq 0.1$ we have

$$|D'| + |N(D)| \leq |V(G) \setminus D| \leq (0.5 + 3\eta)|V(G)| \leq (1 - 2\eta)|V(G)| \leq \tau.$$

Hence, $N(D)$ is a nuke, which by the minimality of X implies that $X = N(D)$. \lrcorner

CLAIM 4.10. *For every nuked component C of $G - \Omega$ it holds that $N(C) \setminus X \neq \emptyset$, that is, there exists a non-nuked vertex in the neighbourhood of C .*

PROOF. A direct corollary from the facts that $X = N(D)$, D is connected, and contains vertices of Ω . \lrcorner

CLAIM 4.11. *For every $x \in X$ there exists a shelter D' of $G - X$ that is different from D and contains a vertex adjacent to x .*

PROOF. If that is not the case, then $X \setminus \{x\}$ is a nuke as well. Indeed, conditions (i) and (ii) of Definition 4.1 hold trivially. For condition (iii), note that each shelter of X is still a shelter of $X \setminus \{x\}$, except the shelter D , which becomes a shelter $D \cup \{x\}$. However, $|D \cup \{x\}| + |X \setminus \{x\}| = |D| + |X| \leq \tau$. Hence, $X \setminus \{x\}$ is a nuke, contradicting the minimality of X . \lrcorner

Our goal is now to exhibit a restricted structure of the nuked components of $G - \Omega$, using the fact that G is P_6 -free. Intuitively, every nuked component gives rise to a potential P_3 or even P_4 sticking into such a component; by combining two such paths we should obtain a forbidden P_6 . The next four observations assert the existence of such sticking out P_3 s and P_4 s.

CLAIM 4.12. *For every nuked component C of $G - \Omega$, and every $v \in N(C) \setminus X$, there exists a P_3 in G with one endpoint in v and the remaining two vertices in C .*

PROOF. See Fig. 4 for an illustration of the proof. Let $x \in C \cap X$, and let D' be a shelter of $G - X$ different from D and adjacent to x , whose existence is asserted by Claim 4.11. Since $N(C) \setminus X \subseteq D$, we have $D' \subseteq C$. Consequently, $D' \cap N(v) = \emptyset$, in particular C is not contained in $N(v)$. The existence of the asserted P_3 follows from the connectivity of C . \lrcorner

CLAIM 4.13. *For every nuked component C of $G - \Omega$ with $\mu(C) \geq 0.1$ and for every $v \in N(C)$ there exists a P_3 in G with one endpoint in v and the remaining two vertices in C .*

PROOF. If such a P_3 does not exist, by the connectivity of C we have $C \subseteq N(v)$. However, then $\mu(N(v)) \geq \mu(C) \geq 0.1$. \lrcorner

CLAIM 4.14. *For every nuked component C of $G - \Omega$, if there exists a vertex $x \in (C \cap X) \setminus N(\Omega \setminus X)$, then there exists a nonempty set $Z \subseteq N(C) \setminus X$ such that for every $v \in Z$ there exists a P_4 in G with one endpoint in v and the remaining three vertices in $C \setminus N(N(C) \setminus (X \cup Z))$.*

PROOF. See Fig. 5 for an illustration of the proof. Define Z to be the set of vertices in $N(C) \setminus X$ that are reachable from x via a path with all internal vertices in $C \cap D$. The fact that Z is nonempty follows from the facts that D is connected, $\emptyset \neq N(C) \setminus X \subseteq D$, and $x \in X = N(D)$. Note that $Z \subseteq \Omega$, because $Z \subseteq N(C)$ and C is a (nuked) component of $G - \Omega$.




Fig. 5. Illustration of the proof of Claim 4.14.




Fig. 6. Illustration of the proof of Claim 4.15.

Consider any $v \in Z$. Let P be a shortest path from v to x with all internal vertices in $C \cap D$. By the definition of Z , such a path exists. Since P is a shortest path, it is an induced one. Furthermore, since $Z \subseteq \Omega$ while $x \notin N(\Omega \setminus X)$, the path P contains at least three vertices. Prolong P with a neighbour of x in D' , a shelter different than D adjacent to x (whose existence is asserted by Claim 4.11), obtaining a path on at least four vertices with one endpoint in v and remaining vertices in C .

To finish the proof, it suffices to argue that no vertex of P except for v may have a neighbour in $N(C) \setminus (X \cup Z)$. This statement is true for the part of P contained in $C \cap D$, by the definition of Z . By assumptions, x has no neighbour in $\Omega \setminus X$. Finally, no vertex in D' is adjacent to any vertex of $N(C) \setminus (X \cup Z) \subseteq D$. \lrcorner

CLAIM 4.15. *For every nuked component C of $G - \Omega$, for every two vertices $u, v \in N(C) \setminus X$, if there exists a vertex $x \in C \cap X \cap (N(v) \setminus N(u))$, then there exists a P_3 in G with one endpoint in v and the remaining two vertices in $C \setminus N(u)$.*

PROOF. See Fig. 6 for an illustration of the proof. Let D' be a shelter of $G - X$, different from D and adjacent to x , whose existence is asserted by Claim 4.11. For the required P_3 , take the vertices v, x , and any vertex of $N(x) \cap D'$. \lrcorner

We now study the possible relations between the neighbourhoods of nuked components. The following observation serves as a starting point.

CLAIM 4.16. *For every two nuked components C_1, C_2 of $G \setminus \Omega$ it holds that $N(C_1) \setminus X \subseteq N(C_2) \setminus X$ or $N(C_2) \setminus X \subseteq N(C_1) \setminus X$.*

PROOF. See Fig. 7 for an illustration of the proof. By contradiction, assume that there exists $v_i \in N(C_i) \setminus (X \cup N(C_{3-i}))$ for $i = 1, 2$. For $i = 1, 2$, let P^i be a P_3 with endpoint in v_i and other vertices in C_i , whose existence is asserted by Claim 4.12. If $v_1 v_2 \in E(G)$, then concatenated paths P^1 and P^2 form a P_6 , a contradiction. Otherwise, by Theorem 2.3 there exists a component C of $G \setminus \Omega$ with $v_1, v_2 \in N(C)$. Clearly, $C \neq C_i$ for $i = 1, 2$. Hence, by concatenating P^1 , a shortest path from v_1 to v_2 through C , and P^2 , we obtain an induced path on at least 7 vertices, a contradiction. \lrcorner

Claim 4.16 allows us to order the nuked components of $G - \Omega$ as C_1, C_2, \dots, C_r , so that

$$N(C_1) \setminus X \supseteq N(C_2) \setminus X \supseteq \dots \supseteq N(C_r) \setminus X.$$




Fig. 7. Illustration of the proof of Claim 4.16. The dotted connection between v_1 and v_2 may be realized through a third component.




Fig. 8. Illustration of the proof of Claim 4.17. The existence of the dotted connection between v and u is implied by the linkedness of C_i and C_j .

By Claim 4.7, $r \geq 2$.

We say that two nuked components C_i and C_j , $1 \leq i, j \leq r$, $i \neq j$ are *linked* if for every choice of $u \in N(C_i) \setminus X$ and $v \in N(C_j) \setminus X$ there exists an induced path in G with endpoints u and v and all internal vertices in $V(G) \setminus N[C_i \cup C_j]$. We remark that if $u = v$ or $uv \in E(G)$, then the last assertion is true, as we can take an one- or two-vertex path, respectively.

In the next few observations we investigate the properties of linked components.

CLAIM 4.17. *If C_i and C_j are linked, then for every two vertices $u, v \in \Omega \setminus X$, one of the following holds:*

- (1) $N(u) \cap X \cap C_i = N(v) \cap X \cap C_i$,
- (2) $N(u) \cap X \cap C_j = N(v) \cap X \cap C_j$,
- (3) $N(u) \cap X \cap (C_i \cup C_j) \subsetneq N(v) \cap X \cap (C_i \cup C_j)$, or
- (4) $N(v) \cap X \cap (C_i \cup C_j) \subsetneq N(u) \cap X \cap (C_i \cup C_j)$.

PROOF. See Fig. 8 for an illustration of the proof. Assume the contrary. By symmetry, we can consider only the case where $(N(v) \cap C_i \cap X) \setminus N(u) \neq \emptyset$ and $(N(u) \cap C_j \cap X) \setminus N(v) \neq \emptyset$. Clearly, $v \in N(C_i) \setminus X$, $u \in N(C_j) \setminus X$, and $u \neq v$. By applying Claim 4.15 twice, we obtain a $P_3 P^v$ with endpoint in v and the remaining two vertices in $C_i \setminus N(u)$, and a $P_3 P^u$ with endpoint in u and the remaining two vertices in $C_j \setminus N(v)$. These two paths, together with the induced path between u and v promised by the fact that C_i and C_j are linked, yield an induced path on at least six vertices, a contradiction. \square

With every nuked component C_i we associate the family $\mathcal{F}_i := \{N(v) \cap C_i \cap X : v \in \Omega \setminus X\}$.

CLAIM 4.18. *If C_i and C_j are linked, then \mathcal{F}_i or \mathcal{F}_j has a unique maximal element with respect to inclusion.*

PROOF. Assume otherwise. Let $u, v \in \Omega \setminus X$ be such that $A_u := N(u) \cap C_i \cap X$ and $A_v := N(v) \cap C_i \cap X$ are two different maximal elements of \mathcal{F}_i , and $p, q \in \Omega \setminus X$ be such that $B_p := N(p) \cap C_j \cap X$ and $B_q := N(q) \cap C_j \cap X$ are two different maximal elements of \mathcal{F}_j . By Claim 4.17 we have $N(u) \cap C_j \cap X = N(v) \cap C_j \cap X$; let us denote this set by B . Similarly, $N(p) \cap C_i \cap X = N(q) \cap C_i \cap X$,

and we denote this set A . If B and B_p are incomparable with respect to inclusion, then Claim 4.17 asserts that $A = A_u$ (for pair u and p) and $A = A_v$ (for pair v and p), a contradiction. By maximality of B_p , we have $B \subseteq B_p$. Similarly we infer that $B \subseteq B_q$. Hence, $B \subseteq B_p \cap B_q$; by the incomparability of B_p and B_q , we infer that $B \subseteq B_p$. However, Claim 4.17 asserts then that $A_u \subseteq A$ (for the pair p, u) and $A_v \subseteq A$ (for the pair p, v). This is a contradiction with the maximality and incomparability of A_u and A_v . \square

CLAIM 4.19. *Let $I \subseteq \{1, 2, \dots, r\}$ be a set of indices such that for any $i, j \in I$, $i \neq j$, C_i and C_j are linked. Then there exists a vertex $v \in \Omega \setminus X$ and an index i_0 such that*

$$X \cap N(\Omega \setminus X) \cap \left(\bigcup_{i \in I \setminus \{i_0\}} C_i \right) \subseteq N(v).$$

PROOF. If $|I| \leq 1$, the claim is straightforward, so assume otherwise. By Claim 4.18, there exists at most one index i_0 such that \mathcal{F}_{i_0} does not admit a unique maximal element. (If no such index exists, we set $i_0 \in I$ arbitrarily).

For $u \in \Omega \setminus X$, we define $I_u \subseteq I \setminus \{i_0\}$ to be the set of indices $i \in I \setminus \{i_0\}$ for which $N(u) \cap C_i \cap X$ is the unique maximal element of \mathcal{F}_i ; note that this unique maximal element exists because $i \in I \setminus \{i_0\}$. Let v be a vertex such that $|I_v|$ is maximized.

We prove that $I_v = I \setminus \{i_0\}$. Assume the contrary: there exists $j \in I \setminus \{i_0\}$ such that $N(v) \cap C_j \cap X$ is not the maximal element of \mathcal{F}_j . Let $w \in \Omega \setminus X$ be such that $N(w) \cap C_j \cap X$ is the maximal element of \mathcal{F}_j . We have $N(v) \cap C_j \cap X \subsetneq N(w) \cap C_j \cap X$. By Claim 4.17, for every $i \in I_v$ we have $N(v) \cap C_i \cap X \subseteq N(w) \cap C_i \cap X$. However, $N(v) \cap C_i \cap X$ is the unique maximal element of \mathcal{F}_i . Consequently, $I_v \subseteq I_w$. However, $j \in I_w \setminus I_v$, which contradicts the choice of v .

Using that $I_v = I \setminus \{i_0\}$, $N(v) \cap C_i \cap X$ is the maximal element of \mathcal{F}_i for each $i \in I \setminus \{i_0\}$, meaning that $N(u) \cap C_i \cap X \subseteq N(v) \cap C_i \cap X \subseteq N(v)$ for each $u \in \Omega \setminus X$ and each $i \in I \setminus \{i_0\}$. Hence, $X \cap N(\Omega \setminus X) \cap \bigcup_{i \in I \setminus \{i_0\}} C_i \subseteq N(v)$, as claimed. \square

Consider now the following corollary of Claim 4.14.

CLAIM 4.20. *For every $2 \leq i \leq r$ we have $X \cap C_i \subseteq N(\Omega \setminus X)$. Furthermore, if $X \cap C_1 \not\subseteq N(\Omega \setminus X)$, then the set Z whose existence is asserted in Claim 4.14 for the component C_1 is completely contained in $N(C_1) \setminus (X \cup N(C_2))$.*

PROOF. See Fig. 9 for an illustration of the proof. Let $1 \leq i \leq r$ and suppose that $X \cap C_i \not\subseteq N(\Omega \setminus X)$ (for $i \geq 2$, this is for the sake of contradiction; for $i = 1$, this is the assumption in the claim statement). Let $Z_i \subseteq N(C_i) \setminus X$ denote the nonempty set whose existence is asserted by Claim 4.14. If $i = 1$, then also assume for the sake of contradiction that $Z_1 \not\subseteq N(C_1) \setminus (X \cup N(C_2))$. Let $j \in \{1, 2\} \setminus \{i\}$ (C_j exists, because $r \geq 2$ by Claim 4.7).

We first observe that $Z_i \cap (N(C_j) \setminus X) \neq \emptyset$. If $i > j$, then this follows from the fact that $Z_i \subseteq N(C_i) \setminus X$ by definition (see Claim 4.14) and that $N(C_i) \setminus X \subseteq N(C_j) \setminus X$ by the ordering on the nuked components. If $i < j$, then $i = 1$, $j = 2$, and Z_1 contains a vertex $u \notin N(C_1) \setminus (X \cup N(C_2))$ by assumption, and thus $u \in N(C_2) \setminus X$, because $Z_1 \subseteq N(C_1) \setminus X$. This proves the observation.

Now let $v \in Z_i \cap (N(C_j) \setminus X)$ (v exists by the above observation). By Claim 4.14 and the fact that $v \in Z_i$, there exists a P_4 in G with one endpoint in v and the remaining three vertices in C_i . By Claim 4.12 and the fact that $v \in N(C_j) \setminus X$, there exists a P_3 in G with one endpoint in v and the remaining two vertices in C_j . The concatenation of these two paths is a P_6 in G , a contradiction. \square

Observe now the following.

CLAIM 4.21. *For every $2 \leq i, j \leq r$, $i \neq j$, the components C_i and C_j are linked.*

PROOF. Two vertices $u \in N(C_i) \setminus X$ and $v \in N(C_j) \setminus X$ can be linked either via a direct edge if it exists in G , or via a shortest path with internal vertices in C_1 . \square

Combining now Claim 4.20 with Claim 4.19 applied to $I = \{2, 3, \dots, r\}$ we obtain that

CLAIM 4.22. *There exists an index $i_0 \in \{2, \dots, r\}$ such that $\mu(C_1 \cup C_{i_0}) \geq 0.8$. In particular, $\mu(C_1), \mu(C_{i_0}) \geq 0.3$.*

PROOF. By Claim 4.19, applied to $I = \{2, 3, \dots, r\}$, we have an index i_0 and a vertex v adjacent to all vertices of $X \cap C_j \cap N(\Omega \setminus X)$ for $j \notin \{1, i_0\}$. However, by Claim 4.20, these are actually all vertices of $X \cap C_j$. Since $\mu(N(v)) < \gamma \leq 0.1$ and $\mu(\Omega) \leq 0.1$, the first claim follows. The second claim follows from the choice of Ω : $\mu(C) \leq 0.5$ for every connected component C of $G - \Omega$. \square




Fig. 9. Illustration of the proof of Claim 4.20. The set $X \cap \Omega$ is omitted in order to keep the picture clear.




Fig. 10. Illustration of the proof of Claim 4.23. The dotted connection between v and u may be realized through a third component.

Fix the index i_0 from Claim 4.22.

CLAIM 4.23. $N(C_1) \cup N(C_{i_0}) \neq \Omega$.

PROOF. See Fig. 10 for an illustration of the proof. By contradiction, assume that $N(C_1) \cup N(C_{i_0}) = \Omega$. By Theorem 2.3, neither $N(C_1)$ nor $N(C_{i_0})$ equals the whole Ω , thus there exists $v \in N(C_1) \setminus N(C_{i_0})$ and $u \in N(C_{i_0}) \setminus N(C_1)$. By Claim 4.13, there exist a P_3 P^v with endpoint in v and remaining two vertices in C_1 , and a P_3 P^u with endpoint in u and remaining two vertices in C_{i_0} . If $uv \in E(G)$, then these two paths together give a P_6 in G , a contradiction. Otherwise, by Theorem 2.3, there exists a component C of $G - \Omega$ such that $u, v \in N(C)$. Clearly, $C \notin \{C_1, C_{i_0}\}$. However, then a concatenation of P^v , a shortest path from v to u with internal vertices in C , and P^u , yields an induced path in G on at least 7 vertices, a contradiction. \square

CLAIM 4.24. C_1 and C_{i_0} are linked.

PROOF. Consider any $v \in N(C_1) \setminus X$ and $u \in N(C_{i_0}) \setminus X$; we are going to exhibit an induced path from v to u with internal vertices in $V(G) \setminus N[C_1 \cup C_{i_0}]$. If $v = u$ or $vu \in E(G)$, then we are done with the one- or two-vertex path. If there exists a connected component C of $G - \Omega$ different than C_1 or C_{i_0} such that $u, v \in N(C)$, then we can choose a shortest path from v to u with all internal vertices in C .

Otherwise, let w be an arbitrary vertex of $\Omega \setminus N(C_1 \cup C_{i_0})$, whose existence follows from Claim 4.23, and we route a path through the vertex w . Let $P^v = vw$ if $vw \in E(G)$, and otherwise let P^v be a shortest path from v to w with internal vertices in a connected component C covering the nonedge vw ; note that $C \notin \{C_1, C_{i_0}\}$ as $w \notin N(C_1 \cup C_{i_0})$. Similarly define the path P^u from u




Fig. 11. Illustration of the proof of Claim 4.25. The vertex u is not adjacent to any of the three vertices in C_1 since $u \notin Z$ (Claim 4.20). The existence of a connection between u and v is guaranteed by the linkedness of C_1 and C_{i_0} .

to w . Since no component different than C_1 or C_{i_0} has both u and v in their neighbourhood, the concatenation of P^v and P^u forms the desired path. \lrcorner

In the next two claims we exhibit the final contradiction.

CLAIM 4.25. $C_1 \cap X \subseteq N(\Omega \setminus X)$.

PROOF. See Fig. 11 for an illustration of the proof. Assume the contrary. By Claim 4.20, the set Z whose existence is asserted by Claim 4.14 for the component C_1 is completely contained in $N(C_1) \setminus (X \cup N(C_2)) \subseteq N(C_1) \setminus (X \cup N(C_{i_0}))$. Consider any $v \in Z$ and $u \in N(C_{i_0}) \setminus X \subseteq N(C_1) \setminus (X \cup Z)$. By Claim 4.14, there exists a $P_4 P^v$ with endpoint in v and internal vertices in $C_1 \setminus N(u)$. Furthermore, by Claim 4.13, there exists a $P_3 P^u$ with endpoint u and internal vertices in C_{i_0} . Recall that $v \notin N(C_{i_0})$, thus P^u does not contain any neighbour of v , except for possibly u . Hence, the paths P^v and P^u , together with the path between v and u whose existence is asserted by the fact that C_1 and C_{i_0} are linked following Claim 4.24, form an induced path in G on at least seven vertices, a contradiction. \lrcorner

CLAIM 4.26. *There exists a vertex v such that $C_1 \cap X \subseteq N(v)$ or $C_{i_0} \cap X \subseteq N(v)$.*

PROOF. By Claim 4.19 applied to $I = \{1, i_0\}$, we obtain a vertex v such that $C_1 \cap X \cap N(\Omega \setminus X) \subseteq N(v)$ or $C_{i_0} \cap X \cap N(\Omega \setminus X) \subseteq N(v)$. However, $C_{i_0} \cap X \subseteq N(\Omega \setminus X)$ due to Claim 4.20 and $C_1 \cap X \subseteq N(\Omega \setminus X)$ due to Claim 4.25. \lrcorner

We now complete the proof of Theorem 4.2. Recall that by Claim 4.22, we have $\mu(C_1), \mu(C_{i_0}) \geq 0.3$. Moreover, by Claim 4.26, there exists a vertex v such that $C_1 \cap X \subseteq N(v)$ or $C_{i_0} \cap X \subseteq N(v)$. Since μ is a probability measure over X , this implies that $\mu(N(v)) \geq 0.3$. This contradicts our assumption that $\mu(N(u)) < \gamma \leq 0.1$ for every $u \in V(G)$, and thus finishes the proof.

5 THE ALGORITHM FOR MAXIMUM WEIGHT INDEPENDENT SET

We now make use of Theorems 1.4 and 4.2 to design an algorithm that solves MAXIMUM WEIGHT INDEPENDENT SET in n -vertex P_6 -free graphs in $n^{O(\log^2 n)}$ time.

5.1 Description of the algorithm

The algorithm consists of two recursive procedures, FindIS and FindISNuke, which both aim to find an independent set of maximum weight in a given connected vertex-weighted P_6 -free graph G . The procedure FindIS is the ‘base’ procedure, which we call on the graph G . Both procedures make recursive calls to themselves and to each other. We describe each procedure, and then analyze their running time.

5.1.1 Procedure FindIS. The input for the procedure FindIS is just a connected P_6 -free graph G . As a base case, if the input graph consists of one vertex, FindIS returns the weight of this vertex. Otherwise, it checks if there exists a vertex of degree at least $0.05\beta|V(G)|$, where the constant β comes from Theorem 1.4.

If such a vertex v exists, then the procedure branches on the vertex v . In one branch, we seek a solution not containing v , and we call FindIS independently on every connected component of $G - v$. In the second branch, we seek a solution containing v , and we call FindIS independently on every connected component of $G - N(v)$.

Otherwise, that is if all vertices are of degree less than $0.05\beta|V(G)|$, the algorithm takes an arbitrary minimal triangulation \widehat{G} of G (see e.g. [37] for algorithms that find such a triangulation), constructs its clique tree, and finds a maximal clique Ω in \widehat{G} such that every connected component of $G - \Omega$ has at most $|V(G)|/2$ vertices (such a maximal clique can be found in polynomial time by the arguments in Claim 4.6). We observe the following:

CLAIM 5.1. $|\Omega| < 0.05|V(G)|$

PROOF. If $|\Omega| \geq 0.05|V(G)|$, then Theorem 1.4 applied to Ω with the uniform measure, implies that there exists a vertex v with $|N(v)| \geq |N(v) \cap \Omega| \geq 0.05\beta|V(G)|$, a contradiction. \square

By the choice of Ω , if we set $\tau = 0.8|V(G)|$, then Ω is a $(0.1, \tau)$ -nuke in G (with a lot of slack in the inequalities in condition (ii) and (iii) of Definition 4.1). The algorithm passes the graph G , the threshold τ , and the nuke Ω to the procedure FindISNuke.

5.1.2 Procedure FindISNuke. The input for the procedure FindISNuke is a connected P_6 -free graph G , a threshold τ , and a set $X \subseteq V(G)$ with the promise that for every connected component C of $G - X$ it holds that $|C| + |X| \leq \tau$.

The algorithm first checks if G contains at least two vertices and X is a non-empty $(0.1, \tau)$ -nuke of G (note that X is such a nuke when FindISNuke is invoked by FindIS). If this is not the case, then the algorithm invokes FindIS on the graph G , forgetting about τ and X . Otherwise, it finds any inclusion-wise minimal $(0.1, \tau)$ -nuke $Y \subseteq X$, and finds a vertex v with $|N(v) \cap Y| \geq \gamma|Y|$; the existence of such vertex is guaranteed by applying Theorem 4.2 to Y with the uniform measure. The algorithm branches on vertex v as usual. That is, in one branch, we seek a solution not containing v , and we call FindISNuke independently on every connected component of $G - v$. In the second branch, we seek a solution containing v , and we call FindISNuke independently on every connected component of $G - N(v)$. In every subcall, we pass the same threshold τ , and the set Y restricted to the vertex set of the connected component in question. Clearly, since we delete only vertices from G or reduce X to a minimal sub-nuke, in the subcalls we maintain the promise that for every connected component C of $G - X$ it holds that $|C| + |X| \leq \tau$.

5.2 Analysis

As the algorithm performs exhaustive branching, it clearly returns an optimum solution. Also, the polynomial space bound is immediate. It remains to argue about the running time.

Consider the recursion tree \mathcal{T}_0 of the algorithm, and focus on one call c to FindIS(G) that resulted in a subcall FindISNuke(G, τ, X); here $\tau = 0.8|V(G)|$ and X is a potential maximal clique in G of size at most $0.05|V(G)|$ (by Claim 5.1). Every call to FindISNuke results either in branching and multiple calls to the same procedure (call it a *branching call*), or in a single call to FindIS (call it a *fallback call*). Let \mathcal{T} be a maximal subtree at \mathcal{T}_0 , rooted at the chosen call c to FindIS, that contains (apart from the root) only calls to FindISNuke. That is, we put into \mathcal{T} all recursive calls that originated from c , and stop whenever we encounter a fallback call; in particular, all leaves of \mathcal{T} are fallback calls.

CLAIM 5.2. \mathcal{T} has $|V(G)|^{O(\log |V(G)|)}$ leaves.

PROOF. In every branch of FindISNuke, we either A) delete one vertex from G , or B) delete a constant fraction of the minimal sub-nuke Y of X (recall that the vertex v we branch on was chosen so that $|N(v) \cap Y| \geq \gamma|Y|$). Color the edges of \mathcal{T} either by 'A' or 'B' depending on the type of branch. Note that the path P_ℓ from a leaf ℓ of \mathcal{T} to the root of \mathcal{T} can contain at most $\log |V(G)|$ edges of \mathcal{T} colored 'B', because FindISNuke stops when the nuke is empty and each 'B' branch reduces the size of the nuke by a constant fraction. Moreover, each path P_ℓ has length at most $|V(G)|$, because each branch reduces the number of vertices of G by at least 1, and has a unique coloring, because the edges to the children of a vertex of \mathcal{T} are colored distinctly. Hence, \mathcal{T}

has $|V(G)|^{O(\log |V(G)|)}$ leaves. The fact that the algorithm independently considers every connected component only helps in the analysis. \square

Let $\text{FindISNuke}(G', \tau', X')$ be a leaf of \mathcal{T} . We claim the following.

CLAIM 5.3. $|V(G')| < \frac{8}{9}|V(G)|$.

PROOF. Since we are considering a fallback call, either $|V(G')| = 1$, $|X'| = 0$, or X' is not a $(0.1, \tau')$ -nuke of G' . In the first case, since $|V(G)| > 1$, the claim is obvious. In the second case, by the fact that the input graph of any call to FindISNuke is connected and the promise maintained in the course of algorithm that for every connected component C of $G' - X'$ it holds that $|C| + |X'| \leq \tau$, the claim also holds.

In the third case, consider the reasons why X' may not be a $(0.1, \tau')$ -nuke of G' . Clearly, $\tau' = \tau = 0.8|V(G)|$ and, by the promise maintained in the course of algorithm that for every connected component C of $G' - X'$ it holds that $|C| + |X'| \leq \tau$. Furthermore, $(1 - 2 \cdot 0.1)|V(G')| \leq (1 - 2 \cdot 0.1)|V(G)| = \tau$. Hence, either $(1 - 0.1)|V(G')| < \tau = 0.8|V(G)|$ or $|X'| > 0.1|V(G')|$. In the first case $|V(G')| < \frac{8}{9}|V(G)|$, while in the second case $|V(G')| \leq |V(G)|/2$, because $|X| \leq 0.05|V(G)|$ and $X' \subseteq X$. \square

By Claim 5.3, if we contract every such subtree \mathcal{T} to a single super-node of the recursion tree $\overline{\mathcal{T}}_0$, then at each such super-node we branch into $|V(G)|^{O(\log |V(G)|)}$ subcases, and in each subcase decrease the number of vertices by a multiplicative factor.

Now focus on a call \mathfrak{c} to FindIS that branches on a vertex $v \in V(G)$ of degree at least $0.05\beta|V(G)|$. Observe that at most one recursive subcall of \mathfrak{c} is invoked on a graph with at least $(1 - 0.05\beta)|V(G)|$ vertices: the one for the largest connected component of $G - v$. Mark the edges of the recursion tree that correspond to such subcalls. The marked edges form vertex-disjoint top-bottom paths in the recursion tree. If we contract them (along with the aforementioned subtrees \mathcal{T}), we obtain a recursion tree where every node has $|V(G)|^{O(\log |V(G)|)}$ subcases and where in each subcase the number of vertices decreases by a constant factor. Consequently, the size of the recursion tree is $|V(G)|^{O(\log^2 |V(G)|)}$. This finishes the analysis of the algorithm, and concludes the proof of Theorem 1.1.

6 THE ALGORITHM FOR MAXIMUM WEIGHT EFFICIENT DOMINATING SET

In this section we prove Theorem 1.2. The overall approach is as follows: we take any minimal triangulation of the input graph G , and perform the standard dynamic programming algorithm on the clique tree of this triangulation (which is a tree decomposition of G). In this standard dynamic programming algorithm, every state at bag B keeps information about which vertices of B are contained in the constructed efficient dominating set, and which vertices of B have already been dominated by the forgotten parts of the graph.

The main insight is that we can use Theorem 1.4, together with technical insight from the proof of Theorem 1.1, to show that in P_6 -free graphs there are only polynomially many reasonable states for the aforementioned dynamic programming algorithm. This then yields the claimed polynomial-time algorithm.

6.1 Bounding the Number of States

Before we state this main result formally, we need the following definition. Let Ω be a potential maximal clique in G , and let C be the set of connected components of $G - \Omega$. A *state* is a function $f : \Omega \rightarrow C \cup \{\Omega, \perp\}$. A state f is *consistent* with an efficient dominating set X if $X \cap \Omega = f^{-1}(\perp)$ and furthermore, for every $v \in \Omega \setminus X$, the unique vertex of $N(v) \cap X$ belongs to the vertex set of $f(v)$.

THEOREM 6.1. *Given a P_6 -free graph G and a potential maximal clique Ω in G , one can in polynomial time compute a family \mathcal{S} of states of polynomial size, such that for every efficient dominating set X in G , there exists a state $f \in \mathcal{S}$ consistent with X .*

This subsection is devoted to the proof of Theorem 6.1. We describe the algorithm as a branching algorithm that outputs a state at every leaf of the branching tree, and every leaf-to-root path of the branching tree contains $O(\log n)$ nodes of constant degree and $O(1)$ nodes of degree polynomial in n . Furthermore, it will be straightforward to perform the computation required at every node of the branching tree in polynomial time. These properties give the promised polynomial bounds on the size of the output and the total running time.

Every node of the branching tree is labeled with two vertex sets X_0 and Y , and the goal of the subtree rooted at the node labeled (X_0, Y) is to output a family of states such that for every efficient dominating set X with $X_0 \subseteq X$ and $(X \setminus X_0) \subseteq Y$ (henceforth called an efficient dominating set *consistent* with (X_0, Y)) there exists an output consistent state. In every branching step, in every subcase, the algorithm puts some vertices into X_0 and/or removes some vertices from Y . Since every two elements of an efficient dominating set are within distance at least three, we implicitly assume that if the algorithm puts a vertex v into X_0 , it at the same time removes from Y all vertices within distance at most two from v . Furthermore, we immediately terminate a branch if two vertices of X_0 are within distance less than three, or if there exists $v \in V(G)$ with $N[v] \cap (X_0 \cup Y) = \emptyset$.

The algorithm terminates branching at nodes labeled (X_0, Y) where for every $v \in \Omega$ either $N[v]$ contains a vertex of X_0 , or $N[v] \cap Y$ is contained in a single component of C . For such a label (X_0, Y) , we define a state f as follows: $f(v) = \perp$ for $v \in X_0 \cap \Omega$, $f(v) = \Omega$ for $v \in N(X_0) \cap \Omega$, and otherwise $f(v)$ is the unique component of C that contains vertices of $N[v] \cap (X_0 \cup Y)$. It is straightforward to verify that if X is consistent with (X_0, Y) , then f is well-defined and f is also consistent with X . Consequently, the algorithm outputs the function f in this leaf node of the branching tree.

At the root of the branching tree we have $X_0 = \emptyset$ and $Y = V(G)$.

6.1.1 Guessing Vertices from the Solution Inside the Potential Maximal Clique. We start with the following observation.

LEMMA 6.2. *For every P_7 -free graph G , every potential maximal clique Ω in G , and every efficient dominating set X in G , we have $|\Omega \cap X| \leq 1/\beta$, where the constant β comes from Theorem 1.4.*

PROOF. Without loss of generality, we can assume that G is connected (we can consider every component independently) and contains at least two vertices (for one-vertex graphs the statement is trivial).

Let $\ell = |\Omega \cap X|$. Consider a measure μ on Ω such that $\mu(v) = 1/\ell$ if $v \in \Omega \cap X$ and $\mu(v) = 0$ otherwise. By Theorem 1.4, there exists a vertex u with $\mu(N(u)) \geq \beta$. However, by the definition of an efficient dominating set, we have $|N(u) \cap X| \leq 1$. Consequently, $\mu(N(u)) \leq 1/\ell$, hence $\ell \leq 1/\beta$. \square

By Lemma 6.2, our algorithm can, as a first step, guess all vertices from the solution that lie in Ω . More formally, the algorithm branches into a subcase for every subset $X_\Omega \subseteq \Omega$ of size at most $1/\beta$; we label the subcase corresponding to X_Ω by $(X_\Omega, V(G) \setminus (N^2[X_\Omega] \cup \Omega))$. We emphasize here that we not only removed from Y all vertices within distance at most two from X_Ω , but also all vertices from Ω . Thus, from this point, we have that $Y \cap \Omega = \emptyset$.

6.1.2 Reduction Rule. Fix a node of the branching tree labeled (X_0, Y) with $Y \cap \Omega = \emptyset$. We say that a component $C \in C$ is *active* if $C \cap Y \neq \emptyset$. Let $A = \Omega \setminus N[X_0]$ be the set of vertices that are not yet dominated by the vertices from X_0 . Let $B \subseteq A$ be the set of these vertices v such that the vertices of $N(v) \cap Y$ appear in at least two connected components of C . Note that the algorithm terminates branching and outputs a state if $B = \emptyset$; the main goal in the branching step is to shrink the set B as much as possible.

We start by introducing a reduction rule, aimed at shrinking the set Y without performing any branching. For a vertex $u \in V(G) \setminus \Omega$, let $C(u)$ be the component of C that contains u . Assume that for some vertex $v \in B$ there exists $u \in N(v) \cap Y$ such that $N[u] \cap Y \subseteq N(v)$. Let X be an efficient dominating set consistent with (X_0, Y) . Since $Y \cap N[X_0] = \emptyset$, the vertex u is dominated by some vertex $w \in N[u] \cap (X \setminus X_0) \subseteq N[u] \cap Y$. By our assumption, w also dominates v . Consequently, in every efficient dominating set consistent with (X_0, Y) , the vertex v is dominated by an element of $C(u)$, and we can introduce the following reduction rule.

Reduction Rule. If there exist vertices $v \in B$ and $u \in N(v) \cap Y$ such that $N[u] \cap Y \subseteq N(v)$, then remove from Y all vertices of $N(v) \setminus C(u)$.

Note that the Reduction Rule removes neither v , nor u , nor the vertex dominating them in any efficient dominating set consistent with (X_0, Y) . Furthermore, an application of the Reduction Rule on a vertex $v \in B$ effectively removes v from the set B , as it leaves in Y only the neighbors of v that are in the component $C(u)$. Also observe that the Reduction Rule triggers on any vertex of B that is fully adjacent to an active component (recall that $\Omega \cap Y = \emptyset$).

In what follows we assume that at every node of the recursion tree, the Reduction Rule is applied exhaustively. Observe that if this rule is not applicable, then for every $v \in B$ and $u \in N(v) \cap Y$, there exists a vertex $w \in (Y \cap N(u)) \setminus N(v)$; note that $w \in C(u)$




Fig. 12. Illustration of the proof of Lemma 6.3. The dotted connection between v^1 and v^2 may be realized through a third component.

and $\{v, u, w\}$ induce a P_3 in G . The main intuition of the remaining proof is that the graph needs to be highly structured in order to not to allow two such P_3 's to "glue" together into a P_6 in G .

6.1.3 *Structure of B -Neighbourhoods.* As a first application of this principle, observe the following.

LEMMA 6.3. *If C^1, C^2 are two different components of C , then $N(C^1) \setminus N(C^2)$ is fully adjacent to C^1 , or $N(C^2) \setminus N(C^1)$ is fully adjacent to C^2 .*

PROOF. See Fig. 12 for an illustration of the proof. Assume the contrary. Let $v^i \in N(C^i) \setminus N(C^{3-i})$ be a vertex that is not fully adjacent to C^i for $i = 1, 2$. Since v^i is not fully adjacent to C^i , but $v^i \in N(C^i)$ and C^i is connected, there exists an induced P_3 with one endpoint v^i and other vertices in C^i ; denote this P_3 as P^i . Furthermore, by Theorem 2.3, either $v^1 v^2 \in E(G)$ or there exists a component $C \in C$ such that $v^1, v^2 \in N(C)$. Clearly, $C \notin \{C^1, C^2\}$. Consequently, by concatenating P^1, P^2 , and the edge $v^1 v^2$ or a shortest path between v^1 and v^2 with internal vertices in C , we obtain an induced path on at least 6 vertices, a contradiction. \square

Since our Reduction Rule removes from B vertices that are fully adjacent to some active component, we infer from Lemma 6.3 that we can enumerate active components as C_1, C_2, \dots, C_r such that $N_B(C_i) \supseteq N_B(C_j)$ for every $i \leq j$. Furthermore, since every element in B has neighbours in Y in at least two components by definition, we have that $N_B(C_1) = N_B(C_2) = B$. Summing up,

$$B = N_B(C_1) = N_B(C_2) \supseteq N_B(C_3) \supseteq \dots \supseteq N_B(C_r). \quad (6.1)$$

6.1.4 *Obtaining Linkedness.* In order to "glue" two P_3 's, we use the following notion. We say that two active components C^1 and C^2 are *linked* if for every two vertices $v^1 \in B \cap N(C^1)$, $v^2 \in B \cap N(C^2)$ there exists an induced path in G with endpoints v^1 and v^2 and all internal vertices in $V(G) \setminus (N[C^1] \cup N[C^2])$. We explicitly allow 1-vertex and 2-vertex paths here (if $v^1 = v^2$ or $v^1 v^2 \in E(G)$).

We start by observing the following:

LEMMA 6.4. *Every pair of active components is linked, except for possibly the pair $\{C_1, C_2\}$.*

PROOF. By (6.1), for every other pair $\{C_i, C_j\}$, we can use either C_1 or C_2 to route the desired path. \square

Our goal now is to ensure that also $\{C_1, C_2\}$ are linked. The following lemma uses essentially the same arguments as Claim 4.24 of Section 4.

LEMMA 6.5. *If two active components C^1 and C^2 satisfy $N(C^1) \cup N(C^2) \neq \Omega$, then they are linked.*

PROOF. Let $w \in \Omega \setminus (N(C^1) \cup N(C^2))$ and consider two vertices $v^1 \in B \cap N(C^1)$, $v^2 \in B \cap N(C^2)$. If $v^1 = v^2$ or $v^1 v^2 \in E(G)$, then we are trivially done. Furthermore, if there exists a component $C \in C \setminus \{C^1, C^2\}$ with $v^1, v^2 \in N(C)$, then we are done as well by taking a shortest path from v^1 to v^2 with all internal vertices in C .

In the remaining case, we start with connecting for $i = 1, 2$ the vertex v^i with w by an induced path P^i as follows: if $v^i w \in E(G)$, then we take P^i to be this edge only, while otherwise we take a component $D^i \in C$ with $v^i, w \in N(D^i)$ (whose existence is

promised by Theorem 2.3) and take as P^i a shortest path from v^i to w with internal vertices in D^i . Note that $D^i \notin \{C^1, C^2\}$, since $w \in N(D^i)$. Furthermore, $v^{3-i} \notin D^i$, as no component other than C^1 and C^2 can neighbour both v^1 and v^2 . Consequently, $D^1 \neq D^2$, and the concatenation of P^1 and P^2 gives the desired path from v^1 to v^2 . \square

By Lemma 6.5, the pair $\{C_1, C_2\}$ is linked unless $N(C_1) \cup N(C_2) = \Omega$. However, if this is the case, then by Theorem 2.3 we have that both $N(C_1) \setminus N(C_2)$ and $N(C_2) \setminus N(C_1)$ are nonempty. By Lemma 6.3, there exists $i \in \{1, 2\}$ and a vertex $v^i \in \Omega$ that is fully adjacent to C_i . Consequently, every efficient dominating set consistent with (X_0, Y) contains exactly one vertex of C_i : it needs to contain at least one to dominate $Y \cap C_i$, but at most one since every vertex of C_i dominates v^i .

We branch into $|Y \cap C_i|$ directions, guessing the vertex from $Y \cap C_i$ that belongs to the solution, and putting it into X_0 . Furthermore, in every branch we remove from Y all vertices of $Y \cap C_i$. In every subcase, C_i is no longer an active component, but witnesses that every two other components that remain active are linked: since $B \subseteq N(C_i)$, we can always route a path between the desired endpoints through C_i .

By the above analysis and branching step, we can assume henceforth that any pair of active components is linked.

6.1.5 Branching on Bad Vertices. Partition Y into $Y_1 = \{y \in Y : |N_B(y)| \geq |B|/16\}$ and $Y_2 = Y \setminus Y_1$. Let Y_1^* be the set of vertices $y \in Y_1$ for which the addition of y to the solution (i.e., to X_0) and the subsequent exhaustive application of the Reduction Rule reduces B to an empty set. Let $Y_1^\circ = Y_1 \setminus Y_1^*$.

If we would know that some vertex of Y_1^* belongs to the solution, then we could just guess it and the Reduction Rule would reduce the set B completely (we analyze this more precisely in the next subsection). In this subsection we focus on the analysis of the set of “bad” vertices Y_1° , showing that any such vertex also gives ground to a good branching — but in a completely different fashion.

Let $y \in Y_1^\circ$ be chosen arbitrarily. Define Y° and B° as follows: if we add y to X_0 and exhaustively apply the Reduction Rule, then let Y° be the set to which Y is shrunk and $B^\circ \neq \emptyset$ the set to which B is shrunk. It is important to note that we do not actually add y to X_0 ; we only aimed to define Y° and B° . We claim the following:

LEMMA 6.6. *For every $z \in N[y] \cap Y$, it holds that $N_B(y) \subseteq N(z)$ or $B^\circ \subseteq N(z)$.*

PROOF. See the left panel of Fig. 13 for an illustration of the proof. Fix a vertex z as in the statement, and assume the contrary: there exist $p \in N_B(y) \setminus N(z)$ and $q \in B^\circ \setminus N(z)$. Note that $z \neq y$, as $p \in N_B(y)$ and $p \notin N(z)$. Since $q \in B^\circ$, there exists at least two components of C that contain vertices of $N(q) \cap Y^\circ$. Let C_q be one of these components that is different from $C(y)$, and let $s \in N(q) \cap C_q \cap Y^\circ$. Since the Reduction Rule does not trigger on q and s after y has been put into X_0 , there exists $t \in (N(s) \cap Y^\circ) \setminus N(q)$; clearly, t lies also in C_q .

Observe that q, s, t induce a P_3 , and that s and t are not adjacent to $p \in N_B(y)$ as $s, t \in Y^\circ$. Furthermore, p, y, z induce a P_3 , while y and z are not adjacent to q . Since C_q and $C(y)$ are linked, we can connect p and q by an induced path avoiding $N[C_q] \cup N[C(y)]$, giving together with the aforementioned P_3 ’s an induced path on at least six vertices, a contradiction. \square

Lemma 6.6 allows us to branch into two directions, deciding whether the element of the sought efficient dominating set that dominates the vertex y also dominates the set $N_B(y)$ or the set B° . That is, in the first subcase, we delete from Y all vertices of $N(N_B(y)) \setminus C(y)$, while in the second subcase, we delete from Y all vertices of $N(B^\circ) \setminus C(y)$. We claim that in both subcases, after exhaustively applying the Reduction Rule, the size of B decreased at least by a multiplicative factor of $1 - 1/16$.

Clearly this is the case in the first subcase, as then $N_B(y)$ is removed from B and $|N_B(y)| \geq |B|/16$ since $y \in Y_1$. We claim the following:

LEMMA 6.7. *In the second subcase, the Reduction Rule also removes the entire set $N_B(y)$ from B .*

PROOF. See the right panel of Fig. 13 for an illustration of the proof. Assume that this is not the case. Let B', Y' be the reduced sets in the second subcase, and let $p \in N_B(y) \cap B'$. Since $p \in B'$, there exist at least two components of C that contain vertices of $N(p) \cap Y'$; let C_p be such component different from $C(y)$, and let $a \in C_p \cap N(p) \cap Y'$. Since the Reduction Rule does not trigger on p and a , given sets B' and Y' , there exists $b \in (N(a) \cap Y') \setminus N(p)$; clearly also $b \in C_p$.




Fig. 13. Illustration of the proof of Lemma 6.6 (left) and of Lemma 6.7 (right).

Consider now any $q \in B^\circ$. Since $a, b \in Y' \setminus C(y)$, we have that a and b are not adjacent to q . Furthermore, since $q \in B^\circ$, there exist at least two components of C that contain vertices of $N(q) \cap Y^\circ$; let C_q be such a component different than C_p , and let $s \in C_q \cap N(q) \cap Y^\circ$. Since the Reduction Rule does not trigger on q and s given sets B° and Y° , there exists $t \in (N(s) \cap Y^\circ) \setminus N(q)$; clearly also $t \in C_q$. Furthermore, p is adjacent to neither s nor t , as $s, t \in Y^\circ$ and $p \in N_B(y)$.

Consequently, the vertices p, a, b, q, s, t , together with a path between p and q promised by the fact that C_p and C_q are linked, induce a path on at least six vertices, a contradiction. \square

We infer that in both subcases at least a constant fraction of the set B has been reduced. Consequently, in the branching tree, every leaf-to-root path contains only $O(\log n)$ nodes with a branching described in this subsection.

6.1.6 Final Branch. We are left with cases (X_0, Y) when $Y_1^\circ = \emptyset$. Consider the following natural branch: we guess whether there exists an element of the solution in Y_1^* or not. That is, in one branch we remove $Y_1^* = Y_1$ from Y . In the second branch, we immediately branch again into $|Y_1^*|$ directions, picking a vertex $y \in Y_1^*$ and putting it into X_0 . By the definition of Y_1^* , in the latter subcases B is reduced to an empty set, and branching terminates. Our main claim is that in the first branch, the size of B shrinks by at least a half.

LEMMA 6.8. *In the first branch, if B' and Y' are the sets B and Y after exhaustive application of the Reduction Rule, then $|B'| \leq |B|/2$.*

PROOF. Assume the contrary. Since $Y' \subseteq Y_2$, we have that for every $y \in Y'$ it holds that

$$|N_{B'}(y)| \leq |N_B(y)| < |B|/16 < |B'|/8.$$

For every $p \in B'$, pick two components C_p^1 and C_p^2 that contain a vertex of $N(p) \cap Y'$. Furthermore, for every $i = 1, 2$, pick a vertex $v_p^i \in N(p) \cap Y' \cap C_p^i$ and a vertex $w_p^i \in (N(v_p^i) \cap Y') \setminus N(p)$; the existence of the latter is guaranteed by the fact that the Reduction Rule does not trigger on p and v_p^i , given the sets B' and Y' . Clearly, p, v_p^i, w_p^i induce a P_3 in G .

Consider the following random experiment: choose two vertices $p, q \in B'$ uniformly independently at random. Since the choice of p and q is independent, while all vertices $v_p^i, w_p^i, v_q^i, w_q^i$ belong to Y' , the probability that q is adjacent to v_p^i is less than $1/8$. Consequently, with positive probability q is fully anti-adjacent to $\{v_p^1, w_p^1, v_p^2, w_p^2\}$, while p is fully anti-adjacent to $\{v_q^1, w_q^1, v_q^2, w_q^2\}$.

Let p, q be a pair for which the aforementioned event happens. By potentially swapping the top indices, we may assume $C_p^1 \neq C_q^1$. However, then $p, v_p^1, w_p^1, q, v_q^1, w_q^1$, together with an induced path between p and q whose existence is promised by the fact that C_p^1 and C_q^1 are linked, gives an induced path in G on at least six vertices, a contradiction. \square

By Lemma 6.8, on every leaf-to-root path a branching node described in this subsection may appear only $O(\log |V(G)|)$ times. This finishes the description of the algorithm, and concludes the proof of Theorem 6.1.

6.2 The Actual Algorithm

As described in the beginning of the section, the actual algorithm for MAXIMUM WEIGHT EFFICIENT DOMINATING SET is a standard dynamic programming algorithm, using Theorem 6.1 as the source of its state space. For the sake of the analysis, we fix X_0 to be a maximum weight efficient dominating set in G (if such a set exists).

We first pick any minimal triangulation \widehat{G} of G (see [37] for algorithms finding such a triangulation), and compute its clique tree, which is at the same time a tree decomposition of G . In other words, we compute a tree decomposition (T, Π) of G , where for every node $t \in V(T)$ the bag $\Pi(t)$ is a potential maximal clique of G .

We root T at an arbitrary vertex r , and for a node t we denote by $\Gamma(t)$ the union of all bags $\Pi(s)$, where s ranges over all descendants of t in the tree T . Note that the properties of a tree decomposition ensure that every connected component of $G - \Pi(t)$ is either completely contained in or completely disjoint from $\Gamma(t)$.

For every $t \in V(T)$, we invoke Theorem 6.1 on $\Omega = \Pi(t)$, obtaining a family \mathcal{S}_t .

For a node t , a set $Y \subseteq \Gamma(t)$ is called a *partial solution* if $N[u] \cap N[v] = \emptyset$ for every distinct $u, v \in X$ and $N[X]$ contains $\Gamma(t) \setminus \Pi(t)$. Clearly, if X is an efficient dominating set in G , then $X \cap \Gamma(t)$ is a partial solution. A partial solution Y is *consistent* with a state $f \in \mathcal{S}_t$ if $Y \cap \Pi(t) = f^{-1}(\perp)$, every vertex $v \in \Pi(t) \cap N(Y)$ is dominated by an element of Y in $f(v)$, and for every vertex $v \in \Pi(t) \setminus N[Y]$ the component $f(v)$ is disjoint from $\Gamma(t)$.

Our goal is to compute, in bottom-up fashion, for every node $t \in V(T)$ and every state $f \in \mathcal{S}_t$ a partial solution $Y(t, f)$ consistent with f (or $Y(t, f) = \perp$, meaning that no such set has been found), with the following property: if X_0 exists and f is consistent with $X_0 \cap \Gamma(t)$, then $Y(t, f)$ exists and has weight at least the weight of $X_0 \cap \Gamma(t)$. Note that Theorem 6.1 ensures that if X_0 exists then for every node t there exists a state f_0^t consistent with X_0 , and thus also consistent with the partial solution $X_0 \cap \Gamma(t)$. Consequently, if X_0 exists, then $Y(r, f_0^r)$ is a maximum weight efficient dominating set in G .

It remains to describe the computation for fixed values of t and f and to prove the aforementioned property. For every child t' of t , and every $f' \in \mathcal{S}_{t'}$, we say that the set $Y' := Y(t', f')$ is *partially consistent* with f if $Y' \cap \Pi(t) \cap \Pi(t') = f^{-1}(\perp) \cap \Pi(t')$, every vertex $v \in \Pi(t) \cap N(Y')$ is dominated by an element of Y' in $f(v)$, and for every vertex $v \in \Pi(t) \setminus N[Y']$ the component $f(v)$ is disjoint from $\Gamma(t')$ or equals Ω . For every child t' of t , we compute a maximum weight set $Y_{t'}$ among all sets $Y(t', f')$ for $f' \in \mathcal{S}_{t'}$ that are partially consistent with f ; we terminate the computation and set $Y(t, f) = \perp$ if for some child t' the set $Y_{t'}$ does not exist (i.e., we picked the maximum over an empty set). A direct check shows that if all sets $Y_{t'}$ have been computed, then the union of all sets $Y_{t'}$ is a partial solution consistent with f , and we pick it as $Y(t, f)$.

Consider now the state f_0^t , and assume that for every child t' of t , the set $Y(t', f_0^{t'})$ exists and has weight at least the weight of $X_0 \cap \Gamma(t')$. Observe that $Y(t', f_0^{t'})$ is also partially consistent with f_0^t . A direct check from the definition of consistency shows that for any $Y(t', f')$ partially consistent with f_0^t , the set $X'_0 := (X_0 \setminus \Gamma(t')) \cup Y(t', f')$ is also an efficient dominating set. Since $Y_{t'}$ is chosen to be a set $Y(t', f')$ of maximum weight that is partially consistent with f_0^t , and $Y(t', f_0^{t'})$ is one of the candidates, X'_0 is a maximum-weight efficient dominating set. By repeating this replacement argument for every child t' of t , we infer that the computed value $Y(t, f_0^t)$ has weight at least the weight of $X_0 \cap \Gamma(t)$.

Since the computations are polynomial in the size of G and the sizes of the families \mathcal{S}_t , using Theorem 6.1 we conclude the proof of Theorem 1.2.

7 CONCLUSIONS

We have developed a quasi-polynomial-time algorithm for MAXIMUM WEIGHT INDEPENDENT SET and a polynomial-time algorithm for MAXIMUM WEIGHT EFFICIENT DOMINATING SET in P_6 -free graphs. Our algorithms rely on a detailed analysis of the interactions

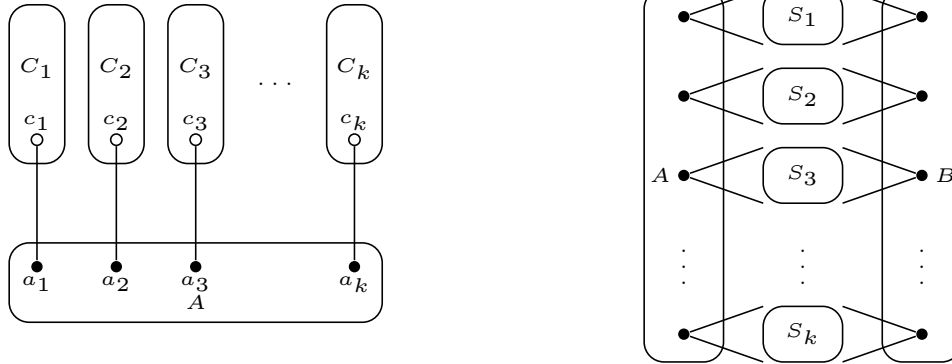


Fig. 14. Counterexamples to generalizations of Theorem 4.2 to P_7 -free graphs (left panel, the nuke vertices are white) and of Theorem 1.3 to P_8 -free graphs (right panel). Every rectangle denotes a clique on k vertices.

between minimal separators, potential maximal cliques, and vertex neighborhoods in P_6 -free graphs. In light of these developments, a few open questions seem natural for the MAXIMUM WEIGHT INDEPENDENT SET problem.

First, can MAXIMUM WEIGHT INDEPENDENT SET on P_6 -free graphs be solved in polynomial time? Very recently, Grzesik et al. [34] answered this question affirmatively.

Second, can our quasipolynomial-time algorithm be generalized to P_7 -free graphs? We remark that polynomial-time algorithms are known for INDEPENDENT SET on various subclasses of P_7 -free graphs [17, 30, 43, 44, 49]. Moreover, Theorems 1.3 and 1.4 work for P_7 -free graphs. However, Theorem 4.2 does not, as can be seen on the following example. Consider the graph G consisting of $k + 1$ pairwise disjoint cliques on k vertices each, denoted A, C_1, C_2, \dots, C_k , with a vertex c_i distinguished in every clique C_i and made adjacent to a private vertex $a_i \in A$ (see the left panel of Fig. 14). G contains many P_6 's with middle two vertices in A , but no P_7 . The set $X = \{c_1, c_2, \dots, c_k\}$ is a nuke in G , but no vertex of G is adjacent to more than one vertex of X . Furthermore, if one adds a new vertex y to G that is adjacent to X , then $X \cup \{y\}$ becomes a potential maximal clique. Recall that the algorithm for MAXIMUM WEIGHT INDEPENDENT SET works by picking a central PMC as a pivot nuke, and branching on vertices adjacent to a constant fraction of the pivot nuke. Hence, with a similar approach on P_7 -free graphs the algorithm may end up with such a seemingly useless nuke as X in G .

Considering P_8 -free graphs and beyond, we note that polynomial-time algorithms are known for INDEPENDENT SET on various subclasses [30, 43]. Recently, Bacsó et al. [3] and Brause [21] independently proved subexponential-time exact algorithms for INDEPENDENT SET on P_k -free graphs for any fixed k . However, our approach breaks down even further: not only does Theorem 4.2 not seem to generalize, as explained previously, but also Theorems 1.3 and 1.4 do not seem to generalize to less restrictive graph classes. We can see this as follows. Consider a graph G consisting of $k + 2$ cliques on k vertices each, denoted $A, B, S_1, S_2, \dots, S_k$, with every S_i adjacent to a private vertex $a_i \in A$ and $b_i \in B$. The set $S := \bigcup_{i=1}^k S_i$ is a minimal separator in G of size k^2 with A and B as full components, yet no vertex of G contains more than k vertices of S in its neighborhood. Furthermore, although G contains many P_7 's with endpoints and middle vertex in S , it does not contain a P_8 nor an E -graph (a P_5 with an additional degree-1 vertex attached to the middle vertex of the path).

Finally, we remark that although the polynomial-time algorithm for MAXIMUM WEIGHT EFFICIENT DOMINATING SET on P_6 -free graphs seems to close a research direction (as the problem is NP-hard on P_7 -free chordal graphs), it would be interesting to see if one can obtain the same end result using the approach of [9, 13], that is, by either obtaining a polynomial-time algorithm for MAXIMUM WEIGHT INDEPENDENT SET in hole-free graphs or showing that the square of a P_6 -free graph having an efficient dominating set is perfect.

ACKNOWLEDGMENTS

The second author acknowledges discussions with Krzysztof Choromański, Dvir Falik, Anita Liebenau, and Viresh Patel on the usage of minimal separators and potential maximal cliques in study of the Erdős-Hajnal conjecture; in particular, the current

proofs of Theorems 1.3 and 1.4, (that replaced our previous proofs more heavy on case analysis) were partially inspired by some quantitative proof attempts in the study of subcases of the Erdős-Hajnal conjecture.

The first author was supported by the BeHard grant under the recruitment programme of the Bergen Research Foundation.

The second author has received funding from the European Research Council under the European Union's Seventh Framework Programme (FP/2007-2013) / ERC Grant Agreement n. 267959 (when M.P. was at the University of Bergen), from the Centre for Discrete Mathematics and its Applications (DIMAP) at the University of Warwick, and from the Warwick-QMUL Alliance in Advances in Discrete Mathematics and its Applications (when M.P. was at the University of Warwick).

The third author was at Max-Planck Institut für Informatik, Saarland Informatics Campus, when most of the work on this paper was done.

REFERENCES

- [1] V. Alekseev. The effect of local constraints on the complexity of determination of the graph independence number. *Combinatorial-algebraic methods in applied mathematics*, pages 3–13, 1982. (in Russian).
- [2] V. E. Alekseev. Polynomial algorithm for finding the largest independent sets in graphs without forks. *Discrete Applied Mathematics*, 135(1–3):3–16, 2004.
- [3] G. Bacsó, D. Marx, and Z. Tuza. H-free graphs, independent sets, and subexponential-time algorithms. In J. Guo and D. Hermelin, editors, *11th International Symposium on Parameterized and Exact Computation, IPEC 2016, August 24–26, 2016, Aarhus, Denmark*, volume 63 of *LIPICs*, pages 3:1–3:12. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2016.
- [4] M. Basavaraju, L. S. Chandran, and T. Karthick. Maximum weight independent sets in hole- and dart-free graphs. *Discrete Applied Mathematics*, 160(16–17):2364–2369, 2012.
- [5] N. Biggs. Perfect codes in graphs. *Journal of Combinatorial Theory, Series B*, 15(3):289–296, 1973.
- [6] H. L. Bodlaender, A. Brandstädt, D. Kratsch, M. Rao, and J. Spinrad. On algorithms for (P_5, gem) -free graphs. *Theoretical Computer Science*, 349(1):2–21, 2005.
- [7] R. Boliac and V. V. Lozin. An augmenting graph approach to the stable set problem in P_5 -free graphs. *Discrete Applied Mathematics*, 131(3):567–575, 2003.
- [8] V. Bouchitté and I. Todinca. Treewidth and minimum fill-in: Grouping the minimal separators. *SIAM J. Comput.*, 31(1):212–232, 2001.
- [9] A. Brandstädt, E. M. Eschen, E. Friese, and T. Karthick. Efficient domination for classes of P_6 -free graphs. *Discrete Applied Mathematics*, 223:15–27, 2017.
- [10] A. Brandstädt and V. Giakoumakis. Weighted efficient domination for $(P_5 + kP_2)$ -free graphs in polynomial time. *CoRR*, abs/1407.4593, 2014.
- [11] A. Brandstädt and V. Giakoumakis. Addendum to: Maximum weight independent sets in hole- and co-chair-free graphs. *Inf. Process. Lett.*, 115(2):345–350, 2015.
- [12] A. Brandstädt and C. T. Hoàng. On clique separators, nearly chordal graphs, and the maximum weight stable set problem. *Theor. Comput. Sci.*, 389(1–2):295–306, 2007.
- [13] A. Brandstädt and T. Karthick. Weighted efficient domination in two subclasses of P_6 -free graphs. *Discrete Applied Mathematics*, 201:38–46, 2016.
- [14] A. Brandstädt, T. Klembt, and S. Mahfud. P_6 - and triangle-free graphs revisited: structure and bounded clique-width. *Discrete Mathematics & Theoretical Computer Science*, 8(1):173–188, 2006.
- [15] A. Brandstädt, M. Milanic, and R. Nevries. New polynomial cases of the weighted efficient domination problem. In *Mathematical Foundations of Computer Science 2013 - 38th International Symposium, MFCS 2013, Klosterneuburg, Austria, August 26–30, 2013. Proceedings*, pages 195–206, 2013.
- [16] A. Brandstädt and R. Mosca. On the structure and stability number of P_5 - and co-chair-free graphs. *Discrete Applied Mathematics*, 132(1–3):47–65, 2003. Stability in Graphs and Related Topics.
- [17] A. Brandstädt and R. Mosca. Maximum weight independent sets for $(P_7, \text{triangle})$ -free graphs in polynomial time. *CoRR*, abs/1511.08066, 2015.
- [18] A. Brandstädt and R. Mosca. Weighted efficient domination for P_6 -free graphs in polynomial time. *CoRR*, abs/1508.07733, 2015.
- [19] A. Brandstädt and R. Mosca. Weighted efficient domination for P_5 -free and P_6 -free graphs. *SIAM J. Discrete Math.*, 30(4):2288–2303, 2016.
- [20] A. Brandstädt and R. Mosca. Weighted efficient domination for P_6 -free and for P_5 -free graphs. In P. Heggernes, editor, *Graph-Theoretic Concepts in Computer Science - 42nd International Workshop, WG 2016, Istanbul, Turkey, June 22–24, 2016, Revised Selected Papers*, volume 9941 of *Lecture Notes in Computer Science*, pages 38–49, 2016.
- [21] C. Brause. A subexponential-time algorithm for the maximum independent set problem in P_7 -free graphs. *Discrete Applied Mathematics*, 231:113–118, 2017.
- [22] H. Broersma, T. Kloks, D. Kratsch, and H. Müller. Independent sets in asteroidal triple-free graphs. *SIAM J. Discrete Math.*, 12(2):276–287, 1999.
- [23] P. Buneman. A characterisation of rigid circuit graphs. *Discrete Mathematics*, 9(3):205–212, 1974.
- [24] D. Corneil, H. Lerchs, and L. S. Burlingham. Complement reducible graphs. *Discrete Applied Mathematics*, 3(3):163–174, 1981.
- [25] H. de Ridder et al. Information system on graph classes and their inclusions (ISGCI), <http://www.graphclasses.org>.
- [26] R. Diestel. *Graph Theory, 4th Edition*, volume 173 of *Graduate texts in mathematics*. Springer, 2012.
- [27] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. Series of Books in the Mathematical Sciences. W. H. Freeman and Co., 1979.
- [28] F. Gavril. Algorithms for minimum coloring, maximum clique, minimum covering by cliques, and maximum independent set of a chordal graph. *SIAM J. Comput.*, 1(2):180–187, 1972.
- [29] F. Gavril. The intersection graphs of subtrees in trees are exactly the chordal graphs. *Journal of Combinatorial Theory, Series B*, 16(1):47–56, 1974.
- [30] M. U. Gerber, A. Hertz, and V. V. Lozin. Stable sets in two subclasses of banner-free graphs. *Discrete Applied Mathematics*, 132(1–3):121–136, 2003.
- [31] M. U. Gerber and V. V. Lozin. On the stable set problem in special P_5 -free graphs. *Discrete Applied Mathematics*, 125(2–3):215–224, 2003.
- [32] M. C. Golumbic. *Algorithmic Graph Theory and Perfect Graphs*. Academic Press, 1980.
- [33] M. Grötschel, L. Lovász, and A. Schrijver. The ellipsoid method and its consequences in combinatorial optimization. *Combinatorica*, 1:169–197, 1981.
- [34] A. Grzesik, T. Klimosova, M. Pilipczuk, and M. Pilipczuk. Polynomial-time algorithm for maximum weight independent set on P_6 -free graphs. *CoRR*, abs/1707.05491, 2017.
- [35] M. Habib and C. Paul. A survey of the algorithmic aspects of modular decomposition. *Computer Science Review*, 4:41–59, 2010.
- [36] T. W. Haynes, S. Hedetniemi, and P. Slater. *Fundamentals of domination in graphs*. CRC Press, 1998.

- [37] P. Heggernes. Minimal triangulations of graphs: A survey. *Discrete Mathematics*, 306(3):297–317, 2006.
- [38] R. M. Karp. Reducibility among combinatorial problems. In *Complexity of Computer Computations*, pages 85–103, 1972.
- [39] T. Karthick. Weighted independent sets in a subclass of P_6 -free graphs. *Discrete Mathematics*, 339(4):1412–1418, 2016.
- [40] T. Karthick and F. Maffray. Weighted independent sets in classes of P_6 -free graphs. *Discrete Applied Mathematics*, 209:217–226, 2016.
- [41] D. Lokshтанov, M. Vatshelle, and Y. Villanger. Independent set in P_5 -free graphs in polynomial time. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2014, Portland, Oregon, USA, January 5-7, 2014*, pages 570–581, 2014.
- [42] V. V. Lozin and M. Milanic. A polynomial algorithm to find an independent set of maximum weight in a fork-free graph. *J. Discrete Algorithms*, 6(4):595–604, 2008.
- [43] V. V. Lozin and D. Rautenbach. Some results on graphs without long induced paths. *Inf. Process. Lett.*, 88(4):167–171, 2003.
- [44] F. Maffray and L. Pastor. Maximum weight stable set in (P_7, bull) -free graphs. *CoRR*, abs/1611.09663, 2016.
- [45] F. Maffray and L. Pastor. The maximum weight stable set problem in (P_6, bull) -free graphs. In P. Heggernes, editor, *Graph-Theoretic Concepts in Computer Science - 42nd International Workshop, WG 2016, Istanbul, Turkey, June 22-24, 2016, Revised Selected Papers*, volume 9941 of *Lecture Notes in Computer Science*, pages 85–96, 2016.
- [46] G. J. Minty. On maximal independent sets of vertices in claw-free graphs. *Journal of Combinatorial Theory, Series B*, 28(3):284–304, 1980.
- [47] R. Mosca. Polynomial algorithms for the maximum stable set problem on particular classes of P_5 -free graphs. *Information Processing Letters*, 61:137–143, 1997.
- [48] R. Mosca. Stable sets in certain P_6 -free graphs. *Discrete Applied Mathematics*, 92(2-3):177–191, 1999.
- [49] R. Mosca. Stable sets of maximum weight in (P_7, banner) -free graphs. *Discrete Mathematics*, 308(1):20–33, 2008.
- [50] R. Mosca. Independent sets in $(P_6, \text{diamond})$ -free graphs. *Discrete Mathematics & Theoretical Computer Science*, 11(1):125–140, 2009.
- [51] R. Mosca. Some results on stable sets for k -colorable P_6 -free graphs and generalizations. *Discrete Mathematics & Theoretical Computer Science*, 14(2):37–56, 2012.
- [52] R. Mosca. Stable sets for $(P_6, K_{2,3})$ -free graphs. *Discussiones Mathematicae Graph Theory*, 32(3):387–401, 2012.
- [53] R. Mosca. Maximum weight independent sets in $(P_6, \text{co-banner})$ -free graphs. *Inf. Process. Lett.*, 113(3):89–93, 2013.
- [54] B. Randerath and I. Schiermeyer. On maximum independent sets in P_5 -free graphs. *Discrete Applied Mathematics*, 158(9):1041–1044, 2010.
- [55] N. Sbihi. Algorithme de recherche d'un stable de cardinalite maximum dans un graphe sans etoile. *Discrete Mathematics*, 29(1):53–76, 1980.
- [56] Y. Shibata. On the tree representation of chordal graphs. *Journal of Graph Theory*, 12(3):421–428, 1988.
- [57] J. R. Walter. *Representations of rigid cycle graphs*. PhD thesis, Wayne State University, 1972.
- [58] C. Yen and R. C. T. Lee. The weighted perfect domination problem and its variants. *Discrete Applied Mathematics*, 66(2):147–160, 1996.
- [59] I. Zverovich and O. Zverovich. Stability number in subclasses of P_5 -free graphs. *Appl. Math. J. Chinese Univ. Ser. B*, 19(2):125–132, 2004.