Subexponential-time Algorithms for Maximum Independent Set in P_t-free and Broom-free Graphs [*]

Gábor Bacsó[†]   Daniel Lokshtanov[‡]   Dániel Marx[§]   Marcin Pilipczuk[¶]

Zsolt Tuza[‖]   Erik Jan van Leeuwen [**]

Abstract

In algorithmic graph theory, a classic open question is to determine the complexity of the MAXIMUM INDEPENDENT SET problem on P_t-free graphs, that is, on graphs not containing any induced path on t vertices. So far, polynomial-time algorithms are known only for $t \leq 5$ [Lokshtanov et al., SODA 2014, 570–581, 2014], and an algorithm for $t = 6$ announced recently [Grzesik et al. Arxiv 1707.05491, 2017]. Here we study the existence of subexponential-time algorithms for the problem: we show that for any $t \geq 1$, there is an algorithm for MAXIMUM INDEPENDENT SET on P_t-free graphs whose running time is subexponential in the number of vertices. Even for the weighted version MWIS, the problem is solvable in $2^{\mathcal{O}(\sqrt{tn \log n})}$ time on P_t-free graphs. For approximation of MIS in broom-free graphs, a similar time bound is proved.

SCATTERED SET is the generalization of MAXIMUM INDEPENDENT SET where the vertices of the solution are required to be at distance at least d from each other. We give a complete characterization of those graphs H for which d-SCATTERED SET on H-free graphs can be solved in time subexponential in the *size of the input* (that is, in the number of vertices plus the number of edges):

- If every component of H is a path, then d-SCATTERED SET on H-free graphs with n vertices and m edges can be solved in time $2^{\mathcal{O}(|V(H)|\sqrt{n+m}\log(n+m))}$, even if d is part of the input.

- Otherwise, assuming the Exponential-Time Hypothesis (ETH), there is no $2^{o(n+m)}$-time algorithm for d-SCATTERED SET for any fixed $d \geq 3$ on H-free graphs with n-vertices and m-edges.

1 Introduction

There are some problems in discrete optimization that can be considered fundamental. The MAXIMUM INDEPENDENT SET problem (MIS, for short) is one of them. It takes a graph G as input, and asks for the maximum number $\alpha(G)$ of mutually nonadjacent (i.e., independent) vertices in G. On unrestricted input, it is not only NP-hard (its decision version "Is $\alpha(G) \geq k$?" being NP-complete), but APX-hard as well, and, in fact, not even approximable within $\mathcal{O}(n^{1-\varepsilon})$ in polynomial time for any $\varepsilon > 0$ unless P=NP, as proved by Zuckerman [27]. For this reason, those classes of graphs on which MIS becomes tractable are of definite

interest. One direction of this area is to study the complexity of MIS on *H-free graphs*, that is, on graphs not containing any *induced* subgraph isomorphic to a given graph H.

For the majority of the graphs H, we know a negative answer on the complexity question. It is easy to see that if G' is obtained from G by subdividing each edge with $2t$ new vertices, then $\alpha(G') = \alpha(G) + t|E(G)|$ holds. This can be used to show that MIS is NP-hard on H-free graphs whenever H is not a forest, and also if H contains a tree component with at least two vertices of degree larger than 2 (first observed in [2], see, e.g., [19]). As MIS is known to be NP-hard on graphs of maximum degree at most 3, the case when H contains a vertex of degree at least 4 is also NP-hard.

The above observations do not cover the case when every component of H is either a path, or a tree with exactly one degree-3 vertex c with three paths of arbitrary lengths starting from c. There are no further unsolved classes but even this collection means infinitely many cases. For decades, on these graphs H only partial results have been obtained, proving polynomial-time solvability in some cases. A classical algorithm of Minty [21] and its corrected form by Sbihi [24] solved the problem when H is a claw (3 paths of length 1 in the model above). This happened in 1980. Much later, in 2004, Alekseev [3] generalized this result by an algorithm for H isomorphic to a fork (2 paths of length 1 and one path of length 2).

The seemingly easy case of P_t-free graphs is poorly understood (where P_t is the path on t vertices). MIS on P_t-free graphs is not known to be NP-hard for any t; for all we know, it could be polynomial-time solvable for every fixed $t \geq 1$. P_4-free graphs (also known as cographs) have a very simple structure, which can be used to solve MIS with a linear-time recursion, but this does not generalize to P_t-free graphs for larger t. In 2010, it was a breakthrough when Randerath and Schiermeyer [22] stated that MIS on P_5-free graphs was solvable in subexponential time, more precisely within $\mathcal{O}(C^{n^{1-\varepsilon}})$ for any constants $C > 1$ and $\varepsilon < 1/4$. Designing an algorithm based on deep results, Lokshtanov et al. [19] finally proved that MIS is polynomial-time solvable on P_5-free graphs. More recently, a *quasipolynomial* ($n^{\log^{\mathcal{O}(1)} n}$-time) algorithm was found for P_6-free graphs [18] and finally a polynomial-time algorithm for P_6-free graphs was announced [13].

We explore MIS and some variants on H-free graphs from the viewpoint of *subexponential-time algorithms* in this work. That is, instead of aiming for algorithms with running time $n^{\mathcal{O}(1)}$ on n-vertex graphs, we ask if $2^{o(n)}$ algorithms are possible. Very recently, Brause [8] and independently the conference version of this paper [4] observed that the subexponential algorithm of Randerath and Schiermeyer [22] can be generalized to arbitrary fixed $t \geq 5$ with running time roughly $2^{\mathcal{O}(n^{1-1/t})}$. Our first result shows a significantly improved subexponential-time algorithm for every t.

Theorem 1.1. *For every fixed $t \geq 5$,* MIS *on n-vertex P_t-free graphs can be solved in subexponential time, namely, it can be solved by a $2^{\mathcal{O}(\sqrt{n \log n})}$-time algorithm.*

The algorithm is based on the combination of two ideas. First, we generalize the observation of Randerath and Schiermeyer [22] stating that in a large connected P_5-free graph there exists a high-degree vertex. Namely, we prove that such a vertex always exists in a large connected P_t-free graph for general $t \geq 5$ and it can be used for efficient branching. Next we prove the combinatorial result that a P_t-free graph of maximum degree Δ has treewidth $\mathcal{O}(t\Delta)$; the proof is inspired by Gyárfás' proof of the χ-boundedness of P_t-free graphs [14]. Thus if the maximum degree drops below a certain threshold during the branching procedure, then we can use standard algorithmic techniques exploiting bounded treewidth.

While our algorithm works for P_t-free graphs with arbitrary large t, it does not seem to be extendable to H-free graphs where H is the subdivision of a $K_{1,3}$. Hence, the existence of subexponential-time algorithms on such graphs remains an open question. However, we are able to give a subexponential-time constant-factor approximation algorithm for the case when H is a (d,t)-broom. A (d,t)-*broom* $B_{d,t}$ is a graph consisting of a path P_t and d additional vertices of degree one, all adjacent to one of the endpoints of the path. In other words, $B_{d,t}$ is a star $K_{1,d+1}$ with one of the edges subdivided to make it a path with t vertices. For $d = 2$, we obtain the *generalized forks* and $t = 3$, $d = 2$ yields the traditional *fork*. We prove the following theorem; here d and t are considered constants, hidden in the big-$\mathcal{O}$ notation.

Theorem 1.2. *Let $d, t \geq 2$ be fixed integers. One can find a d-approximation to* MAXIMUM INDEPENDENT SET *on an n-vertex $B_{d,t}$-free graph G in time $2^{\mathcal{O}(n^{3/4} \log n)}$.*

Let us remark that on $K_{1,d+1}$-free graphs, a folklore linear-time (and very simple) d-approximation algorithm exists for MAXIMUM INDEPENDENT SET; better $d/2$-approximation algorithms also exist [5, 6, 15, 26]. On fork-free graphs, INDEPENDENT SET can be solved in polynomial time [3]. For general graphs, we do not expect that a constant-factor approximation can be obtained in subexponential time for the problem. Strong evidence for this was given by Chalermsook et al. [9], who showed that the existence of such an algorithm would violate the Exponential-Time Hypothesis (ETH) of Impagliazzo, Paturi, and Zane, which can be informally stated as n-variable 3SAT cannot be solved in $2^{o(n)}$ time (see [10, 16, 17]).

SCATTERED SET (also known under other names such as dispersion or distance-d independent set [1, 7, 11, 20, 23, 25]) is the natural generalization of MIS where the vertices of the solution are required to be at distance at least d from each other; the size of the largest such set will be denoted by $\alpha_d(G)$. We can consider with d being part of the input, or assume that $d \geq 2$ is a fixed constant, in which case we call the problem d-SCATTERED SET. Clearly, MIS is exactly the same as 2-SCATTERED SET. Despite its similarity to MIS, the branching algorithm of Theorem 1.1 cannot be generalized: we give evidence that there is no subexponential-time algorithm for 3-SCATTERED SET on P_5-free graphs.

Theorem 1.3. *Assuming the ETH, there is no $2^{o(n)}$-time algorithm for d-SCATTERED SET with $d = 3$ on P_5-free graphs with n vertices.*

In light of the negative result of Theorem 1.3, we slightly change our objective by aiming for an algorithm that is subexponential in the *size of the input,* that is, in the total number of vertices and edges of the graph G. As the number of edges of G can be up to quadratic in the number of vertices, this is a weaker goal: an algorithm that is subexponential in the number of edges is not necessarily subexponential in the number of vertices. We give a complete characterization when such algorithms are possible for SCATTERED SET.

Theorem 1.4. *For every fixed graph H, the following holds.*

1. *If every component of H is a path, then d-SCATTERED SET on H-free graphs with n vertices and m edges can be solved in time $2^{\mathcal{O}(|V(H)|\sqrt{n+m}\log(n+m))}$, even if d is part of the input.*

2. *Otherwise, assuming the ETH, there is no $2^{o(n+m)}$-time algorithm for d-SCATTERED SET for any fixed $d \geq 3$ on H-free graphs with n-vertices and m-edges.*

The algorithmic side of Theorem 1.4 is based on the combinatorial observation that the treewidth of P_t-free graphs is sublinear in the number of edges, which means that standard algorithms on bounded-treewidth graphs can be invoked to solve the problem in time subexponential in the number of edges. It has not escaped our notice that this approach is completely generic and could be used for many other problems (e.g., HAMILTONIAN CYCLE, 3-COLORING, and so on), where $2^{\mathcal{O}(t)} \cdot n^{\mathcal{O}(1)}$ or perhaps $2^{t \cdot \log^{\mathcal{O}(1)} t} \cdot n^{\mathcal{O}(1)}$-time algorithms are known on graphs of treewidth t. For the lower-bound part of Theorem 1.4, we need to examine only two cases: claw-free graphs and C_t-free graphs (where C_t is the cycle on t vertices); the other cases then follow immediately.

The paper is organized as follows. Section 2 introduces basic notation and contains some technical tools for bounding the running time of recursive algorithms. Section 3 contains the combinatorial results that allow us to bound the treewidth of P_t-free graphs. The algorithmic results for MAXIMUM INDEPENDENT SET (Theorems 1.1 and 1.2) appear in Section 4. The upper and lower bounds for d-SCATTERED SET, which together prove Theorem 1.4, are proved in Section 5.

2 Preliminaries

Simple undirected graphs are investigated here throughout. The vertex set of graph G will be denoted by $V(G)$, the edge set by $E(G)$. The notation $d_G(x, y)$ for distance, $G[X]$ for the subgraph induced by the vertex set X, will have the usual meaning, similarly as $N_G[X]$ and $N_G(X)$ for the closed and open neighborhood respectively of vertex set X in G. $\Delta(G)$ is the maximum degree in G. For a vertex set X in G, $G - X$ means the induced subgraph $H := G[V - X]$. P_t (C_t) is the chordless path (cycle) on t vertices. Finally, a graph is H-free if it does not contain H as an induced subgraph.

A *distance-d (d-scattered) set* in a graph G is a vertex set $S \subseteq V(G)$ such that for every pair of vertices in S, the distance between them is at least d in the graph. For $d = 2$, we obtain the traditional notion of independent set (stable set). For $d > c$, a distance-d set is a distance-c set as well, for example, for $d \geq 2$, any distance-d set is an independent set.

The algorithmic problem MAXIMUM WEIGHT INDEPENDENT SET is the problem of maximizing the sum of the weights in an independent set of a graph with nonnegative vertex weights w. The maximum is denoted by $\alpha_w(G)$. For a weight w function that has value 1 everywhere, we obtain the usual problem MAXIMUM INDEPENDENT SET (MIS) with maximum $\alpha(G)$.

An algorithm A is *subexponential* in parameter $p > 1$ if the number of steps executed by A is a subexponential function of the parameter p. We will use here this notion for graphs, mostly in the following cases: p is the number n of vertices, the number m of edges, or $p = n + m$ (which is considered to be the size of the input generally). Several different definitions are used in the literature under the name *subexponential function*. Each of them means some condition: this function (with variable $p > 1$, called the parameter) may not be larger than some bound, depending on p. Here we use two versions, where the bound is of type $exp(o(p))$ and $exp(p^{1-\epsilon})$ respectively, with some $\epsilon > 0$. (Clearly, the second one is the more strict.) Throughout the paper, we state our results emphasizing which version we mean. A problem Π is *subexponential* if there exists some *subexponential* algorithm solving Π.

2.1 Time analysis of recursive algorithms

To formally reason about time complexities, we will need the following technical lemma.

Lemma 2.1. *Let $\Delta : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ be a concave and nondecreasing function with $\Delta(0) = 0$, $\Delta(x) \leq x$ for every $x \geq 1$, and $\Delta(x) \leq \Delta(x/2) \cdot (2 - \gamma)$ for some $\gamma > 0$ and every $x \geq 2$. Let $S, T : \mathbb{N} \to \mathbb{N}$ be two nondecreasing functions such that we have $S(0) = T(0) = 0$, moreover, for some universal constant c and $S(1), T(1) \leq c$ and for every $n \geq 2$:*

$$T(n) \leq 2^{cn \log n / \Delta(n)} + \max(S(n), T(n-1) + T(n - \lceil \Delta(n) \rceil),$$
$$\max_{1 \leq k \leq \lfloor \frac{n}{\Delta(n)} \rfloor} 2^k \cdot n \cdot T(n - \lceil k\Delta(n) \rceil)). \tag{1}$$

Then, for some constant c' depending only on c and γ, for every $n \geq 1$ it holds that

$$T(n) \leq 2^{c'n \log n / \Delta(n)} \cdot (S(n) + 1).$$

We will use Lemma 2.1 as a shortcut to argue about time complexities of our branching algorithms; let us now briefly explain its intuition. The function $T(n)$ will be the running time bound of the discussed algorithm. The term $2^{cn \log n / \Delta(n)}$ in (1) corresponds to a processing time at a single step of the algorithm; note that this is at least polynomial in n as $\Delta(n) \leq n$. The terms in the max in (1) are different branching options chosen by the algorithm. The first one, $S(n)$, is a subcall to a different procedure, such as bounded treewidth subroutine. The second one, $T(n) + T(n - \lceil \Delta(n) \rceil)$, corresponds to a two-way branching on a single vertex of degree at least $\Delta(n)$. The last one corresponds to an exhaustive branching on a set $X \subseteq V(G)$ of size k, such that every connected component of $G - X$ has at most $n - k\Delta(n)$ vertices.

Proof of Lemma 2.1. For notational convenience, it will be easier to assume that the functions S and T is defined on the whole half-line $\mathbb{R}_{\geq 0}$ with $S(x) = S(\lfloor x \rfloor)$ and $T(x) = T(\lfloor x \rfloor)$.

First, let us replace max with addition in the assumed inequality. After some simplifications, this leads to the following.

$$T(n) \leq T(n-1) + S(n) + 2^{cn \log n / \Delta(n)} + 2n \cdot \sum_{k=1}^{\lfloor \frac{n}{\Delta(n)} \rfloor} 2^k \cdot T(n - k\Delta(n)). \tag{2}$$

From the concavity of $\Delta(n)$ it follows that

$$n - i - \Delta(n - i) \leq n - \Delta(n).$$

4

Furthermore, the assumptions on Δ, namely the fact that Δ is nondecreasing, concave, with $\Delta(0) = 0$, implies that for any $0 < y < x$ we have

$$\frac{y}{x}\Delta(x) \geq \Delta(x) - \Delta(x - y).$$

After simple algebraic manipulation, this is equivalent to

$$\frac{x}{\Delta(x)} \geq \frac{x - y}{\Delta(x - y)}.$$

That is, $x \mapsto x/\Delta(x)$ is a nondecreasing function.

Using the fact that $S(n)$ and $T(n)$ are nondecreasing and the facts above, we iteratively apply (2) n times to the first summand, obtaining the following.

$$T(n) \leq n \cdot \left(S(n) + 2^{cn \log n/\Delta(n)} + 2n \cdot \sum_{k=1}^{\lfloor \frac{n}{\Delta(n)} \rfloor} 2^k \cdot T(n - k\Delta(n)) \right). \tag{3}$$

We now show the following.

Claim 2.2. *Consider a sequence $n_0 = n$ and $n_{i+1} = n_i - \Delta(n_i)$. Then $n_i = \mathcal{O}(1)$ for $i = \mathcal{O}(n/\Delta(n))$. Here, the big-$\mathcal{O}$-notation hides constants depending on γ.*

Proof. By the concavity of Δ we have $\Delta(n'/2) \geq \Delta(n')/2$, thus as long as $n_i > n_0/2$ we have that $n_{i+1} \leq n_i - \Delta(n)/2$. Consequently, for some $j = \mathcal{O}(n/\Delta(n))$ we have $n_j < n_0/2$. We infer that we obtain $n_i = \mathcal{O}(1)$ at position

$$i = \mathcal{O}\left(\frac{n}{\Delta(n)} + \frac{n/2}{\Delta(n/2)} + \frac{n/4}{\Delta(n/4)} + \dots \right).$$

By the assumption that $\Delta(x) \leq \Delta(x/2) \cdot (2 - \gamma)$ for some constant $\gamma > 0$ and every $x \geq 2$, the sum above can be bounded by a geometric sequence, yielding $i = \mathcal{O}(n/\Delta(n))$. ⌟

The above claim implies that if we iteratively apply (3) to itself, we obtain

$$T(n) \leq (2n)^{\mathcal{O}(n/\Delta(n))} \cdot \left(S(n) + 2^{cn \log n/\Delta(n)} \right).$$

This finishes the proof of the lemma. □

3 Gyárfás' path-growing argument

The main (technical but useful) result of this section is the following adaptation of Gyárfás' proof that P_t-free graphs are χ-bounded [14].

Lemma 3.1. *Let $t \geq 2$ be an integer, G be a connected graph with a distinguished vertex $v_0 \in V(G)$ and maximum degree at most Δ, such that G does not contain an induced path P_t with one endpoint in v_0. Then, for every weight function $w : V(G) \to \mathbb{Z}_{\geq 0}$, there exists a set $X \subseteq V(G)$ of size at most $(t - 1)\Delta + 1$ such that every connected component C of $G - X$ satisfies $w(C) \leq w(V(G))/2$. Furthermore, such a set X can be found in polynomial time.*

Proof. In what follows, a connected component C of an induced subgraph H of G is *big* if $w(C) > w(V(G))/2$. Note that there can be at most one big connected component in any induced subgraph of G.

If $G - \{v_0\}$ does not contain a big component, we can set $X = \{v_0\}$. Otherwise, let $A_0 = \{v_0\}$ and B_0 be the big component of $G - A_0$. As G is connected, every component of $G - A_0$ is adjacent to A_0, thus $v_0 \in N(B_0)$ holds. We will inductively define vertices $v_1, v_2, v_3, \dots$ such that $v_0, v_1, v_2, \dots$ induce a path in G.

5

Given vertices $v_0, v_1, v_2, \ldots, v_i$, we define sets A_{i+1} and B_{i+1} as follows. We set $A_{i+1} = N_G[v_0, v_1, \ldots, v_i]$. If $G - A_{i+1}$ does not contain a big connected component, we stop the construction. Otherwise, we set B_{i+1} to be the big connected component of $G - A_{i+1}$. During the process we maintain the invariant that B_i is the big component of $G - A_i$ and that $v_i \in N(B_i)$. Note that this is true for $i = 0$ by the choice of A_0 and B_0.

It remains to show how to choose v_{i+1}, given vertices $v_0, v_1, \ldots, v_i$ and sets A_{i+1} and B_{i+1}. Note that $A_{i+1} = A_i \cup N_G[v_i]$ and $v_i \in N(B_i)$, so B_{i+1} is the big connected component of $G[(B_i \setminus N_G(v_i))]$. Consequently, we can choose some $v_{i+1} \in B_i \cap N_G(B_{i+1}) \cap N_G(v_i)$ that satisfies all the desired properties.

Since G does not contain an induced P_t with one endpoint in v_0, the aforementioned process stops after defining a set A_{i+1} for some $i < t - 1$, when $G - A_{i+1}$ does not contain a big component. Observe that

$$|A_{i+1}| \leq (\Delta + 1) + i \cdot \Delta = (i+1)\Delta + 1 \leq (t-1)\Delta + 1.$$

Consequently, the set $X := A_{i+1}$ satisfies the desired properties.

For the algorithmic claim, note that the entire proof can be made algorithmic in a straightforward manner. $\qquad\square$

It is well known that if graph G has a set X of size k for every weight function $w : V(G) \to \mathbb{Z}_{\geq 0}$ such that every connected component C of $G - X$ satisfies $w(C) \leq w(V(G))/2$, then G has treewidth $\mathcal{O}(w)$ (see, e.g., [12, Theorem 11.17(2)]). Thus Lemma 3.1 implies a treewidth bound of $\mathcal{O}(t\Delta)$. Algorithmically, it is also a standard consequence of Lemma 3.1 that a tree decomposition of width $\mathcal{O}(t\Delta)$ can be obtained in polynomial time. What needs to be observed is that standard 4-approximation algorithms for treewidth, which run in time exponential in treewidth, can be made to run in polynomial time if we are given a polynomial-time subroutine for finding the separator X as in Lemma 3.1. For completeness, we sketch the proof here.

Corollary 3.2. *A P_t-free graph with maximum degree Δ has treewidth $\mathcal{O}(t\Delta)$. Furthermore, a tree decomposition of this width can be computed in polynomial time.*

Proof. We follow standard constant approximation algorithm for treewidth, as described in [10, Section 7.6]. This algorithm, given a graph G and an integer k, either correctly concludes that $\mathrm{tw}(G) > k$ or computes a tree decomposition of G of width at most $4k + 4$.

Let G be a P_t-free graph with maximum degree at most Δ. We may assume that G is connected, otherwise we can handle the connected components separately. Let us start by setting $k := (t-1)\Delta$ so that any application of Lemma 3.1 gives a set of size at most $k + 1$.

The only step of the algorithm that runs in exponential time is the following. We are given an induced subgraph $G[W]$ of G and a set $S \subseteq W$ with the following properties:

1. $|S| \leq 3k + 4$ and $W \setminus S \neq \emptyset$;

2. both $G[W]$ and $G[W \setminus S]$ are connected;

3. $S = N_G(W \setminus S)$.

The goal is to compute a set $S \subsetneq \widehat{S} \subseteq W$ such that $|\widehat{S}| \leq 4k + 5$ and every connected component of $G[W \setminus \widehat{S}]$ is adjacent to at most $3k + 4$ vertices of $\widehat{S}$.

The construction of $\widehat{S}$ is trivial for $|S| < 3k+4$, as we can take $\widehat{S} = S \cup \{v\}$ for an arbitrary $v \in W \setminus S$. The crucial step happens for sets S of size exactly $3k+4$. Instead of the exponential search of [10, Section 7.6], we invoke Lemma 3.1 on the graph $G[W]$ and a function $w : W \to \{0, 1\}$ that puts $w(v) = 1$ if and only if $v \in S$. The lemma returns a set $X \subseteq W$ of size at most $k + 1$ such that every connected component C of $G[W \setminus X]$ contains at most $3k/2 + 2$ vertices of S. Since $G[W \setminus S]$ is connected and $(3k/2 + 2) + (k + 1) < 3k + 4$, we cannot have $X \subseteq S$. Consequently, $\widehat{S} := S \cup X$ satisfies all the requirements.

The algorithm of [10, Section 7.6] returns that $\mathrm{tw}(G) > k$ only if at some step it encounters pair (W, S) for which it cannot construct the set $\widehat{S}$. However, our method of constructing $\widehat{S}$ works for every choice of (W, S), and executes in polynomial time. Consequently, the modified algorithm of [10, Section 7.6] always computes a tree decomposition of width at most $4k + 4 = \mathcal{O}(t\Delta)$ in polynomial time, as desired. $\qquad\square$

4 Subexponential algorithms based on the path-growing argument

The goal of this section is to use Corollary 2.2 to prove Theorems 1.1 and 1.2 stated in the Introduction.

4.1 INDEPENDENT SET on graphs without long paths

We first prove the following statement, which implies Theorem 1.1.

Theorem 4.1. *The* MAXIMUM-WEIGHT INDEPENDENT SET *problem on an n-vertex P_t-free graph can be solved in time $2^{\mathcal{O}(\sqrt{tn \log n})}$.*

Proof. Let G be an n-vertex P_t-free graph. We set a threshold $\Delta = \Delta(n) := \sqrt{\frac{n \log(n+1)}{t}}$. If the maximum degree of G is at most Δ, we invoke Corollary 3.2 to obtain a tree decomposition of G of width $\mathcal{O}(t\Delta) = \mathcal{O}(\sqrt{tn \log n})$. By standard techniques on graphs of bounded treewidth (cf. [10]), we solve MAXIMUM-WEIGHT INDEPENDENT SET on G in time $2^{\mathcal{O}(\sqrt{tn \log n})}$.

Otherwise, G contains a vertex of degree greater than Δ. We choose (arbitrarily) such a vertex v and we branch on v: either v is contained in the maximum independent set or not. In the first case we delete $N_G[v]$ from G, in the second we delete only v from G. This gives the following recursion for the time complexity $T(n)$ of the algorithm.

$$T(n) \le \max\left(T(n-1) + T(n - \lceil \Delta(n) \rceil) + \mathcal{O}(n^2), 2^{\mathcal{O}(\sqrt{tn \log n})}\right). \tag{4}$$

Observe that we have $T(n) = 2^{\mathcal{O}(\sqrt{tn \log n})}$ by Lemma 2.1 with $S(n) = 2^{\mathcal{O}(\sqrt{tn \log n})}$; it is straightforward to check that $\Delta(n) = \sqrt{\frac{n \log(n+1)}{t}}$ satisfies all the prerequisites of Lemma 2.1. This finishes the proof of the theorem. $\square$

4.2 Approximation on broom-free graphs

We now extend the argumentation in Theorem 4.1 to (d, t)-*brooms*—however, this time we are able to obtain only an approximation algorithm. Recall that a (d, t)-broom $B_{d,t}$ is a graph consisting of a path P_t and d additional vertices of degree one, all adjacent to one of the endpoints of the path.

We now prove Theorem 1.2 from the introduction.

Proof of Theorem 1.2. Let $\Delta(n) = \frac{1}{2dt} \cdot n^{1/4}$; note that such a definition fits the prerequisites of $\Delta(n)$ for Lemma 2.1. In the complexity analysis, we will use Lemma 2.1 with this $\Delta(n)$ and without any function $S(n)$; this will give the promised running time bound. In what follows, whenever we execute a branching step of the algorithm we argue that it fits into one of the subcases of the max in (1) of Lemma 2.1.

As in the proof of Theorem 4.1, as long as there exists a vertex in G of degree larger than Δ, we can branch on such a vertex v: in one subcase, we consider independent sets not containing v (and thus delete v from G), in the other subcase, we consider independent sets containing v (and thus delete $N(v)$ from G). Such a branching step can be conducted in polynomial time, and fits in the second subcase of max in (1). Thus, we can assume henceforth that the maximum degree of G is at most Δ.

We also assume that G is connected and $n > (2dt)^4$, as otherwise we can consider every connected component independently and/or solve the problem by brute-force.

Later, we will also need a more general branching step. If, in the course of the analysis, we identify a set $X \subseteq V(G)$ such that every connected component of $G - X$ has size at most $n - \frac{|X| n^{1/4}}{2dt}$, then we can exhaustively branch on all vertices of X and independently resolve all connected components of the remaining graph. Such a branching fits into the last case of the max in (1), and hence it again leads to the desired time bound $2^{\mathcal{O}(n^{3/4} \log n)}$ by Lemma 2.1.

We start with greedily constructing a set A_0 with the following properties: $G[A_0]$ is connected and $n^{1/2} \le |N[A_0]| \le n^{1/2} + \Delta$. We start with A_0 being a single arbitrary vertex and, as long as $|N[A_0]| < n^{1/2}$,

we add an arbitrary vertex of $N(A_0)$ to A_0 and continue. Since G is connected, the process ends when $|N[A_0]| \geq n^{1/2}$; since the maximum degree of G is at most Δ, we have $|N[A_0]| \leq n^{1/2} + \Delta < 2n^{1/2}$.

Let B be the vertex set of the largest connected component of $G - N[A_0]$. If $|B| < n - n^{3/4}$, we exhaustively branch on $X := N[A_0]$, as X is of size at most $2n^{1/2}$, but every connected component of $G - X$ is of size at most $n - n^{3/4} \leq n - \frac{1}{2}|X|n^{1/4}$. Hence, we are left with the case $|B| > n - n^{3/4}$.

Let $S = N(B)$. Note that A_0 is disjoint from $N[B]$. Let A_1 be the connected component of $G - S$ that contains A_0. Since $S \subseteq N(A_0)$, we have that $N[A_1] \supseteq N[A_0]$; in particular, $|N[A_1]| \geq n^{1/2}$ while, as $|B| > n - n^{3/4}$, we have $|N[A_1]| \leq n^{3/4}$. Furthermore, since $S \subseteq N(A_0)$ and $A_0 \subseteq A_1$, we have $N(A_1) = S$.

Consider now the following case: there exists $v \in S$ such that $N(v) \cap B$ contains an independent set L of size d. Observe that such a vertex v can be found by an exhaustive search in time $n^{d + \mathcal{O}(1)}$.

For such a vertex v and independent set L, define D to be the vertex set of the connected component of $G - (N[L] \setminus \{v\})$ that contains A_1. Note that as $L \subseteq B$ we have $N[L] \cap A_1 = \emptyset$, and thus such a component D exists. Furthermore, as $N(A_1) = S$, D contains $S \setminus (N(L) \setminus \{v\})$. In particular, D contains v, and

$$|D| \geq |(A_1 \cup S) \setminus N(L)| \geq |N[A_1]| - \Delta \cdot |L| \geq n^{1/2} - dn^{1/4} \geq \frac{1}{2}n^{1/2}.$$

If $|D| < n - n^{1/2}$, then we exhaustively branch on the set $X := N[L] \setminus \{v\}$, as $|X| \leq d\Delta \leq \frac{1}{2}n^{1/4}$ while every connected component of $G - X$ is of size at most $n - \frac{1}{2}n^{1/2}$ due to D being of size at least $\frac{1}{2}n^{1/2}$ and at most $n - n^{1/2}$. Consequently we can assume $|D| \geq n - n^{1/2}$.

Observe that $G[D]$ does not contain a path P_t with one endpoint in v, as such a path, together with the set L, would induce a $B_{d,t}$ in G. Consequently, we can apply Lemma 3.1 to the graph $G[D]$ with the vertex $v_0 = v$ and uniform weight $w(u) = 1$ for every $u \in D$, obtaining a set $X_D \subseteq D$ of size $|X_D| \leq (t-1)\Delta + 1 \leq \frac{1}{2}n^{1/4}$ such that every connected component of $G[D \setminus X]$ has size at most $n/2$. We branch exhaustively on the set $X = X_D \cup (N[L] \setminus \{v\})$: this set is of size at most $n^{1/4}$, while every connected component of $G - X$ is of size at most $n/2$ due to the properties of X_D and the fact that $|D| \geq n - n^{1/2}$. This finishes the description of the algorithm in the case when there exists $v \in S$ and an independent set $L \subseteq N(v) \cap B$ of size d.

We are left with the complementary case, where for every $v \in S$, the maximum independent set in $N(v) \cap B$ is of size less than d. We perform the following operation: by exhaustive search, we find a maximum independent set I_A in $G - B$ and greedily take it to the solution; that is, recurse on $G - N[I_A]$ and return the union of I_A and the independent set found by the recursive call in $G - N[I_A]$. Since $|B| > n - n^{3/4}$, the exhaustive search runs in $2^{n^{3/4}} n^{\mathcal{O}(1)}$ time, fitting the first summand of the right hand side in (1). As a result, the graph reduces by at least one vertex, and hence the remaining running time of the algorithm fits into the second case of the max in (1). This gives the promised running time bound. It remains to argue about the approximation ratio; to this end, it suffices to show the following claim.

Claim 4.2. *If I is a maximum independent set in G and I' is a maximum independent set in $G - N[I_A]$, then $|I| - |I'| \leq d|I_A|$.*

Proof. Let $J = I \setminus N[I_A]$. Clearly, J is an independent set in $G - N[I_A]$, and thus $|J| \leq |I'|$. It suffices to show that $|I| - |J| \leq d|I_A|$, that is, $|I \cap N[I_A]| \leq d|I_A|$.

The maximality of I_A implies that $V(G) \setminus B \subseteq N[I_A]$. As I_A is a maximum independent set in $G - B$, we have that $|I \setminus B| \leq |I_A|$. For every $w \in I \cap N[I_A] \cap B$, pick a neighbor $f(w) \in I_A \cap N(w)$. Note that we have $f(w) \in S$. Since for every vertex $v \in S$, the size of the maximum independent set in $N(v) \cap B$ is less than d, we have $|f^{-1}(v)| < d$ for every $v \in S \cap I$. Consequently,

$$|I \cap N[I_A] \cap B| \leq (d-1)|I_A \cap S| \leq (d-1)|I_A|.$$

Together with $|I \setminus B| \leq |I_A|$, we have $|I \cap N[I_A]| \leq d|I_A|$, as desired. ⌟

This finishes the proof of Theorem 1.2. □

5 Scattered Set

We prove Theorem 1.4 in this section. The algorithm for SCATTERED SET for P_t-free graphs hinges on the following combinatorial bound.

Lemma 5.1. *For every $t \geq 2$ and for every P_t-free graph with m edges, we have that G has treewidth $\mathcal{O}(t\sqrt{m})$.*

Proof. Let X be the set of vertices of G with degree at least $\sqrt{m}$. The sum of the degrees of the vertices in X is at most $2m$, hence we have $|X| \leq 2m/\sqrt{m} = 2\sqrt{m}$. By the definition of X, the graph $G - X$ has maximum degree less than $\sqrt{m}$. Thus by Corollary 3.2, the treewidth of $G - X$ is $\mathcal{O}(t\sqrt{m})$. As removing a vertex can decrease treewidth at most by one, it follows that G has treewidth at most $\mathcal{O}(t\sqrt{m}) + |X| = \mathcal{O}(t\sqrt{m})$. $\square$

It is known that SCATTERED SET can be solved in time $d^{\mathcal{O}(w)} \cdot n^{\mathcal{O}(1)}$ on graphs of treewidth w using standard dynamic programming techniques (cf. [20, 25]). By Lemma 5.1, it follows that SCATTERED SET on P_t-free graphs can be solved in time $d^{\mathcal{O}(t\sqrt{m})} \cdot n^{\mathcal{O}(1)}$. If d is a fixed constant, then this running time can be bounded as $2^{\mathcal{O}(t\sqrt{m}) + \mathcal{O}(\log n)} = 2^{\mathcal{O}(t\sqrt{n+m})}$. If d is part of the input, then (taking into account that we may assume $d \leq n$) the running time is

$$d^{\mathcal{O}(t\sqrt{m})} \cdot n^{\mathcal{O}(1)} = 2^{\mathcal{O}(t\sqrt{m}\log n) + \mathcal{O}(\log n)} = 2^{\mathcal{O}(t\sqrt{n+m}\log(n+m))}.$$

Observe that if every component of a fixed graph H is a path, then H is an induced subgraph of $P_{2|V(H)|}$, which implies that H-free graphs are $P_{2|V(H)|}$-free. Thus the algorithm described here for P_t-free graphs implies the first part of Theorem 1.4.

5.1 Lower bounds for SCATTERED SET

A standard consequence of the ETH and the so-called Sparsification Lemma is that there is no subexponential-time algorithm for MIS even on graphs of bounded degree (see, e.g., [10]):

Theorem 5.2. *Assuming the ETH, there is no $2^{o(n)}$-time algorithm for MIS on n-vertex graphs of maximum degree 3.*

A very simple reduction can reduce MIS to 3-SCATTERED SET for P_5-free graphs, showing that, assuming the ETH, there is no algorithm subexponential in the number of vertices for the latter problem. This proves Theorem 1.3 stated in the Introduction.

Proof of Theorem 1.3. Given an n-vertex m-edge graph G with maximum degree 3 and an integer k, we construct a P_5-free graph G' with $n + m = \mathcal{O}(n)$ vertices such that $\alpha(G) = \alpha_3(G')$. This reduction proves that a $2^{o(n)}$-time algorithm for 3-SCATTERED SET could be used to obtain a $2^{o(n)}$-time algorithm for MIS on graphs of maximum degree 3, and this would violate the ETH by Theorem 5.2.

We may assume that G has no isolated vertices. The graph G' contains one vertex for each vertex of G and additionally one vertex for each edge of G. The m vertices of G' representing the edges of G form a clique. Moreover, if the endpoints of an edge $e \in E(G)$ are $u, v \in V(G)$, then the vertex of G' representing e is connected with the vertices of G' representing u and v. This completes the construction of G'. It is easy to see that G' is P_5-free: an induced path of G' can contain at most two vertices of the clique corresponding to $E(G)$ and the vertices of G' corresponding to the vertices of G form an independent set.

If S is an independent set of G, then we claim that the corresponding vertices of G' are at distance at least 3 from each other. Indeed, no two such vertices have a common neighbor: if $u, v \in S$ and the corresponding two vertices in G' have a common neighbor, then this common neighbor represents an edge e of G whose endpoints are u and v, violating the assumption that S is independent. Conversely, suppose that $S' \subseteq V(G')$ is a set of k vertices with pairwise distance at least 3 in G'. If $k \geq 2$, then all these vertices represent vertices of G: observe that for every edge e of G, the vertex of G' representing e is at distance at most 2 from every other (non-isolated) vertex of G'. We claim that S' corresponds to an independent set of G. Indeed, if $u, v \in S'$ and there is an edge e in G' with endpoints u and v, then the vertex of G' representing e is a common neighbor of u and v, a contradiction. $\square$

Next we give negative results on the existence of algorithms for SCATTERED SET that have running time subexponential in the number of edges. To rule out such algorithms, we construct instances that have bounded degree: then being subexponential in the number of vertices or the number of edges are the same. We consider first claw-free graphs. The key insight here is that SCATTERED SET with $d = 3$ in line graphs (which are claw-free) is essentially the INDUCED MATCHING problem, for which it is easy to prove hardness results.

Theorem 5.3. *Assuming the ETH, d-SCATTERED SET does not have a $2^{o(n)}$ algorithm on n-vertex claw-free graphs of maximum degree 6 for any fixed $d \geq 3$.*

Proof. Given an n-vertex graph G with maximum degree 3, we construct a claw-free graph G' with $\mathcal{O}(dn)$ vertices and maximum degree 4 such that $\alpha_d(G') = \alpha(G)$. Then by Theorem 5.2, a $2^{o(n)}$-time algorithm for d-SCATTERED SET for n-vertex claw-free graphs of maximum degree 4 would violate the ETH.

The construction is slightly different based on the parity of d; let us first consider the case when d is odd. Let us construct the graph G^+ by attaching a path Q_v of $\ell = (d-1)/2$ edges to each vertex $v \in V(G)$; let us denote by $e_{v,1}, \ldots, e_{v,\ell}$ the edges of this path such that $e_{v,1}$ is incident with v. The graph G' is defined as the line graph of G^+, that is, each vertex of G' represents an edge of G^+ and two vertices of G' are adjacent if the corresponding two vertices share an endpoint. It is well known that line graphs are claw-free. As G^+ has $\mathcal{O}(dn)$ edges and maximum degree 4 (recall that G has maximum degree 3), the line graph G' has maximum degree 6 with $\mathcal{O}(dn)$ vertices an edges. Thus an algorithm for SCATTERED SET with running time $2^{o(n)}$ on n-vertex claw-free graphs of maximum degree 3 could be used to solve MIS on n-vertex graphs with maximum degree 3 in time $2^{o(n)}$, contradicting the ETH.

If there is an independent set S of size k in G, then we claim that the set $S' = \{e_{v,\ell} \mid v \in S\}$ is a d-scattered set of size k in G'. To see this, suppose for a contradiction that there are two vertices $u, v \in S$ such that the vertices of G' representing $e_{u,\ell}$ and $e_{v,\ell}$ are at distance at most $d - 1$ from each other. This implies that there is a path in G^+ that has at most d edges and whose first and last edges are $e_{u,\ell}$ and $e_{v,\ell}$, respectively. However, such a path would need to contain all the ℓ edges of path Q_u and all the ℓ edges of Q_v, hence it can contain at most $d - 2\ell = 1$ edges outside these two paths. But u and v are not adjacent in G^+ by assumption, hence more than one edge is needed to complete Q_u and Q_v to a path, a contradiction.

Conversely, let S' be a distance-d scattered set in G', which corresponds to a set S^+ of edges in G^+. Observe that for any $v \in V(G)$, at most one edge of S^+ can be incident to the vertices of Q_v: otherwise, the corresponding two vertices in the line graph G' would have distance at most $\ell < d$. It is easy to see that if S^+ contains an edge incident to a vertex of Q_v, then we can always replace this edge with $e_{v,\ell}$, as this can only move it farther away from the other edges of S^+. Thus we may assume that every edge of S^+ is of the form $e_{v,\ell}$. Let us construct the set $S = \{v \mid e_{v,\ell} \in S^+\}$, which has size exactly k. Then S is independent in G: if $u, v \in S$ are adjacent in G, then there is a path of $2\ell + 1 = d$ edges in G^+ whose first an last edges are $e_{v,\ell}$ and $e_{u,\ell}$, respectively, hence the vertices of G' corresponding to them have distance at most $d - 1$.

If $d \geq 4$ is even, then the proof is similar, but we obtain the graph G^+ by first subdividing each edge and attaching paths of length $\ell = d/2 - 1$ to each original vertex. The proof proceeds in a similar way: if u and v are adjacent in G, then G^+ has a path of $2\ell + 2 = d$ edges whose first and last edges are $e_{v,\ell}$ and $e_{u,\ell}$, respectively, hence the vertices of G' corresponding to them have distance at most $d - 1$. □

There is a well-known and easy way of proving hardness of MIS on graphs with large girth: subdividing edges increases girth and the size of the largest independent set changes in a controlled way.

Lemma 5.4. *If there is an $2^{o(n)}$-time algorithm for MIS on n-vertex graphs of maximum degree 3 and girth more than g for any fixed $g > 0$, then the ETH fails.*

Proof. Let g be a fixed constant and let G be a simple graph with n vertices, m edges, and maximum degree 3 (hence $m = \mathcal{O}(n)$). We construct a graph G' by subdividing each edge with $2g$ new vertices. We have that G' has $n' = \mathcal{O}(n + gm) = \mathcal{O}(n)$ vertices, maximum degree 3, and girth at least $3(2g + 1) > g$. It is known and easy to show that subdividing the edges this way increases the size of the maximum independent set exactly by gm. Thus a $2^{o(n')}$- time algorithm for n'-vertex graphs of maximum degree 3 and girth at least

g could be used to give a $2^{o(n)}$-time algorithm for n-vertex graphs of maximum degree 3, hence the ETH would fail by Theorem 5.2. □

We use the lower bound of Lemma 5.4 to prove lower bounds for SCATTERED SET on C_t-free graphs.

Theorem 5.5. *Assuming the ETH, d-SCATTERED SET does not have a $2^{o(n)}$ algorithm on n-vertex C_t-free graphs with maximum degree 3 for any fixed $t \geq 3$ and $d \geq 2$.*

Proof. Let G be an n-vertex m-edge graph of maximum degree 3 and girth more than t. We construct a graph G' the following way: we subdivide each edge of G with $d-2$ new vertices to create a path of length $d-1$, and attach a path of length $d-1$ to each of the $(d-2)m = \mathcal{O}(dn)$ new vertices created. The resulting graph has maximum degree 3, $\mathcal{O}(d^2 n)$ vertices and edges, and girth more than $(d-1)t$ (hence it is C_t-free). We claim that $\alpha_d(G') = \alpha(G) + m(d-2)$ holds. This means that an $2^{o(n')}$-time algorithm for SCATTERED SET n'-vertex C_t-free graphs with maximum degree 3 would give a $2^{o(n)}$-time algorithm for n-vertex graphs of maximum degree 3 and girth more than t and this would violate the ETH by Lemma 5.4.

To see that $\alpha_d(G') = \alpha(G) + m(d-2)$ holds, consider first an independent set S of G. When constructing G', we attached $m(d-2)$ paths of length $d-1$. Let S' contain the degree-1 endpoints of these $m(d-2)$ paths, plus the vertices of G' corresponding to the vertices of S. It is easy to see that any two vertices of S' has distance at least d from each other: S is an independent set in G, hence the corresponding vertices in G' are at distance at least $2(d-1) \geq d$ from each other, while the degree-1 endpoints of the paths of length $d-1$ are at distance at least d from every other vertex that can potentially be in S'. This shows $\alpha_d(G') \geq \alpha(G) + m(d-2)$. Conversely, let S' be a set of vertices in G' that are at distance at least d from each other. The set S' contains two types of vertices: let S'_1 be the vertices that correspond to the original vertices of G and let S'_2 be the vertices that come from the $m(d-2)d$ new vertices introduced in the construction of G'. Observe that S'_2 can be covered by $m(d-2)$ paths of length $d-1$ and each such path can contain at most one vertex of S', hence at most $m(d-2)$ vertices of S' can be in S'_2. We claim that S'_1 can contain at most $\alpha(G)$ vertices, as $S' \cap S'_1$ corresponds to an independent set of G. Indeed, if u and v are adjacent vertices of G, then the corresponding two vertices of G' are at distance $d-1$, hence they cannot be both present in S'. This shows $\alpha_d(G') \leq \alpha(G) + m(d-2)$, completing the proof of the correctness of the reduction. □

As the following corollary shows, putting together Theorems 5.3 and 5.5 implies Theorem 1.4(2).

Corollary 5.6. *If H is a graph having a component that is not a path, then, assuming the ETH, d-SCATTERED SET has no $2^{o(n+m)}$-time algorithm on n-vertex m-edge H-free graphs for any fixed $d \geq 3$.*

Proof. Suppose first that H is not a forest and hence some cycle C_t for $t \geq 3$ appears as an induced subgraph in H. Then the class of H-free graphs is a superset of C_t-free graphs, which means that statement follows from Theorem 5.5 (which gives a lower bound for a more restricted class of graphs).

Assume therefore that H is a forest. Then it must have a component that is a tree, but not a path, hence it has a vertex v of degree at least 3. The neighbors of v are independent in the forest H, which means that the claw $K_{1,3}$ appears in H as an induced subgraph. Then the class of H-free graphs is a superset of claw-free graphs, which means that statement follows from Theorem 5.3 (which gives a lower bound for a more restricted class of graphs). □

References

[1] G. Agnarsson, P. Damaschke, and M. M. Halldórsson. Powers of geometric intersection graphs and dispersion algorithms. *Discrete Applied Mathematics*, 132(1-3):3–16, 2003.

[2] 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).

[3] V. E. Alekseev. Polynomial algorithm for finding the largest independent sets in graphs without forks. *Discrete Applied Mathematics*, 135(1–3):3–16, 2004.

[4] 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.

[5] V. Bafna, B. Narayanan, and R. Ravi. Nonoverlapping local alignments (weighted independent sets of axis-parallel rectangles). *Discrete Applied Mathematics*, 71:41–53, 1996.

[6] P. Berman. A $d/2$ approximation for maximum weight independent set in d-claw free graphs. In M. Halldórsson, editor, *Algorithm Theory – SWAT 2000, 7th Scandinavian Workshop on Algorithm Theory, Bergen Norway, July 2000, Proceedings*, pages 214–219. Springer, 2000.

[7] B. K. Bhattacharya and M. E. Houle. Generalized maximum independent sets for trees in subquadratic time. In *Algorithms and Computation, 10th International Symposium, ISAAC '99, Chennai, India, December 16-18, 1999, Proceedings*, pages 435–445, 1999.

[8] C. Brause. A subexponential-time algorithm for the Maximum Independent Set problem in P_t-free graphs. *Discrete Applied Mathematics*, 231:113–118, 2017.

[9] P. Chalermsook, B. Laekhanukit, and D. Nanongkai. Independent set, induced matching, and pricing: Connections and tight (subexponential time) approximation hardnesses. In *54th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2013, 26-29 October, 2013, Berkeley, CA, USA*, pages 370–379, 2013.

[10] M. Cygan, F. V. Fomin, L. Kowalik, D. Lokshtanov, D. Marx, M. Pilipczuk, M. Pilipczuk, and S. Saurabh. *Parameterized Algorithms.* Springer, 2015.

[11] H. Eto, F. Guo, and E. Miyano. Distance-d independent set problems for bipartite and chordal graphs. *J. Comb. Optim.*, 27(1):88–99, 2014.

[12] J. Flum and M. Grohe. *Parameterized Complexity Theory.* Texts in Theoretical Computer Science. An EATCS Series. Springer, 2006.

[13] A. Grzesik, T. Klimošová, M. Pilipczuk, and M. Pilipczuk. Polynomial-time algorithm for maximum weight independent set on P_6-free graphs. *CoRR*, abs/1707.05491, 2017.

[14] A. Gyárfás. Problems from the world surrounding perfect graphs. *Zastowania Matematyki, Applicationes Mathematicae*, XIX(3–4):413–441, 1987.

[15] M. M. Halldórson. Approximating discrete collections via local improvements. In *Proceedings of the sixth annual ACM-SIAM symposium on Discrete algorithms (SODA '95)*, pages 160–169. SIAM, 1995.

[16] R. Impagliazzo, R. Paturi, and F. Zane. Which problems have strongly exponential complexity? *J. Comput. Syst. Sci.*, 63(4):512–530, 2001.

[17] D. Lokshtanov, D. Marx, and S. Saurabh. Lower bounds based on the exponential time hypothesis. *Bulletin of the EATCS*, 105:41–72, 2011.

[18] D. Lokshtanov, M. Pilipczuk, and E. J. van Leeuwen. Independence and efficient domination on P_6-free graphs. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016, Arlington, VA, USA, January 10-12, 2016*, pages 1784–1803, 2016.

[19] D. Lokshtanov, 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.

[20] D. Marx and M. Pilipczuk. Optimal parameterized algorithms for planar facility location problems using Voronoi diagrams. In *Algorithms - ESA 2015 - 23rd Annual European Symposium, Patras, Greece, September 14-16, 2015, Proceedings*, pages 865–877, 2015.

[21] G. J. Minty. On maximal independent sets of vertices in claw-free graphs. *Journal of Combinatorial Theory, Series B*, 28(3):284–304, 1980.

[22] B. Randerath and I. Schiermeyer. On maximum independent sets in P_5-free graphs. *Discrete Applied Mathematics*, 158(9):1041–1044, 2010.

[23] D. J. Rosenkrantz, G. K. Tayi, and S. S. Ravi. Facility dispersion problems under capacity and cost constraints. *J. Comb. Optim.*, 4(1):7–33, 2000.

[24] N. Sbihi. Algorithme de recherche d'un stable de cardinalite maximum dans un graphe sans etoile. *Discrete Mathematics*, 29(1):53–76, 1980.

[25] D. M. Thilikos. Fast sub-exponential algorithms and compactness in planar graphs. In *Algorithms - ESA 2011 - 19th Annual European Symposium, Saarbrücken, Germany, September 5-9, 2011. Proceedings*, pages 358–369, 2011.

[26] G. Yu and O. Goldschmidt. On locally optimal independent sets and vertex covers. *Naval Research Logistics*,

43:737–748, 1996.

[27] D. Zuckerman. Linear degree extractors and the inapproximability of max clique and chromatic number. *Theory of Computing*, 3(1):103–128, 2007.