

Max Weight Independent Set in Graphs with No Long Claws: An Analog of the Gyárfás' Path Argument

Konrad Majewski 

Institute of Informatics, Faculty of Mathematics,
Informatics and Mechanics, University of Warsaw,
Poland

Jana Novotná 

Institute of Informatics, Faculty of Mathematics,
Informatics and Mechanics, University of Warsaw,
Poland

Marcin Pilipczuk 

Institute of Informatics, Faculty of Mathematics,
Informatics and Mechanics, University of Warsaw,
Poland

Marek Sokołowski 


Institute of Informatics, Faculty of Mathematics,
Informatics and Mechanics, University of Warsaw,
Poland

Tomáš Masařík 

Institute of Informatics, Faculty of Mathematics,
Informatics and Mechanics, University of Warsaw,
Poland

Karolina Okrasa 

Faculty of Mathematics and Information Science,
Warsaw University of Technology, Poland
Institute of Informatics, Faculty of Mathematics,
Informatics and Mechanics, University of Warsaw,
Poland

Paweł Rzażewski 

Faculty of Mathematics and Information Science,
Warsaw University of Technology, Poland
Institute of Informatics, Faculty of Mathematics,
Informatics and Mechanics, University of Warsaw,
Poland

Abstract

We revisit recent developments for the MAXIMUM WEIGHT INDEPENDENT SET problem in graphs excluding a subdivided claw $S_{t,t,t}$ as an induced subgraph [Chudnovsky, Pilipczuk, Pilipczuk, Thomassé, SODA 2020] and provide a subexponential-time algorithm with improved running time $2^{\mathcal{O}(\sqrt{n} \log n)}$ and a quasipolynomial-time approximation scheme with improved running time $2^{\mathcal{O}(\varepsilon^{-1} \log^5 n)}$.

The Gyárfás' path argument, a powerful tool that is the main building block for many algorithms in P_t -free graphs, ensures that given an n -vertex P_t -free graph, in polynomial time we can find a set P of at most $t - 1$ vertices, such that every connected component of $G - N[P]$ has at most $n/2$ vertices. Our main technical contribution is an analog of this result for $S_{t,t,t}$ -free graphs: given an n -vertex $S_{t,t,t}$ -free graph, in polynomial time we can find a set P of $\mathcal{O}(t \log n)$ vertices and an extended strip decomposition (an appropriate analog of the decomposition into connected components) of $G - N[P]$ such that every particle (an appropriate analog of a connected component to recurse on) of the said extended strip decomposition has at most $n/2$ vertices.

2012 ACM Subject Classification Theory of computation → Graph algorithms analysis; Mathematics of computing → Graph algorithms; Mathematics of computing → Approximation algorithms

Keywords and phrases Max Independent Set, subdivided claw, QPTAS, subexponential-time algorithm

Digital Object Identifier 10.4230/LIPIcs.ICALP.2022.93

Category Track A: Algorithms, Complexity and Games

Related Version *Preprint Version*: <https://arxiv.org/abs/2203.04836> [22]



© Konrad Majewski, Tomáš Masařík, Jana Novotná, Karolina Okrasa, Marcin Pilipczuk, Paweł Rzażewski, and Marek Sokołowski;
licensed under Creative Commons License CC-BY 4.0

49th International Colloquium on Automata, Languages, and Programming (ICALP 2022).

Editors: Mikołaj Bojańczyk, Emanuela Merelli, and David P. Woodruff;

Article No. 93; pp. 93:1–93:19



Leibniz International Proceedings in Informatics
LIPIcs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany





Funding This research is part of projects that has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme Grant Agreement 714704 (JN, KO, MP, PRz) and 948057 (KM, TM, MS).

Paweł Rzqżewski: Partially supported by Polish National Science Centre grant no. 2018/31/D/ST6/00062.

1 Introduction

The complexity of the MAXIMUM WEIGHT INDEPENDENT SET problem (MWIS for short), one of the classic combinatorial optimization problems, varies depending on the restrictions imposed on the input graph from polynomial-time solvable (e.g., in bipartite or chordal graphs) through known to admit a quasipolynomial-time algorithm (graphs with bounded longest induced path [15]), a polynomial-time approximation scheme and a fixed-parameter algorithm (planar graphs [8]), a quasipolynomial-time approximation scheme (graphs excluding a fixed subdivided claw as an induced subgraph [10, 11]), to being NP-hard and hard to approximate within $n^{1-\varepsilon}$ factor in general graphs [20, 26]. A methodological study of this behavior leads to the following question:

For which structures in the input graph, the assumption of their absence from the input graph makes MWIS easier and by how much?

The “absence of structures” notion can be made precise by specifying the forbidden structure and the containment relation, for example as a minor, topological minor, induced minor, subgraph, or induced subgraph. The last one – induced subgraph relation – is the weakest one, and thus the most expressible. This leads to the study of the complexity of MWIS in various hereditary graph classes, that is, graph classes closed under vertex deletion and thus definable by a (possibly infinite) list of forbidden induced subgraphs.

While a general classification of *all* hereditary graph classes with regards to the complexity of MWIS (or other classic graph problems) may be too complex, classifying graph classes with one forbidden induced subgraph looks more feasible. That is, we focus on H -free graphs, graphs excluding a fixed graph H as an induced subgraph. Furthermore, the complexity of a given problem (here, MWIS) in H -free graphs may indicate the impact of forbidding H as an induced subgraph on the complexity of MWIS in more general settings.

As observed by Alekseev [5, 6], the fact that MWIS remains NP-hard and APX-hard in subcubic graphs, together with the observation that subdividing every edge twice in a graph increases the size of the maximum independent set by exactly the number of edges of the original graph, leads to the conclusion that MWIS remains NP-hard and APX-hard in H -free graphs unless every connected component of H is a path or a tree with three leaves.

In what follows, for integers $t, a, b, c > 0$, by P_t we denote the path on t vertices, and by $S_{a,b,c}$ we denote the tree with three leaves within distance a, b , and c from the unique vertex of degree 3 of the tree. Since 1980s, it has been known that MWIS is polynomial-time solvable in P_4 -free graphs (because of their strong structural properties) and in $S_{1,1,1}$ -free graphs [23, 25] (because the notion of an augmenting path from the matching problem generalizes to MWIS in $S_{1,1,1}$ -free, i.e., claw-free graphs). For many years, only partial results in subclasses were obtained until the area started to develop rapidly around 2014.

Lokshtanov, Vatshelle, and Villanger [21] adapted the framework of potential maximal cliques [9] to show a polynomial-time algorithm for MWIS in P_5 -free graphs; this was later generalized to P_6 -free graphs [17] and other related graph classes [3, 4]. More importantly for this work, Bacsó et al. [7] observed that the classic Gyárfás’ path argument, developed to

show that for every fixed t the class of P_t -free graphs is χ -bounded [18, 19], also easily gives a subexponential-time algorithm for MWIS in P_t -free graphs. The crucial corollary of the Gyárfás' path argument lies in the following.

► **Theorem 1.** *Given an n -vertex graph G , one can in polynomial time find an induced path Q in G such that every connected component of $G - N[V(Q)]$ has at most $n/2$ vertices.*

For P_t -free graphs the said path Q has at most $t - 1$ vertices. Bacsó et al. [7] observed that branching either on the highest degree vertex (if this degree is larger than \sqrt{n}) or on the whole set $N[V(Q)]$ for the path Q coming from Theorem 1 (otherwise) gives an algorithm with running time bound exponential in $\sqrt{n} \cdot \text{poly}(t, \log n)$.

Chudnovsky, Pilipczuk, Pilipczuk, and Thomassé [10, 11] added to the mix an observation that a simple branching algorithm is able to get rid of *heavy* vertices: vertices of the input graph whose neighborhood contains a large fraction of the sought independent set. Once this branching is executed and the graph does not have heavy vertices, the set $N[Q]$ from Theorem 1 contains only a small fraction of the sought solution and, if one aims for an approximation algorithm, can be just sacrificed, yielding a quasipolynomial-time approximation scheme (QPTAS) for MWIS in P_t -free graphs. Using this as a starting point and leveraging on the celebrated *three-in-a-tree* theorem of Chudnovsky and Seymour [13], they developed a much more involved QPTAS and a subexponential algorithm (with running time bound $2^{n^{8/9} \text{poly}(\log n, t)}$) for MWIS in $S_{t,t,t}$ -free graphs.

Consider the following simple template for a branching algorithm for MWIS: if the current graph is disconnected, solve independently every connected component; otherwise, pick a vertex (*pivot*) v and branch whether v is in the sought independent set (recurring on $G - N[v]$) or not (recurring on $G - v$). The performance of such an algorithm highly depends on how we choose the pivot v . Theorem 1 suggests that in P_t -free graphs the vertices of Q may be good choices: there is only a bounded number of them, and the deletion of *the whole* neighborhood $N[V(Q)]$ splits G into multiplicatively smaller pieces. In a breakthrough result, Gartland and Lokshtanov [15] showed how to choose the pivot and measure the progress of the algorithm, obtaining a quasipolynomial-time algorithm for MWIS in P_t -free graphs. Later, Pilipczuk, Pilipczuk, and Rzażewski [24] provided an arguably simpler measure, leading to an improved (but still quasipolynomial) running time bound. These developments have been subsequently generalized to a larger class of problems beyond MWIS and to $C_{>t}$ -free graphs (graphs without induced cycle of length more than t) [16].

This progress suggests that MWIS may be actually solvable in polynomial time in H -free graphs for all open cases, that is, whenever H is a forest whose every connected component has at most three leaves. However, we seem still far from proving it: not only we do not know how to improve the quasipolynomial bounds of [15, 24] to polynomial ones, but also it remains unclear how to merge the approach of [15, 24] with the way how [10, 11] used the three-in-a-tree theorem [13].

In this work, we make a step in this direction, providing an analog of Theorem 1 for $S_{t,t,t}$ -free graphs. Before we state it, let us briefly discuss what we can hope for in the class of $S_{t,t,t}$ -free graphs.

Consider an example of a graph G being the line graph of a clique K . The graph G is $S_{1,1,1}$ -free, but does not admit any (balanced in any useful sense) separator of the form $N[P]$ for a small set $P \subseteq V(G)$. The MWIS problem on G translates back to the maximum weight matching problem in the clique K ; this problem is polynomial-time solvable, but with very different methods than branching. In particular, we are not aware of any way of solving maximum weight matching in a clique in quasipolynomial time by simple branching.

Thus, we expect that an algorithm for MWIS in $S_{t,t,t}$ -free graphs, given such a graph G , will discover that it is actually working with the line graph of a clique and apply maximum weight matching techniques to the preimage graph K .

Chudnovsky and Seymour, in their project to understand claw-free graphs [12], developed a good way of describing that a graph “looks like a line graph” by the notion of an *extended strip decomposition*. The formal definition can be found in Section 2. Here, we remark that in an extended strip decomposition of a graph, one can distinguish *particles* being induced subgraphs of the graph; an algorithm for MWIS can recurse on individual particles, compute the maximum weight independent sets there, and combine the results into a maximum weight independent set in the whole graph using a maximum weight matching algorithm on an auxiliary graph (cf. [10, 11]). Thus, an extended strip decomposition of a graph with particles of multiplicatively smaller size is very useful for recursion; it can be seen as an analog of splitting into connected components of multiplicatively smaller size, as it is in the case of the components of $G - N[V(Q)]$ in Theorem 1.

With the above discussion in mind, we can now state our main technical result.

- **Theorem 2.** *Given an n -vertex graph G and $t \geq 1$, one can in polynomial time either:*
- *output an induced copy of $S_{t,t,t}$ in G , or*
 - *output a set \mathcal{P} consisting of at most $11 \log n + 6$ induced paths in G , each of length at most $t + 1$, and a rigid extended strip decomposition of $G - N[\bigcup_{P \in \mathcal{P}} V(P)]$ whose every particle has at most $n/2$ vertices.*

Combining Theorem 2 with previously known algorithmic techniques, we derive two algorithms for MWIS in $S_{t,t,t}$ -free graphs. Actually, our algorithms work in a slightly more general setting. For integers $s, t \geq 1$, by $sS_{t,t,t}$ we denote the graph with s connected components, each isomorphic to $S_{t,t,t}$. Recall that by the observation of Alekseev [5, 6] the only graphs H , for which we can hope for tractability results for MWIS in H -free graphs, are forests whose every component has at most three leaves. We observe that each such H is contained in $sS_{t,t,t}$, for some s and t depending on H . Thus algorithms for $sS_{t,t,t}$ -free graphs, for every s and t , cover *all* potential positive cases.

First, we observe that the statement of Theorem 2 seamlessly combines with the method how [7] obtained a subexponential-time algorithm for MWIS in P_t -free graphs. As a result, we obtain a subexponential-time algorithm for MWIS in $sS_{t,t,t}$ -free graphs with improved running time as compared to [10, 11].

- **Theorem 3.** *Let $s, t \geq 1$ be constants. Given an n -vertex $sS_{t,t,t}$ -free graph G with weights on vertices, one can in time exponential in $\mathcal{O}(\sqrt{n} \log n)$ compute an independent set in G of maximum possible weight.*

Second, we observe that the statement of Theorem 2 again seamlessly combines with the method how [10, 11] obtained a QPTAS for MWIS in P_t -free graphs, obtaining an arguably simpler QPTAS for MWIS in $sS_{t,t,t}$ -free graphs with improved running time (compared to [10, 11]).

- **Theorem 4.** *Let $s, t \geq 1$ be constants. Given an n -vertex $sS_{t,t,t}$ -free graph G with weights on vertices, and a real $\varepsilon > 0$, one can in time exponential in $\mathcal{O}(\varepsilon^{-1} \log^5 n)$ compute an independent set in G that is within a factor of $(1 - \varepsilon)$ of the maximum possible weight.*

After preliminaries in Section 2, we prove Theorem 2 in Section 3. Proofs of Theorems 3 and 4 are provided in Section 4. Finally, we discuss future steps in Section 5.

2 Preliminaries

Notation. For a family \mathcal{Q} of sets, by $\bigcup \mathcal{Q}$ we denote $\bigcup_{Q \in \mathcal{Q}} Q$. If the base of a logarithmic function is not specified, we mean the logarithm of base 2, i.e., $\log n := \log_2 n$. For a function $\mathfrak{w} : V \rightarrow \mathbb{Z}_{\geq 0}$ and subset $V' \subseteq V$, we denote $\mathfrak{w}(V') := \sum_{v \in V'} \mathfrak{w}(v)$.

Let G be a graph. For $X \subseteq V(G)$, by $G[X]$ we denote the subgraph of G induced by X , i.e., $(X, \{uv \in E(G) : u, v \in X\})$. If the graph G is clear from the context, we will often identify induced subgraphs with their vertex sets. The sets $X, Y \subseteq V(G)$ are *complete* to each other if for every $x \in X$ and $y \in Y$ the edge xy is present in G . Note that this, in particular, implies that X and Y are disjoint. We say that two sets X, Y *touch* if $X \cap Y \neq \emptyset$ or there is an edge with one end in X and another in Y .

For a vertex v , by $N_G(v)$ we denote the set of neighbors of v , and by $N_G[v]$ we denote the set $N_G(v) \cup \{v\}$. For a set $X \subseteq V(G)$, we also define $N_G(X) := \bigcup_{v \in X} N_G(v) - X$, and $N_G[X] = N_G(X) \cup X$. If it does not lead to confusion, we omit the subscript and write simply $N(\cdot)$ and $N[\cdot]$.

By $T(G)$, we denote the set of all triangles in G . Similarly to writing $xy \in E(G)$, we will write $xyz \in T(G)$ to indicate that $G[\{x, y, z\}] \simeq K_3$.

Extended strip decompositions. Now let us define a certain graph decomposition which will play an important role in the paper. An *extended strip decomposition* of a graph G is a pair (H, η) that consists of:

- a simple graph H ,
- a set $\eta(x) \subseteq V(G)$ for every $x \in V(H)$,
- a set $\eta(xy) \subseteq V(G)$ for every $xy \in E(H)$, and its subsets $\eta(xy, x), \eta(xy, y) \subseteq \eta(xy)$,
- a set $\eta(xyz) \subseteq V(G)$ for every $xyz \in T(H)$,


which satisfy the following properties (also see Figure 1):

1. $\{\eta(o) \mid o \in V(H) \cup E(H) \cup T(H)\}$ is a partition of $V(G)$,
2. for every $x \in V(H)$ and every distinct $y, z \in N_H(x)$, the set $\eta(xy, x)$ is complete to $\eta(xz, x)$,
3. every $uv \in E(G)$ is contained in one of the sets $\eta(o)$ for $o \in V(H) \cup E(H) \cup T(H)$, or is as follows:
 - $u \in \eta(xy, x), v \in \eta(xz, x)$ for some $x \in V(H)$ and $y, z \in N_H(x)$, or
 - $u \in \eta(xy, x), v \in \eta(x)$ for some $xy \in E(H)$, or
 - $u \in \eta(xyz)$ and $v \in \eta(xy, x) \cap \eta(xy, y)$ for some $xyz \in T(H)$.

Note that for an extended strip decomposition (H, η) of a graph G , the number of vertices of H can be much larger than the number of vertices of G . However, in such case many sets $\eta(\cdot)$ are empty and thus H is “unnecessarily complicated.” An extended strip decomposition (H, η) is *rigid* if (i) for every $xy \in E(H)$ it holds that $\eta(xy, x) \neq \emptyset$, and (ii) for every $x \in V(H)$ such that x is an isolated vertex it holds that $\eta(x) \neq \emptyset$. Observe that if we *restrict* η to $V' \subseteq V(G)$, i.e. we keep in η only vertices of V' , (H, η) after the restriction remains an extended strip decomposition, but it might not be rigid any more.

► **Observation 5.** Let (H, η) be a rigid extended strip decomposition of an n -vertex graph G . Then $|E(H)| \leq n$ and $|V(H)| \leq 2n$.

Proof. Recall that since (H, η) is rigid, for every $xy \in E(H)$ we have that $\emptyset \neq \eta(xy, x) \subseteq \eta(xy)$, and for every isolated vertex x of H we have $\eta(x) \neq \emptyset$.



■ **Figure 1** A graph H and an extended strip decomposition (H, η) of a graph G . Sets $\eta(\cdot)$ corresponding to vertices, edges, and the triangle of H are marked green, blue, and orange, respectively. The edges between distinct sets are drawn thick if they must exist, and thin if they may exist.

Let V_0 and V_+ denote, respectively, the sets of vertices of H with degree 0 and more than 0. As the family $\{\eta(xy) \mid xy \in E(H)\} \cup \{\eta(x) \mid x \in V_0\}$ consists of pairwise disjoint nonempty subsets of $V(G)$, we conclude that $|E(H)| + |V_0| \leq n$ and therefore $|E(H)| \leq n$.

Note that by the handshaking lemma we have $|E(H)| \geq |V_+|/2$, and so $|V(H)| = |V_0| + |V_+| \leq |V_0| + 2|E(H)| \leq 2n$ by the previous argument. ◀

We say that a vertex $v \in V(G)$ is *peripheral* in (H, η) if there is a degree-one vertex x of H , such that $\eta(xy, x) = \{v\}$, where y is the (unique) neighbor of x in H . For a set $Z \subseteq V(G)$, we say that (H, η) is an *extended strip decomposition* of (G, Z) if H has $|Z|$ degree-one vertices and each vertex of Z is peripheral in (H, η) .

The following theorem by Chudnovsky and Seymour [13] is a slight strengthening of their celebrated solution of the famous *three-in-a-tree* problem. We will use it as a black-box to build extended strip decompositions.

► **Theorem 6** (Chudnovsky, Seymour [13, Section 6]). *Let G be an n -vertex graph and consider $Z \subseteq V(G)$ with $|Z| \geq 2$. There is an algorithm that runs in time $\mathcal{O}(n^5)$ and returns one of the following:*

- an induced subtree of G containing at least three elements of Z ,
- a rigid extended strip decomposition (H, η) of (G, Z) .

Let us point out that actually, an extended strip decomposition produced by Theorem 6 satisfies more structural properties, but of our purpose, we will only use the fact that it is rigid.

Particles of extended strip decompositions. Let (H, η) be an extended strip decomposition of a graph G . We introduce some special subsets of $V(G)$ called *particles*, divided into five types.

- vertex particle: $A_x := \eta(x)$ for each $x \in V(H)$,
- edge interior particle: $A_{xy}^\perp := \eta(xy) - (\eta(xy, x) \cup \eta(xy, y))$ for each $xy \in E(H)$,
- half-edge particle: $A_{xy}^x := \eta(x) \cup \eta(xy) - \eta(xy, y)$ for each $xy \in E(H)$,
- full edge particle: $A_{xy}^{xy} := \eta(x) \cup \eta(y) \cup \eta(xy) \cup \bigcup_{z : xyz \in T(H)} \eta(xyz)$ for each $xy \in E(H)$,
- triangle particle: $A_{xyz} := \eta(xyz)$ for each $xyz \in T(H)$.

Observe that the number of all particles of (H, η) is at most $\mathcal{O}(|V(H)|^3)$. However, the number of nonempty particles is linear in the number of vertices of G .

► **Observation 7.** *Let (H, η) be an extended strip decomposition of an n -vertex graph. Then the number of nonempty particles of (H, η) is bounded by $4n$.*

Proof. Let V', E', T' , respectively, be the subsets consisting of those elements o of $V(H)$, $E(H)$, or $T(H)$, for which $\eta(o) \neq \emptyset$. Observe that each $o \in V' \cup T'$ gives rise to one nonempty particle A_o , and each $xy \in E'$ gives rise to at most four nonempty particles: $A_{xy}^\perp, A_{xy}^x, A_{xy}^y, A_{xy}^{xy}$. Moreover, since $\{\eta(o) \mid o \in V' \cup E' \cup T'\}$ are pairwise disjoint subsets of $V(G)$, we have that $|V'| + |E'| + |T'| \leq n$. Hence, the number of nonempty particles is bounded by $|V'| + |T'| + 4|E'| = (|V'| + |T'| + |E'|) + 3|E'| \leq 4n$. ◀

A vertex particle A_x is *trivial* if x is an isolated vertex in H . Similarly, an extended strip decomposition (H, η) is *trivial* if H is an edgeless graph. The following observation follows immediately from the definitions of an extended strip decomposition and particles.

► **Observation 8.** *Let (H, η) be an extended strip decomposition of a graph G . For each $xy \in E(H)$ the following hold:*

1. $A_{xy}^\perp \subseteq A_{xy}^x \subseteq A_{xy}^{xy}$,
2. for any $v_x \in \eta(xy, x)$ and $v_y \in \eta(xy, y)$ we have $N(A_{xy}^{xy}) = N(v_x) \cup N(v_y) - A_{xy}^{xy}$.

We conclude this section by recalling an important property of particles of extended strip decompositions, observed by Chudnovsky et al. [10].

► **Theorem 9** (Chudnovsky et al. [10, Lemma 6.8]). *Let (H, η) be an extended strip decomposition of G . Suppose P_1, P_2, P_3 are three induced paths in G that do not touch each other, and moreover each of P_1, P_2, P_3 has an endvertex that is peripheral in (H, η) . Then in (H, η) there is no particle that touches each of P_1, P_2, P_3 .*

3 Main result

In this section, we prove our main result, i.e., Theorem 2. Let us first give an overview of our approach. We present a recursive algorithm that, for a given graph G , will return one of the outcomes of Theorem 2. Let $n := |V(G)|$ be the number of vertices in the input graph; the value of n will not change throughout the recursive steps of the algorithm. We start with finding a Gyárfás path Q navigating towards the largest component in G . That is, by Theorem 1, we find Q such that each connected component of $G - N[Q]$ is of size at most $\frac{n}{2}$. Finding such a small connected component is a great outcome as we can readily include it as a small trivial vertex particle of an extended strip decomposition we are constructing. We say that a particle is *small* if its size is at most $\frac{n}{2}$, and an extended strip decomposition is *refined* if all its particles are small. Observe that if $|Q| \leq 3t + 1$, we immediately get the

desired refined extended strip decomposition of G . Otherwise, we proceed to the main part of the algorithm. At each step, we will remove some vertices from Q , and will measure the progress of our algorithm in the number of the remaining vertices of Q .

Formally, we create a set \mathcal{Q} of at most two non-touching induced paths such that $\bigcup \mathcal{Q} \subseteq Q$. At each step of recursion we obtain a set $\hat{\mathcal{Q}}$ with $|\bigcup \hat{\mathcal{Q}}| \leq \frac{2}{3}|\bigcup \mathcal{Q}|$ that represents \mathcal{Q} for the next step. Hence, in $11 \log n$ recursive steps, $|\bigcup \mathcal{Q}|$ drops below $3t + 1$. In the base case of the recursion, when $|\bigcup \mathcal{Q}| \leq 3t + 1$, we return the refined trivial extended strip decomposition ensured by maintaining the property that $G - N[\bigcup \mathcal{Q}]$ has connected components of size at most $\frac{n}{2}$ throughout the recursive steps. In each step of recursion, we further split the induced path(s) in \mathcal{Q} so we are able to use Theorem 6 to obtain an extended strip decomposition (H, η) . If (H, η) is already refined, then we are done. Otherwise, it contains a particle A that is not small. We use Theorem 9 to select at most two paths touching A . Then it is easy to separate A with the respective touching paths from the rest of the graph. The graph induced by A and the touching paths form a smaller instance, i.e., an instance where $|\bigcup \mathcal{Q}|$ drops by a factor of $\frac{2}{3}$. We need to ensure that at every recursive step, we include only a constant number of paths of length $t + 1$ into \mathcal{P} (i.e., the set of paths in the second outcome of Theorem 2). We now prove the core recursive formulation of the algorithm formally.

► **Lemma 10** (Recursion). *Given a graph G and a set \mathcal{Q} of at most two induced paths (vertex disjoint non-adjacent), and a refined extended strip decomposition of $G - N[\bigcup \mathcal{Q}]$. In polynomial time, we can output one of the following:*

- an induced copy of $S_{t,t,t}$ in G , or
- \mathcal{P} , $X \subseteq N[\bigcup \mathcal{P}]$, and a refined extended strip decomposition (H, η) of $G - X$, so that $|\mathcal{P}| \leq 6 \log_{3/2}(|\bigcup \mathcal{Q}|) + 6$ and the longest path in \mathcal{P} has at most $t + 1$ vertices.

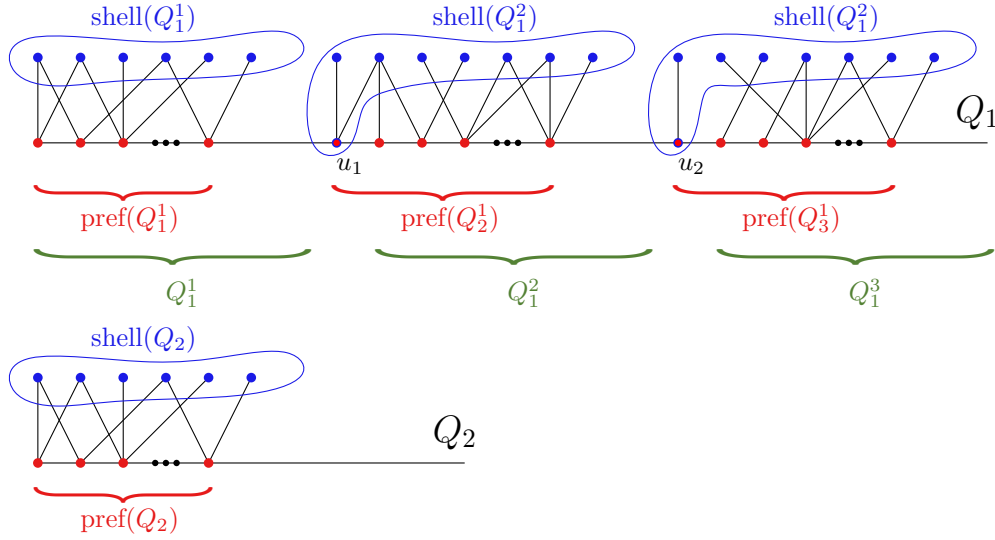
Proof. If the longest path of \mathcal{Q} has at most $3t + 1$ vertices, return $\mathcal{P} := \mathcal{Q}$ where each path in \mathcal{P} may be further split in at most three paths on at most $t + 1$ vertices, and $X := N[\bigcup \mathcal{P}]$. Hence, we output the extended strip decomposition we were given by the assumptions of the lemma.

Otherwise, let Q_1 be the longest path in \mathcal{Q} . Let u_1 and u_2 be the $\left(\left\lfloor \frac{|Q_1|}{3} \right\rfloor + 1\right)$ -th and the $\left(2 \left\lfloor \frac{|Q_1|}{3} \right\rfloor + 2\right)$ -th vertex of Q_1 , respectively. The removal of u_1 and u_2 from Q_1 divides the path into three induced non-touching subpaths Q_1^1 , Q_1^2 , and Q_1^3 , each of length at least t . Let Q_2 be the remaining path of \mathcal{Q} , should it exist. We define $\mathcal{S} := \{Q_1^1, Q_1^2, Q_1^3, Q_2\}$ if Q_2 exists, or $\mathcal{S} := \{Q_1^1, Q_1^2, Q_1^3\}$, otherwise. Consult Figure 2 to see an overview of the definitions described in this paragraph. For each path $P \in \mathcal{S}$ we define $\text{pref}(P)$ as the set comprising:

- first $t - 1$ vertices of P (or all vertices of P if $|P| < t - 1$), and
- the separating vertex of Q_1 directly preceding P if $P \in \{Q_1^2, Q_1^3\}$.

It can be easily seen that the set of vertices $\text{pref}(P)$ forms an induced path of length at most t . We finally define *shells* of paths in \mathcal{S} . Given a path $P \in \mathcal{S}$, we set $\text{shell}(P) := N[\text{pref}(P)] - \bigcup \mathcal{S}$ if $|P| \geq t$ and $\text{shell}(P) := N[\text{pref}(P)]$ otherwise. Intuitively, if $|P| < t$, the shell of P takes the whole neighborhood as we do not have a use for a short path in the next stage of our algorithm. For a long enough path P ($|P| > t$), the shell of P intersects all short paths (shorter than t) connecting the first vertex of P with the rest of the graph. In other words, each path from the first vertex of P to any vertex of $G - \text{shell}(P)$ outside of P will have length at least t . To ease the notation, we define $\mathcal{S}_{\geq t} := \{P \in \mathcal{S} \mid |P| \geq t\}$, $\text{shell}(\mathcal{S}) := \bigcup_{P \in \mathcal{S}} \text{shell}(P)$, and $\text{pref}(\mathcal{S}) := \bigcup_{P \in \mathcal{S}} \text{pref}(P)$.

Now, we use the algorithm from Theorem 6 on Z being the set of the first vertices of paths in $\mathcal{S}_{\geq t}$ and the graph defined as $G - \text{shell}(\mathcal{S})$. If Theorem 6 produced an induced tree with three leaves among Z , it contains an induced $S_{t,t,t}$, since those must have been



■ **Figure 2** Definitions of $\text{pref}(S)$ and $\text{shell}(S)$ in case of $|Q_2| \geq t$.

induced branches at least t vertices long in $G - \text{shell}(S)$. Hence, we obtained an extended strip decomposition (H', η') of $G - \text{shell}(S)$. If the obtained decomposition is refined, we return $\mathcal{P} := \text{pref}(S)$, $X := \text{shell}(S)$, and the extended strip decomposition $(H := H', \eta := \eta')$.

Therefore, the obtained extended strip decomposition (H', η') of $G - \text{shell}(S)$ contains a particle A which is not small, i.e., A is composed of at least $\frac{n}{2}$ vertices. As every vertex in Z is peripheral in (H', η') , we know that no three paths in $\mathcal{S}_{\geq t}$ touch one particle by Theorem 9. Therefore, we take the set $\hat{\mathcal{Q}}$ of at most two paths, say P_1 and P_2 , touching A (for convenience, let P_1 or P_2 be an empty set if it does not exist). We now compute the maximum proportion of $\bigcup \mathcal{Q}$ put to $\hat{\mathcal{Q}}$. If both $P_1, P_2 \subseteq Q_1$, then this fraction is at most $\frac{2}{3}$ as by the definition $|Q_1^i| \leq \frac{|Q_1|}{3}$, for $i \in \{1, 2, 3\}$. If one is Q_2 and the other comes from Q_1 , then we estimate $a + \frac{1-a}{3} = \frac{2a+1}{3} \leq \frac{2}{3}$ for $a = |Q_2|/|\bigcup \mathcal{Q}| \leq \frac{1}{2}$. Hence, we know that $|\bigcup \hat{\mathcal{Q}}| \leq \frac{2}{3} |\bigcup \mathcal{Q}|$. We define $\hat{G} := A \cup P_1 \cup P_2$ to use Lemma 10 on a smaller instance. Now, we need to verify that the assumption of the lemma holds. We claim the following:

▷ **Claim 11.** $\hat{G} - N[\bigcup \hat{\mathcal{Q}}]$ has a refined extended strip decomposition.

Proof. As \hat{G} is an induced subgraph of G and $G - N[\bigcup \mathcal{Q}]$ has a refined extended strip decomposition, we know that $\hat{G} - N[\bigcup \mathcal{Q}]$ has a refined extended strip decomposition. First, recall that $N[u_1] - (Q_1^1 \cup Q_1^2) \subseteq \text{shell}(Q_1^2)$, which is disjoint with $V(\hat{G})$. Analogously $N[u_2] - (Q_1^2 \cup Q_1^3)$ is disjoint with $V(\hat{G})$. Also, if $|Q_2| < t$ then Q_2 is disjoint with $V(\hat{G})$ as well. Hence, $\hat{G} - N[\bigcup \mathcal{Q}] \simeq \hat{G} - N[\bigcup \mathcal{S}_{\geq t}]$. Also, recall that the only paths among $\mathcal{S}_{\geq t}$ that touch A are in $\hat{\mathcal{Q}}$. Hence, observe that $\hat{G} - N[\bigcup \mathcal{S}_{\geq t}] \simeq \hat{G} - N[\bigcup \hat{\mathcal{Q}}]$. \triangleleft

Therefore, we can apply Lemma 10 inductively on \hat{G} and $\hat{\mathcal{Q}}$, obtaining $\hat{\mathcal{P}}$ and \hat{X} , and a refined extended strip decomposition $(\hat{H}, \hat{\eta})$ of $\hat{G} - \hat{X}$. We need to combine the extended strip decomposition obtained from the recursion with the extended strip decomposition (H', η') we obtained earlier.

We can always suppose that particle A is of type A_{xy}^{xy} for some edge $xy \in E(H')$, unless A is of type A_x for an isolated vertex $x \in V(H')$. That is because A_{xy}^{xy} is the superset of all possible particle types. As Theorem 6 gives us that both $\eta'(xy, x)$ and $\eta'(xy, y)$ are nonempty, we can select $v_x \in \eta'(xy, x)$ and $v_y \in \eta'(xy, y)$ (possibly $v_x = v_y$). By Observation 8, the set

$$X' := (N(v_y) \cup N(v_x)) - V(A)$$

separates A from the rest of G . Set $\mathcal{P}' := \{\{v_x\}, \{v_y\}\}$. In the case of A_x such that $x \in V(H)$ is an isolated vertex, we set $\mathcal{P}' := \emptyset$ and $X' := \emptyset$ and still such A is separated from the rest of G by X' . We return:

- $\mathcal{P} := \hat{\mathcal{P}} \cup \mathcal{P}' \cup \text{pref}(\mathcal{S})$,
- $X := \hat{X} \cup X' \cup \text{shell}(\mathcal{S})$,
- an extended strip decomposition (H, η) of $G - X$, where H is \hat{H} with an additional isolated vertex w , and η is $\hat{\eta}$ restricted only to vertices in $A - X$ with an additional trivial vertex particle $\eta(w)$ containing all vertices of $G - X - A$.

Recall that during the recursion, we do not require rigidity, therefore, we do not mind restricting η only to a subset of vertices. Note that indeed, $G - X - A$ may contain parts of P_1 or P_2 , however, $\eta(w)$ does not touch any vertices contained in $\hat{\eta}$ restricted to $A - X$ as $X' \subseteq X$ completely separated $A - X$ from $G - X - A$.

We compute that $|\mathcal{P}| \leq 6 + 6 \log_{3/2} \left(\left| \bigcup \hat{\mathcal{Q}} \right| \right) + 6 \leq 6 \log_{3/2} (|\bigcup \mathcal{Q}|) + 6$ as we added at most six new paths into \mathcal{P} . Observe that the described algorithm runs in polynomial time as we just computed that the depth of recurrence is logarithmic in $|\bigcup \mathcal{Q}| \leq |V(G)|$ and each recursive call takes polynomial time in the size of G . ◀

Proof of Theorem 2. Using Theorem 1 we find a Gyárfás path Q . We get the desired outcome by Lemma 10 on G with $\mathcal{Q} := \{Q\}$. The extended strip decomposition needed by the lemma's assumption is trivial. That is, each connected component of $G - Q$ is represented by a vertex particle of small size. Note that Lemma 10 provides an extended strip decomposition of $G - X$, where $X \subseteq \bigcup \mathcal{P}$ and Theorem 2 only requires an extended strip decomposition of $G - \bigcup \mathcal{P}$, so we can restrict the obtained extended strip decomposition to $V(G) - \bigcup \mathcal{P}$. We conclude the proof of Theorem 2 by the following calculation:

$$6 \log_{3/2} n + 6 \leq 11 \log n + 6.$$

Note that for any extended strip decomposition (H, η) we can easily add the assumption that sets $\eta(xy, x) \neq \emptyset$ for any edge $xy \in E(H)$. As suppose $\eta(xy, x) = \emptyset$; then we can update (H, η) by adding $\eta(xy)$ to $\eta(y)$ and removing xy from H . Moreover, we can simply remove any empty trivial vertex particle from η and the corresponding isolated vertex from H . Therefore, we may suppose that the obtained extended strip decomposition is rigid. ◀

In the following simple corollary we apply Theorem 2 to $sS_{t,t,t}$ -free graphs, for some $s, t \geq 1$.

► **Corollary 12.** *Let $s \geq 1, t \geq 1$ be constants. Let G be an $sS_{t,t,t}$ -free graph on n vertices. Then in polynomial time we can find a set X consisting of at most*

$$(s - 1)(3t + 1) + (11 \log n + 6)(t + 1)$$

vertices and a rigid extended strip decomposition of $G - N[X]$ whose every particle has at most $n/2$ vertices.

Proof. Induction on s . If $s = 1$, then we obtain the result immediately by Theorem 2. Thus let us assume that $s \geq 2$ and the theorem holds for $(s - 1)S_{t,t,t}$ -free graphs.

We exhaustively check if there is some $Y \subseteq V(G)$ with $|Y| = 3t + 1$, such that $G[Y] \simeq S_{t,t,t}$; we can do it in time $n^{3(t+1)+O(1)} = n^{O(1)}$. If such Y does not exist, then we can immediately apply Theorem 2, and the proof is complete. Thus suppose that Y exists.

We observe that the graph $G' := G - N[Y]$ is $(s - 1)S_{t,t,t}$ -free. Denote $n' := |V(G')|$. By the inductive assumption, in time $(n')^{O(1)} = n^{O(1)}$ we can obtain a set $X' \subseteq V(G')$ of size at most $(s - 2)(3t + 1) + (11 \log n' + 6)(t + 1)$ and a rigid extended strip decomposition (H, η) of $G' - N[X']$ whose every particle is of size at most $n'/2$.

We set $X = Y \cup X'$. Now X and (H, η) satisfy the statement of the theorem, as $G' - N[X'] = G - N[X]$ and $n' \leq n$. The total running time is polynomial in n as the depth of the recursion is $s - 1$. \blacktriangleleft

4 Algorithmic applications

In this section we will show how to combine Theorem 2 with the approach of Chudnovsky et al. [10, 11] in order to obtain a QPTAS and a subexponential-time algorithm for MWIS in $S_{t,t,t}$ -free graphs, i.e., we prove Theorems 3 and 4.

Both algorithms follow the same general outline; let us sketch it before we get into the details of each particular case. Each algorithm is a recursive procedure, which consists of two phases. In the first one, we deal with the vertices of G that are *heavy*, which means that their neighborhood is “large”, where the exact meaning of “large” depends on the particular algorithm.

Once there are no heavy vertices, i.e., the neighborhood of each vertex is “small”, we proceed to the second phase. We call Corollary 12 for the current instance G , obtaining a small-sized set X and a rigid extended strip decomposition (H, η) of $G - N[X]$, whose every particle is of small size. The crux is that since we are in the second phase, all vertices in X are not heavy, and since X is of small size, the whole set $N[X]$ is “small”. We treat $N[X]$ separately in a way that depends on the particular algorithm.

Next, for each particle A of (H, η) , we call the algorithm recursively for $G[A]$, obtaining (a good approximation of) a maximum-weight independent set in $G[A]$. Finally, we combine the obtained results to derive (a good approximation of) a maximum-weight independent set in G . This last step is based on the idea of Chudnovsky et al. [10, 11] to reduce the problem to finding a maximum-weight matching in a graph obtained by a simple modification of H . Since the size of H is linear in $|V(G)|$ (by Observation 5), this problem can be solved in time polynomial in $|V(G)|$ using, e.g., the classic algorithm of Edmonds [14]. The last step is encapsulated in the following lemma, whose exact statement comes from Abrishami et al. [1].

► **Lemma 13** (Chudnovsky et al. [10, 11]). *Let $\varsigma \in [0, 1]$ be a real number. Let G be an n -vertex graph equipped with a weight function $\mathbf{w} : V(G) \rightarrow \mathbb{Z}_{\geq 0}$. Suppose that G is given along with an extended strip decomposition (H, η) , where H has N vertices. Let $I_0 \subseteq V(G)$ be a fixed independent set in G . Furthermore, assume that for each particle A of (H, η) we are given an independent set $I(A)$ in $G[A]$ such that $\mathbf{w}(I(A)) \geq \varsigma \cdot \mathbf{w}(I_0 \cap A)$. Then in time polynomial in $n + N$ we can compute an independent set I in G such that $\mathbf{w}(I) \geq \varsigma \cdot \mathbf{w}(I_0)$.*

Let us stress out that the algorithm from Lemma 13 does not need to know the value of ς or the independent set I_0 .

The main difference between our approach and the one of Chudnovsky et al. [11] is that we use Theorem 2 and its consequence, i.e., Corollary 12. The previous algorithms used a similar statement but with a worse (and much more involved) guarantee on the size of X and each particle. Furthermore, the way we obtain our set X is significantly simpler.

4.1 Proof of Theorem 3

Before we proceed to the proof, let us first explain the meaning of “small”, and how to deal with $N[X]$ in this particular case. Here the neighborhood of a vertex is “small” if it has few vertices (more specifically, at most $\sqrt{n/t}$). In the first phase, we deal with heavy vertices v (i.e., of large degree) with simple branching: we guess whether v is included in our optimum solution or not. Since the degree of v is large, in the first branch, we obtain significant progress, which is enough to obtain a subexponential running time.

In the second phase, since $N[X]$ is the neighborhood of $\mathcal{O}(\log n)$ vertices, each of degree $\mathcal{O}(\sqrt{n})$, the total size of $N[X]$ is $\mathcal{O}(\sqrt{n} \log n)$. Thus we can afford to exhaustively guess the intersection of our optimum solution with $N[X]$.

Proof of Theorem 3. Let $s, t \geq 1$ be constants and let (G, \mathfrak{w}) be an instance of MWIS, where G is $sS_{t,t,t}$ -free and has n vertices. We observe that if n is small, i.e., bounded by a constant, then we can solve the problem by brute force. Thus we assume that $n \geq n_0$, where n_0 is a constant (depending on s and t) whose exact value follows from the reasoning below.

First, consider the case that there exists $v \in V(G)$ such that $\deg v \geq \sqrt{n/t}$. We branch on including v in the final solution: we either delete v from G , or we delete $N[v]$ and add v to the solution returned by the recursive call. Then we output the one of these two solutions that has a larger weight. The correctness of this step of the algorithm is straightforward.

Hence, we can assume that for every $v \in V(G)$ it holds that $\deg v \leq \sqrt{n/t}$. By Corollary 12, since G is $sS_{t,t,t}$ -free, we obtain a set X of size $(s-1)(3t+1) + (11 \log n + 6)(t+1) \leq 12(t+1) \log n$ (here we use that n is large), and a rigid extended strip decomposition (H, η) of $G' = G - N[X]$ whose every particle has at most $n/2$ vertices.

We exhaustively guess an independent set $J \subseteq N[X]$; think of it as an intersection of the intended optimum solution with $N[X]$. Consider the graph $G'' := G' - N[J]$. We modify (H, η) by removing the vertices from $N[J]$ from the sets $\eta(\cdot)$. Let us call the obtained strip decomposition (H, η') ; note that it might not be rigid. We call the algorithm recursively for the subgraph $G''[A]$ for every nonempty particle A of (H, η') . Let $I(A)$ be the solution. If $A = \emptyset$, then $I(A) = \emptyset$. By the inductive assumption $I(A)$ is a maximum-weight independent set in $G''[A]$. Then we use Lemma 13 for $\varsigma = 1$ to combine the solutions into a maximum-weight independent set I_J of G'' . Finally, we return the independent set $J \cup I_J$ whose weight is maximum over all choices of J . Note that the correctness of this step is guaranteed by the exhaustive guessing of J and Lemma 13.

Running time. Let $F(n)$ denote the running time of our algorithm for n -vertex instances. We prove that $F(n) = 2^{\mathcal{O}(\sqrt{tn} \log n)}$. If $n < n_0$, then the claim clearly holds. So let us assume that $n \geq n_0$.

In the first case we call the algorithm for two instances, one of size $n-1$ and one of size at most $n - \sqrt{n/t}$. Hence,

$$F(n) \leq F(n-1) + F(n - \sqrt{n/t}) = 2^{\mathcal{O}\left(\frac{n \log n}{\sqrt{n/t}}\right)} \leq 2^{\mathcal{O}(\sqrt{tn} \log n)}.$$

Here we skip the description how this recursion is solved, as it is pretty standard. For a formal proof we refer the reader to Bacsó et al. [7, Lemma 1].

It remains to analyze the running time of the step in which the maximum degree of vertices in G is bounded by $\sqrt{n/t}$. Corollary 12 asserts that we obtain X and the rigid extended strip decomposition (H, η) of $G' = G - N[X]$ in time polynomial in n . There are $2^{\mathcal{O}(\sqrt{n/t} \cdot t \log n)} = 2^{\mathcal{O}(\sqrt{nt} \log n)}$ ways of choosing the set J . In polynomial time we modify (H, η) into (H, η') .

Observe that while (H, η') might not be rigid, it was obtained from a rigid extended strip decomposition (H, η) by deleting some vertices from the sets $\eta(\cdot)$. In particular, both decompositions have the same sets of particles, and every nonempty particle of (H, η') is also a nonempty particle of (H, η) . Thus by Observation 7 we call the algorithm recursively for at most $4n$ nonempty particles, each of size at most $n/2$. By Observation 5, the total number of particles of (H, η') is polynomial in n . Finally, having computed a maximum-weight

independent set contained in each particle, by Lemma 13, we can compute the final solution in time polynomial in n . Hence, there are constants c, c_1, c_2 , where $c \gg c_1, c_2$, such that total running time of this step is bounded by:

$$F(n) \leq 2^{c_1 \cdot \sqrt{nt} \log n} \left(n^{c_2} + 4n \cdot 2^{c \cdot \sqrt{tn/2} \log(n/2)} \right) \stackrel{c \gg c_1, c_2}{\leq} 2^{c \cdot \sqrt{tn} \log n}, \quad (1)$$

and so is the total complexity of the algorithm. \blacktriangleleft

4.2 Proof of Theorem 4

Again let us start with explaining the algorithm-specific details of the outline presented at the start of Section 4.

We will use the notion of β -heavy vertices from [10, 11]. Consider a graph G , a weight function $\mathfrak{w} : V(G) \rightarrow \mathbb{Z}_{\geq 0}$, and an independent set $I \subseteq V(G)$. Let $\beta \in (0, 1/2]$ be a real. We say that a vertex $v \in V(G)$ is β -heavy (with respect to I) if $\mathfrak{w}(N[v] \cap I) > \beta \cdot \mathfrak{w}(I)$. A set J is good for I if $J \subseteq I$ and $N[J]$ contains all vertices that are β -heavy with respect to I .

► **Lemma 14** (Chudnovsky et al. [10, 11]). *Let G be an n -vertex graph for $n > 2$, $\mathfrak{w} : V(G) \rightarrow \mathbb{Z}_{\geq 0}$ be a weight function, $I \subseteq V(G)$ be an independent set, and $\beta \in (0, 1/2]$ be a real. Then there exists a set J of size at most $\lceil \beta^{-1} \log n \rceil$ which is good for I .*

Now the vertex is heavy if it is β -heavy for some carefully chosen parameter β . This means that a neighborhood of a vertex is “large” if it contains a significant ($\geq \beta$) fraction of the weight of I_{OPT} . In the first phase, we exhaustively guess the set J that is good for a fixed optimum solution I_{OPT} . Note that J is of small size and since $J \subseteq I_{\text{OPT}}$, we know that $N(J)$ contains no vertices from I_{OPT} and thus can be safely removed from the graph.

Since J is good for I_{OPT} , we know that $G - N[J]$ contains no heavy vertices, and for this graph we call Corollary 12. Now, as $N[X]$ is a neighborhood of few non-heavy vertices, we know that the total weight of $I_{\text{OPT}} \cap N[X]$ is small and thus can be sacrificed, as we aim for an approximation.

Proof of Theorem 4. Let $s, t \geq 1$ be constants and let (G, \mathfrak{w}) be an instance of MWIS, where G is $sS_{t,t,t}$ -free and has n vertices. Let $\varepsilon \in (0, 1)$ be fixed. Fix a maximum-weight independent set I_{OPT} in G with respect to \mathfrak{w} . We describe a procedure that finds in G an independent set I of weight at least $(1 - \varepsilon) \cdot \mathfrak{w}(I_{\text{OPT}})$.

Let N be the minimum power of two greater than or equal to the size of our initial instance. Note that $n \leq N < 2n$. The value of N will not change throughout the execution of the algorithm.

The algorithm itself is a recursive procedure. The arguments of each call are a graph G' , which is an induced subgraph of G , the weight function on $V(G')$ obtained by restricting the domain of \mathfrak{w} , and an integer h , which can be intuitively understood as the depth of the current call in the recursion tree. Since it does not lead to confusion, we will always denote the weight function by \mathfrak{w} . We will keep the invariant that for each call (G', \mathfrak{w}, h) it holds that $|V(G')| \leq N/2^h$. The initial call, corresponding to the root of the recursion tree, is for $(G, \mathfrak{w}, 0)$.

Consider a call for the instance (G', \mathfrak{w}, h) . If $|V(G')| < n_0$, where n_0 is a constant (depending on s and t) that follows from the reasoning below, then we can solve the problem by brute force. Thus let us assume that $n \geq n_0$. In particular, $N > 1$.

We set

$$\beta(h, \varepsilon) := \frac{\varepsilon}{12(t+1) \log(N/2^h) \cdot ((1-\varepsilon) \log N + \varepsilon(h+1))}. \quad (2)$$

It is straightforward to verify that for $h < \log N$ we have $\beta(h, \varepsilon) \in (0, 1/2]$. On the other hand, if $h \geq \log N$, then G' is of constant size and thus $\beta(h, \varepsilon)$ is not computed for such h .

Let \mathcal{J} be the family of all independent sets in G' of size at most $\lceil \beta(h, \varepsilon)^{-1} \log(N/2^h) \rceil$. For each $J \in \mathcal{J}$ we proceed as follows. If $|V(G' - N[J])| < n_0$, then we compute a maximum-weight independent set I_J in $G' - N[J]$ by brute force. Otherwise, we use Corollary 12, to obtain a set $X_J \subseteq V(G' - N[J])$ and a rigid extended strip decomposition (H, η) of $G' - N[J] - N[X_J]$ such that each particle of (H, η) is of size at most $|V(G' - N[J])|/2$. By Corollary 12, we obtain

$$\begin{aligned} |X_J| &\leq (s-1)(3t+1) + (11 \log |V(G' - N[J])| + 6)(t+1) \\ &\leq 12(t+1) \log |V(G')| \leq 12(t+1) \log(N/2^h). \end{aligned} \quad (3)$$

Let $Y_J := N(J) \cup N[X_J]$. We modify (H, η) into an extended strip decomposition of $G' - Y_J$ as follows. For each $v \in J$, we add to H an isolated vertex x_v , and set $\eta(x_v) = \{v\}$.¹ Let us call this extended strip decomposition (H', η') . Observe that each particle of (H', η') is of size at most $|V(G' - N[J])|/2 \leq |V(G')|/2$. Furthermore, since (H, η) is rigid, so is (H', η') .

For each nonempty particle A of (H', η') we call the algorithm recursively on an instance $(G'[A], \mathfrak{w}, h+1)$. Let $I(A)$ be the value returned by the algorithm. For each empty particle A we set $I(A) := \emptyset$. Finally, we apply the algorithm from Lemma 13, in order to obtain an independent set I_J of $G' - Y_J$ and thus of G' . Recall that the value of ς is not needed to apply Lemma 13; we will define it in the next paragraph when we discuss the approximation guarantee. As the solution, we return the set I_J of maximum weight, over all choices of $J \in \mathcal{J}$.

Approximation guarantee. Consider the recursion tree of our algorithm. We mark some nodes of the recursion tree. First, we mark the root. Now consider some marked node z corresponding to a call (G', \mathfrak{w}, h) , such that z is not a leaf node. Observe that by Lemma 14, there is some $J \in \mathcal{J}$ (for this particular instance) which is good for $I_{\text{OPT}} \cap V(G')$. Fix such J . If there is more than one, we choose one arbitrarily. We mark the children of z that correspond to the calls on the particles of the extended strip decomposition of $G' - Y_J$.

Let \mathcal{T} be the subtree of the recursion tree induced by the marked nodes. Note that each leaf of \mathcal{T} is a leaf of the whole recursion tree, i.e., it corresponds to an instance of constant size. Since at each level of the recursion, the size of the instance drops by at least half, we observe that each instance at level h (where the root is at level 0) is of size at most $N/2^h$. Consequently, the depth of \mathcal{T} is at most $\log N$.

Consider a call for an instance (G', \mathfrak{w}, h) and let J be good for I_{OPT} . Let us estimate $\mathfrak{w}(I_{\text{OPT}} \cap Y_J)$. First, observe that since $J \subseteq I_{\text{OPT}}$, we have that $\mathfrak{w}(I_{\text{OPT}} \cap N(J)) = 0$. Moreover, since J was chosen to be good, there are no $\beta(h, \varepsilon)$ -heavy vertices in $V(G' - N[J])$, and in particular, in $N[X_J]$. Hence,

$$\begin{aligned} \mathfrak{w}(I_{\text{OPT}} \cap Y_J) &= \mathfrak{w}(I_{\text{OPT}} \cap N[X_J]) \leq |X_J| \cdot \beta(h, \varepsilon) \cdot \mathfrak{w}(I_{\text{OPT}} \cap V(G')) \\ &\stackrel{(2) \text{ and } (3)}{\leq} \frac{\varepsilon}{(1-\varepsilon) \log N + \varepsilon(h+1)} \cdot \mathfrak{w}(I_{\text{OPT}} \cap V(G')). \end{aligned} \quad (4)$$

¹ Another possible way of dealing with the set J would be to add it directly in the computed solution. However, we decided to restore J to the graph, so that these vertices are handled by Lemma 13 and do not require any special treatment.

The following claim shows that the solution computed for the instance (G', \mathfrak{w}, h) at each node of \mathcal{T} is a reasonable approximation of $I_{\text{OPT}} \cap V(G')$.

▷ **Claim 15.** Let z be a node of \mathcal{T} , and let (G', \mathfrak{w}, h) be the instance corresponding to z . Let I be the independent set returned by the algorithm for the call at z . Then $\mathfrak{w}(I) \geq \left(1 - \varepsilon + \frac{\varepsilon h}{\log N}\right) \cdot \mathfrak{w}(I_{\text{OPT}} \cap V(G'))$.

Proof. First, observe that if z is a leaf of \mathcal{T} , then the statement of the claim is satisfied. Indeed, in this case I is computed by brute force, and hence $\mathfrak{w}(I) = \mathfrak{w}(I_{\text{OPT}} \cap V(G'))$.

Recall that the algorithm returns the solution of maximum weight among all choices of $J \in \mathcal{J}$, so clearly we have $\mathfrak{w}(I) \geq \mathfrak{w}(I_J)$, where J is good for $I_{\text{OPT}} \cap V(G')$.

We proceed by induction on h . First, consider a node z at the level $h = \log N$. As the depth of \mathcal{T} is at most $\log N$, we observe that z must be a leaf, so the claim follows by the observation above.

Assume that the claim holds for $h + 1 \in [\log N]$ and consider a node z at level h . If z is a leaf, then again, we are done. Otherwise, let \mathcal{A} be the set of nonempty particles of the extended strip decomposition of $G' - Y_J$. For every such particle A , we recursively computed an independent set $I(A)$. By the inductive assumption, we have that $\mathfrak{w}(I(A)) \geq \left(1 - \varepsilon + \frac{\varepsilon(h+1)}{\log N}\right) \mathfrak{w}(I_{\text{OPT}} \cap V(G'[A]))$; note that these recursive calls are at level $h + 1$. Clearly, the same holds for empty particles because \emptyset is there an optimum solution.

Thus, by Lemma 13 applied to I_{OPT} and $\varsigma = 1 - \varepsilon + \frac{\varepsilon(h+1)}{\log N}$, we obtain an independent set I_J in $G' - Y_J$, such that

$$\begin{aligned} \mathfrak{w}(I_J) &\geq \left(1 - \varepsilon + \frac{\varepsilon(h+1)}{\log N}\right) \mathfrak{w}(I_{\text{OPT}} \cap V(G' - Y_J)) \\ &= \left(1 - \varepsilon + \frac{\varepsilon(h+1)}{\log N}\right) \left(\mathfrak{w}(I_{\text{OPT}} \cap V(G')) - \mathfrak{w}(I_{\text{OPT}} \cap Y_J)\right). \end{aligned} \quad (5)$$

Combining (5) with (4) and simplifying the formula, we obtain

$$\mathfrak{w}(I_J) \geq \left(1 - \varepsilon + \frac{\varepsilon h}{\log N}\right) \mathfrak{w}(I_{\text{OPT}} \cap V(G')),$$

which concludes the proof of the claim. \triangleleft

Since the root of the recursion tree belongs to \mathcal{T} , the final result I returned for the call at the root (i.e., for $(G, \mathfrak{w}, 0)$) satisfies

$$\mathfrak{w}(I) \geq (1 - \varepsilon) \cdot \mathfrak{w}(I_{\text{OPT}} \cap V(G)) = (1 - \varepsilon) \cdot \mathfrak{w}(I_{\text{OPT}}).$$

This concludes the discussion of the approximation guarantee.

Running time. Recall that the recursion tree has depth at most $\log N$. Let us show the following claim concerning the running time.

▷ **Claim 16.** Let z be a node of the recursion tree, and let (G', \mathfrak{w}, h) be the instance corresponding to z . Then the algorithm solves this instance in time $2^{\mathcal{O}(\varepsilon^{-1} \log^4 N \log(N/2^{h-1}))}$.

Proof. Let $F(h)$ denote the upper bound for the running time of our algorithm, depending on the level of the call in the recursion tree. We aim to show that there is an absolute constant c , such that for N sufficiently large we have

$$F(h) \leq 2^{c \cdot \varepsilon^{-1} \log^4 N \log(N/2^{h-1})}.$$

Recall that $|V(G')| \leq N/2^h$. If z is a leaf, then the instance is of constant size, and thus the claim holds (assuming that c is sufficiently large). In particular this happens if $h = \log N$. So let us assume that the claim holds for the calls at level $h + 1$ and that $h < \log N$.

Recall that we first enumerate the family \mathcal{J} of all independent sets of size at most $\lceil \beta(h, \varepsilon)^{-1} \log(N/2^h) \rceil$. Observe that

$$|\mathcal{J}| \leq |V(G')|^{\lceil \beta(h, \varepsilon)^{-1} \log(N/2^h) \rceil} \leq 2^{\log(N/2^h) \lceil \beta(h, \varepsilon)^{-1} \log(N/2^h) \rceil},$$

and the family \mathcal{J} can be enumerated in time polynomial in its size.

For each $J \in \mathcal{J}$, using Corollary 12 and modifying its outcome, in polynomial time we obtain a set X_J and a rigid extended strip decomposition (H', η') of $G - Y_J$, where $Y_J = N[X_J] \cup N(J)$.

Next, we call the algorithm recursively for at most $4 \cdot |V(G')| \leq 4 \cdot N/2^h$ instances, each at depth $h + 1$. Finally, use Lemma 13 to obtain our solution in time polynomial in $|V(G')|$ and thus in $N/2^h$.

Thus the running time is bounded by the following expression (here c_1, c_2, c_3 are absolute constants, such that c_1 and c_2 are much smaller than c_3 , and $c_3 = c/12(t + 1)$):

$$\begin{aligned} F(h) &\leq 2^{c_1 \cdot \beta(h, \varepsilon)^{-1} \log^2(N/2^h)} \cdot ((N/2^h)^{c_2} + 4 \cdot (N/2^h) \cdot F(h + 1)) \\ &\stackrel{c_3 \gg c_1, c_2}{\leq} 2^{c_3 \cdot \beta(h, \varepsilon)^{-1} \log^2(N/2^h)} \cdot 2^{c \cdot \varepsilon^{-1} \log^4 N \log(N/2^h)} \\ &= \exp \left\{ c_3 \cdot \beta(h, \varepsilon)^{-1} \log^2(N/2^h) + c \cdot \varepsilon^{-1} \log^4 N \log(N/2^h) \right\} \\ &\leq \exp \left\{ c_3 \cdot 12(t + 1) \cdot \left(\frac{1 - \varepsilon}{\varepsilon} \log N + (h + 1) \right) \log^3(N/2^h) + c \cdot \varepsilon^{-1} \log^4 N \log(N/2^h) \right\} \\ &\stackrel{h < \log N}{\leq} \exp \left\{ c \cdot \varepsilon^{-1} \log^4 N + c \cdot \varepsilon^{-1} \log^4 N \log(N/2^h) \right\} \\ &= \exp \left\{ c \cdot \varepsilon^{-1} \log^4 N (\log(N/2^h) + 1) \right\} = \exp \left\{ c \cdot \varepsilon^{-1} \log^4 N \log(N/2^{h-1}) \right\}. \end{aligned}$$

This completes the proof of the claim. \triangleleft

Now we apply Claim 16 to the initial call $(G, \mathbf{w}, 0)$ and obtain that the overall running time is

$$2^{\mathcal{O}(\varepsilon^{-1} \log^5 N)} = 2^{\mathcal{O}(\varepsilon^{-1} \log^5 n)},$$

as $N < 2n$. This completes the proof. \blacktriangleleft

5 Conclusion

In the QPTAS of Chudnovsky, Pilipczuk, Pilipczuk, and Thomassé [10, 11] it was more convenient to measure the weight of parts of the graph not by the number of vertices, but by the weight of the intersection of the sought solution with the part in question. We observe that we can adapt Theorem 2 to this setting of unknown weight function.

► **Theorem 17.** *Given an n -vertex graph G and an integer t , one can in time $n^{\mathcal{O}(t \log n)}$ either:*

- *output an induced copy of $S_{t,t,t}$ in G , or*
- *output a family \mathcal{F} satisfying the following:*
 1. *every element of \mathcal{F} is a pair of a set \mathcal{P} consisting of at most $11 \log n + 6$ induced paths in G , each of length at most $t + 1$, and an extended strip decomposition of $G - N[\bigcup \mathcal{P}]$;*

2. for every weight function $\mathbf{w} : V(G) \rightarrow \mathbb{Z}_{\geq 0}$ there exists a pair in \mathcal{F} such that every particle in the extended strip decomposition of the pair has weight at most half of the total weight of G ;
3. the size of \mathcal{F} is bounded by $n^{\mathcal{O}(\log n)}$.

Proof sketch. As observed in [10, 11], in G one can identify at most n^2 induced paths such that for every weight function $\mathbf{w} : V(G) \rightarrow \mathbb{Z}_{\geq 0}$, at least one of the identified path is a Gyárfás' path for \mathbf{w} , that is, a path Q such that every connected component of $G - N[Q]$ is of weight at most half of the weight of G . Thus, we can guess the path Q as in the proof in Theorem 2 out of at most n^2 candidates.

Then, in the recursive step in the proof of Theorem 2, instead of choosing the heavy particle to recurse on, we guess which particle is heavy (or that none exists). It is easy to see that any extended strip decomposition in the process will have fewer than n inclusion-wise maximal particles; thus, this gives $n^{\mathcal{O}(\log n)}$ possible outputs to enumerate. ◀

We think the $\log n$ factor in Theorem 2 is an artifact of our technique, and is not necessary. Therefore, we pose the following conjecture.

► **Conjecture 18.** *For every integer $t \geq 1$ there exists a constant $\varepsilon > 0$ and an integer s such that every $S_{t,t,t}$ -free graph G admits a set $P \subseteq V(G)$ of size at most s such that $G - N[P]$ admits a rigid extended strip decomposition whose every particle has at most $(1 - \varepsilon)|V(G)|$ vertices.*

Abrishami, Chudnovsky, Dibek, and Rzażewski [2] very recently announced a polynomial-time algorithm for MWIS in $S_{t,t,t}$ -free graphs of bounded degree. Their argument is quite involved and revisits the proof of the three-in-a-tree theorem [13].

Confirming Conjecture 18 would imply the same result almost immediately, possibly with a better running time. Indeed, one needs to branch on $N[P]$ and recurse on the remainder of every particle of (H, η) . The maximum degree of H is bounded by a function of the maximum degree of G (i.e., is a constant), which ensures that the sum of sizes of all particles is linear in $|V(G)|$. This in turns implies that the total complexity of the algorithm can be bounded by a polynomial function. Note that the same approach using Theorem 2 yields quasipolynomial running time bound.

We see Theorem 2 as the analog of Theorem 1 in the classes of $S_{t,t,t}$ -free graphs: with its help, obtaining a QPTAS or a subexponential algorithm was relatively simple, following the ideas of [7, 10, 11]. We expect it is a first step to get a quasipolynomial-time algorithm for MWIS in $S_{t,t,t}$ -free graphs, similarly as Theorem 1 is an essential ingredient of the algorithms for P_t -free graphs [15, 24]. However, there is a lot of work to be done: the way how [15, 24] measure the progress of the branching algorithm is quite intricate; furthermore, for the class of $C_{>t}$ -free graphs (graphs excluding all cycles of length more than t as induced subgraphs, a proper superclass of P_t -free graphs) while an analog of Theorem 1 is known, the corresponding measure of the progress of the branching algorithm is much more involved [16].

References

- 1 Tara Abrishami, Maria Chudnovsky, Cemil Dibek, and Paweł Rzażewski. Polynomial-time algorithm for maximum independent set in bounded-degree graphs with no long induced claws. *CoRR*, abs/2107.05434, 2021. [arXiv:2107.05434](#).
- 2 Tara Abrishami, Maria Chudnovsky, Cemil Dibek, and Paweł Rzażewski. Polynomial-time algorithm for maximum independent set in bounded-degree graphs with no long induced claws. In Niv Buchbinder Joseph (Seffi) Naor, editor, *Proceedings of the 2022 ACM-SIAM Symposium on Discrete Algorithms, SODA 2022, Virtual Conference, January 9-12, 2022*, pages 1448–1470. SIAM, 2022. doi:10.1137/1.9781611977073.61.

- 3 Tara Abrishami, Maria Chudnovsky, Cemil Dibek, Stéphan Thomassé, Nicolas Trotignon, and Kristina Vušković. Graphs with polynomially many minimal separators. *J. Comb. Theory, Ser. B*, 152:248–280, 2022. doi:10.1016/j.jctb.2021.10.003.
- 4 Tara Abrishami, Maria Chudnovsky, Marcin Pilipczuk, Paweł Rzażewski, and Paul D. Seymour. Induced subgraphs of bounded treewidth and the container method. In Dániel Marx, editor, *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms, SODA 2021, Virtual Conference, January 10 - 13, 2021*, pages 1948–1964. SIAM, 2021. doi:10.1137/1.9781611976465.116.
- 5 Vladimir E. 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.
- 6 Vladimir E. Alekseev. On easy and hard hereditary classes of graphs with respect to the independent set problem. *Discret. Appl. Math.*, 132(1-3):17–26, 2003. doi:10.1016/S0166-218X(03)00387-1.
- 7 Gábor Bacsó, Daniel Lokshtanov, Dániel Marx, Marcin Pilipczuk, Zolt Tuza, and Erik Jan van Leeuwen. Subexponential-time algorithms for Maximum Independent Set in P_t -free and broom-free graphs. *Algorithmica*, 81(2):421–438, 2019. doi:10.1007/s00453-018-0479-5.
- 8 Brenda S. Baker. Approximation algorithms for NP-complete problems on planar graphs. *J. ACM*, 41(1):153–180, 1994. doi:10.1145/174644.174650.
- 9 Vincent Bouchitté and Ioan Todinca. Treewidth and minimum fill-in: Grouping the minimal separators. *SIAM J. Comput.*, 31(1):212–232, 2001. doi:10.1137/S0097539799359683.
- 10 Maria Chudnovsky, Marcin Pilipczuk, Michał Pilipczuk, and Stéphan Thomassé. Quasi-polynomial time approximation schemes for the Maximum Weight Independent Set Problem in H -free graphs. *CoRR*, abs/1907.04585, 2019. arXiv:1907.04585.
- 11 Maria Chudnovsky, Marcin Pilipczuk, Michał Pilipczuk, and Stéphan Thomassé. Quasi-polynomial time approximation schemes for the Maximum Weight Independent Set Problem in H -free graphs. In Shuchi Chawla, editor, *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms, SODA 2020, Salt Lake City, UT, USA, January 5-8, 2020*, pages 2260–2278. SIAM, 2020. doi:10.1137/1.9781611975994.139.
- 12 Maria Chudnovsky and Paul D. Seymour. The structure of claw-free graphs. In Bridget S. Webb, editor, *Surveys in Combinatorics, 2005 [invited lectures from the Twentieth British Combinatorial Conference, Durham, UK, July 2005]*, volume 327 of *London Mathematical Society Lecture Note Series*, pages 153–171. Cambridge University Press, 2005. doi:10.1017/cbo9780511734885.008.
- 13 Maria Chudnovsky and Paul D. Seymour. The three-in-a-tree problem. *Comb.*, 30(4):387–417, 2010. doi:10.1007/s00493-010-2334-4.
- 14 Jack Edmonds. Paths, trees, and flowers. *Canadian Journal of Mathematics*, 17:449–467, 1965. doi:10.4153/CJM-1965-045-4.
- 15 Peter Gartland and Daniel Lokshtanov. Independent set on P_k -free graphs in quasi-polynomial time. In Sandy Irani, editor, *61st IEEE Annual Symposium on Foundations of Computer Science, FOCS 2020, Durham, NC, USA, November 16-19, 2020*, pages 613–624. IEEE, 2020. doi:10.1109/FOCS46700.2020.00063.
- 16 Peter Gartland, Daniel Lokshtanov, Marcin Pilipczuk, Michał Pilipczuk, and Paweł Rzażewski. Finding large induced sparse subgraphs in $C_{>t}$ -free graphs in quasipolynomial time. In Samir Khuller and Virginia Vassilevska Williams, editors, *STOC '21: 53rd Annual ACM SIGACT Symposium on Theory of Computing, Virtual Event, Italy, June 21-25, 2021*, pages 330–341. ACM, 2021. doi:10.1145/3406325.3451034.
- 17 Andrzej Grzesik, Tereza Klimošová, Marcin Pilipczuk, and Michał Pilipczuk. Polynomial-time algorithm for Maximum Weight Independent Set on P_6 -free graphs. In Timothy M. Chan, editor, *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019, San Diego, California, USA, January 6-9, 2019*, pages 1257–1271. SIAM, 2019. doi:10.1137/1.9781611975482.77.

- 18 András Gyárfás. On Ramsey covering-numbers. In *Infinite and finite sets (Colloq., Keszthely, 1973; dedicated to P. Erdős on his 60th birthday), Vol. II*, number 10 in Colloq. Math. Soc. Janos Bolyai, pages 801–816. North-Holland, Amsterdam, 1975.
- 19 András Gyárfás. Problems from the world surrounding perfect graphs. In *Proceedings of the International Conference on Combinatorial Analysis and its Applications, (Pokrzywna, 1985)*, number 19 in Zastos. Mat., pages 413–441, 1987. doi:10.4064/am-19-3-4-413-441.
- 20 Johan Håstad. Clique is hard to approximate within $n^{1-\varepsilon}$. *Acta Math.*, 182(1):105–142, 1999. doi:10.1007/BF02392825.
- 21 Daniel Lokshantov, Martin Vatshelle, and Yngve Villanger. Independent set in P_5 -free graphs in polynomial time. In Chandra Chekuri, editor, *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2014, Portland, Oregon, USA, January 5-7, 2014*, pages 570–581. SIAM, 2014. doi:10.1137/1.9781611973402.43.
- 22 Konrad Majewski, Tomáš Masařík, Jana Novotná, Karolina Okrasa, Marcin Pilipczuk, Paweł Rzażewski, and Marek Sokołowski. Max weight independent set in graphs with no long claws: An analog of the gyárfás’ path argument, 2022. doi:10.48550/ARXIV.2203.04836.
- 23 George J. Minty. On maximal independent sets of vertices in claw-free graphs. *Journal of Combinatorial Theory, Series B*, 28(3):284–304, 1980. doi:10.1016/0095-8956(80)90074-X.
- 24 Marcin Pilipczuk, Michał Pilipczuk, and Paweł Rzażewski. Quasi-polynomial-time algorithm for independent set in P_t -free graphs via shrinking the space of induced paths. In Hung Viet Le and Valerie King, editors, *4th Symposium on Simplicity in Algorithms, SOSA 2021, Virtual Conference, January 11-12, 2021*, pages 204–209. SIAM, 2021. doi:10.1137/1.9781611976496.23.
- 25 Najiba Sbihi. Algorithme de recherche d’un stable de cardinalité maximum dans un graphe sans étoile. *Discrete Mathematics*, 29(1):53–76, 1980. doi:10.1016/0012-365X(90)90287-R.
- 26 David Zuckerman. Linear degree extractors and the inapproximability of Max Clique and Chromatic Number. *Theory of Computing*, 3(1):103–128, 2007. doi:10.4086/toc.2007.v003a006.