

Quasi-polynomial time approximation schemes for the Maximum Weight Independent Set Problem in H -free graphs*

Maria Chudnovsky[†]

Princeton University, Princeton, NJ 08544

Marcin Pilipczuk[‡]

Institute of Informatics, University of Warsaw
Banacha 2, 02-097 Warsaw, Poland

Michał Pilipczuk[§]

Institute of Informatics, University of Warsaw
Banacha 2, 02-097 Warsaw, Poland

Stéphan Thomassé

Institut Universitaire de France

Laboratoire d'Informatique du Parallélisme, UMR 5668 ENS Lyon,
CNRS, UCBL, INRIA, Université de Lyon, France

Abstract

In the MAXIMUM INDEPENDENT SET problem we are asked to find a set of pairwise nonadjacent vertices in a given graph with the maximum possible cardinality. In general graphs, this classical problem is known to be NP-hard and hard to approximate within a factor of $n^{1-\varepsilon}$ for any $\varepsilon > 0$. Due to this, investigating the complexity of MAXIMUM INDEPENDENT SET in various graph classes in hope of finding better tractability results is an active research direction.

In H -free graphs, that is, graphs not containing a fixed graph H as an induced subgraph, the problem is known to remain NP-hard and APX-hard whenever H contains a cycle, a vertex of degree at least four, or two vertices of degree at least three in one connected component. For the remaining cases, where every component of H is a path or a subdivided claw, the complexity of MAXIMUM INDEPENDENT SET remains widely open, with only a handful of polynomial-time solvability results for small graphs H such as P_5 , P_6 , the claw, or the fork.

We prove that for every such “possibly tractable” graph H there exists an algorithm that, given an H -free graph G and an accuracy parameter $\varepsilon > 0$, finds an independent set in G of cardinality within a factor of $(1 - \varepsilon)$ of the optimum in time exponential in a polynomial of $\log |V(G)|$ and ε^{-1} . Furthermore, an independent set of maximum size can be found in subexponential time $2^{\mathcal{O}(|V(G)|^{8/9} \log |V(G)|)}$. That is, we show that for every graph H for which MAXIMUM INDEPENDENT SET is not known to be APX-hard and SUBEXP-hard in H -free graphs, the problem admits a quasi-polynomial time approximation scheme and a subexponential-time exact

*An extended abstract of this work appeared at SODA 2020 [11]. The SODA 2020 version contained only approximation schemes, but not the subexponential algorithms presented here.

[†]Supported by NSF grants DMS-1763817. This material is based upon work supported in part by the U. S. Army Research Office under grant number W911NF-16-1-0404.

[‡]This research is a part of a project that has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme Grant Agreement no. 714704.

[§]This research is a part of a project that has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme Grant Agreement no. 677651.

algorithm in this graph class. Our algorithms work also in the more general weighted setting, where the input graph is supplied with a weight function on vertices and we are maximizing the total weight of an independent set.

1 Introduction

For an undirected graph G , a vertex subset $X \subseteq V(G)$ is *independent* if no two vertices of X are adjacent. The size of the largest independent set in a graph, often denoted as $\alpha(G)$, is one of the fundamental graph parameters studied in graph theory. Therefore, it is natural to study the computational task of computing $\alpha(G)$, given G , which we call the MAXIMUM INDEPENDENT SET problem (MIS). In the weighted generalization, MAXIMUM WEIGHT INDEPENDENT SET (MWIS), the given graph G is supplied with a weight function $\mathbf{w}: V(G) \rightarrow \mathbb{N}$, and we ask for an independent set X in G with the maximum possible total weight $\mathbf{w}(X) = \sum_{x \in X} \mathbf{w}(x)$. MIS is a classic problem that is known not only to be NP-hard, but also hard to approximate within a factor of $n^{1-\varepsilon}$ for every $\varepsilon > 0$, unless $P = NP$ [27, 42].

In light of these lower bounds, a lot of effort has been put into understanding the complexity of MIS and MWIS in restricted graph classes. While the celebrated Baker’s technique yields a polynomial-time approximation scheme (PTAS) for MWIS in planar graphs [5], MIS remains NP-hard in planar graphs of degree at most three and APX-hard in graphs of maximum degree at most three [17, 18, 25]. To extend these lower bounds to other graph classes, the following observation due to Poljak [40] is very useful: if G' is created from G by subdividing one edge twice, then $\alpha(G') = \alpha(G) + 1$. Thus, if we fix any graph H that contains either a cycle, a vertex of degree at least four, or two vertices of degree three in one connected component, then starting from a graph G of maximum degree at most three and subdividing each edge a sufficient number of times, we obtain a graph G' where computing $\alpha(\cdot)$ is equally hard, while G does not contain an induced subgraph isomorphic to H . Now, MIS is known to be APX-hard in graphs of maximum degree at most three, and in this case $\alpha(G)$ is linear in the size of the graph. Moreover, under the Exponential-Time Hypothesis, MIS has no *subexponential-time algorithm* (that is, one with running time $2^{o(n)}$) on graphs of maximum degree at most three; we call this property *SUBEXP-hardness*. This implies that MIS remains APX-hard and SUBEXP-hard in \mathcal{H} -free graphs for every finite family of graphs \mathcal{H} such that every $H \in \mathcal{H}$ is not a disjoint union of paths and subdivided claws.¹

However, when H is a disjoint union of paths and subdivided claws, no hardness result on the complexity of MIS nor MWIS on H -free graphs is known. In fact, it would be consistent with our knowledge if MWIS turns out to be polynomial-time solvable in H -free graphs for all such graphs H . Currently, such a result seems beyond our reach. Let P_t be the path on t vertices and the *claw* be the four-vertex tree with one vertex of degree three and three leaves. The class of P_4 -free graphs (known also as *cographs*) have a very rigid structure (in particular, they have clique-width at most 2), and hence they admit a simple polynomial-time algorithm for MWIS [14]. Claw-free graphs also possess very strong structural properties and inherit many properties of their main subclass: line graphs. In particular, the augmenting-path algorithm for maximum cardinality matching generalizes to a polynomial-time algorithm for MWIS in claw-free graphs [33, 35, 41]. A more modern approach based on the decomposition theorems for claw-free graphs yields a different line of algorithms [15, 16, 36–38]. This, in turn, can be generalized to so-called fork-free graphs [29], where the fork is constructed from the claw by subdividing one edge once. The case of P_5 -free graphs, after being open for a long time,

¹A graph is *H-free* if it does not contain an induced subgraph isomorphic to H . A graph G is \mathcal{H} -free if G is H -free for every $H \in \mathcal{H}$. A *subdivided claw* is a tree with one vertex of degree three and all other vertices of degree at most two.

was resolved positively in 2014 by Lokshtanov, Vatschelle, and Villanger [28] using the framework of *potential maximal cliques*. With a substantially larger technical effort, their approach has been generalized to P_6 -free graphs by Grzesik et al. [24]. The polynomial-time solvability of MWIS on P_7 -free graphs, or T -free graphs where T is any subdivision of the claw other than the fork, remains open. There is a significant body of work concerning the complexity of MWIS in various subclasses of P_t -free or T -free graphs, see e.g. [6, 22, 30, 31, 34].

In 2017, evidence in favor of the tractability of MIS and MWIS at least in P_t -free graphs has been found: there is a subexponential-time algorithm for the problem running in time $2^{\mathcal{O}(\sqrt{nt \log n})}$ on an n -vertex P_t -free graph [4, 7, 23]. The main insight is that the classical Gyárfás' path argument, originally used to show that P_t -free graphs are χ -bounded [26], implies that a P_t -free graph G admits a balanced separator consisting of at most $t - 1$ vertex neighborhoods. Here, a balanced separator is a set of vertices whose removal results in a graph where every connected component has at most $|V(G)|/2$ vertices.

Our results. We provide a new evidence in favor of the tractability of MWIS in all cases of H -free graphs where it is not known to be APX-hard.

Theorem 1.1. *For every graph H whose every connected component is a path or a subdivided claw, there exists an algorithm that, given an H -free graph G with a weight function $\mathbf{w} : V(G) \rightarrow \mathbb{N}$ and an accuracy parameter $\varepsilon > 0$, computes a $(1 - \varepsilon)$ -approximation to MAXIMUM WEIGHT INDEPENDENT SET on (G, \mathbf{w}) in time exponential in a polynomial of $\log |V(G)|$ and ε^{-1} .*

Theorem 1.2. *For every graph H whose every connected component is a path or a subdivided claw, there exists an algorithm that, given an H -free graph G with a weight function $\mathbf{w} : V(G) \rightarrow \mathbb{N}$, solves MAXIMUM WEIGHT INDEPENDENT SET on (G, \mathbf{w}) in time exponential in $\mathcal{O}(|V(G)|^{8/9} \log |V(G)|)$.*

That is, in all the cases when MWIS is not known to be APX-hard or SUBEXP-hard on H -free graphs, we prove that MWIS admits a quasi-polynomial time approximation scheme (QPTAS) and a subexponential-time algorithm.

We remark here that Theorems 1.1 and 1.2 treat H as a constant-sized graph. That is, the polynomial (of $\log |V(G)|$ and ε^{-1}) and the constant factor (hidden in the big- \mathcal{O} notation) in the exponents of the time bounds depend on the graph H . If one follows closely the arguments, the final bound of the running time of the approximation scheme is of the form $\exp(c_H \varepsilon^{-c} \log^c n)$ for some constant c_H depending on H and a universal (independent of H) constant c . Since the constant c is significantly larger than 1, we refrain from precisely analysing this running time bound for the sake of simplicity.

For an insight into the techniques standing behind Theorems 1.1 and 1.2, let us first focus on the case $H = P_t$. A subexponential-time algorithm for this case has been already provided in [4]. For an approximation scheme, let (G, \mathbf{w}) be an input to MWIS with G being P_t -free and let $\varepsilon > 0$ be an accuracy parameter. Let $X \subseteq V(G)$ be an independent set in G of maximum possible weight. Fix a threshold $\beta := \varepsilon^{-1} t \log n$ and say that a vertex $v \in V(G)$ is X -heavy if it contains at least a β^{-1} fraction of the weight of X in its closed neighborhood, that is, $\mathbf{w}(X \cap N[v]) \geq \beta^{-1} \mathbf{w}(X)$. A simple coupon-collecting argument shows that there is a set $Y \subseteq X$ of size $\mathcal{O}(\beta \log n)$ such that all X -heavy vertices are contained in $N[Y]$. We investigate all the $n^{\mathcal{O}(\beta \log n)} = 2^{\mathcal{O}(\varepsilon^{-1} \log^3 n)}$ subcases corresponding to the possible choices of Y . Having fixed Y in a subcase, we can delete $N(Y)$ from the graph and from now on assume that there are no more X -heavy vertices (except for isolated vertices that are easy to deal with).

Now the Gyárfás path argument, like e.g. in [23], asserts that in G there exists a balanced separator $A = N[B]$ for some $|B| \leq t - 1$. We simply delete A from the graph and restart the

whole algorithm on every connected component of G . Since there are no X -heavy vertices, we lose only a fraction of $\beta^{-1}t < \varepsilon/\log n$ of the weight of X in this step. Since every connected component of $G - A$ is of size at most $n/2$, the depth of the recursion is at most $\log n$. Consequently, throughout the recursion the total loss in the weight of the optimum solution X is at most $\varepsilon \cdot \mathbf{w}(X)$. Furthermore, it can be easily seen that the whole recursion tree has size bounded by $2^{\mathcal{O}(\varepsilon^{-1} \log^4 n)}$, giving a quasi-polynomial running time bound of the whole algorithm.

To generalize this argument (and the argument for subexponential-time algorithm of [4]) to the case of H being a subdivided claw, an additional ingredient is needed: the *Three-in-a-Tree* Theorem by Chudnovsky and Seymour [13]. Let G be a graph and let $x, y, z \in V(G)$ be three distinguished vertices. The Three-in-a-Tree Theorem provides a dichotomy: either we can find an induced tree in G that contains x, y , and z , or we can find a suitable decomposition of G that somehow “separates” x, y, z and witnesses that no such tree exists; this decomposition has a similar flavor to the decomposition for claw-free graphs [12]. By carefully combining this result with the Gyárfás path argument, we show that in an H -free graph G one can either find a balanced separator containing a small fraction of the weight of the optimum solution (in the case of a QPTAS) or of small size (in the case of a subexponential-time algorithm), e.g., consisting of a constant number of vertex neighborhoods, or a decomposition coming from the Three-in-a-Tree Theorem where every part is of significantly smaller size. Such a decomposition allows us to recurse on every part independently and then assemble the final result from partial results using a reduction to the maximum weight matching problem.

Having obtained the statements of Theorems 1.1 and 1.2 for H being a path or a subdivided claw, we can generalize it to H being a disjoint union of such graphs in a relatively simple and standard way.

In light of Theorems 1.1 and 1.2, we conjecture the following generalization.

Conjecture 1.3. *For every forest H of maximum degree at most three, MWIS admits a QPTAS and a subexponential-time algorithm in the class of graphs that do not contain any subdivision of H as an induced subgraph.*

Our techniques stop short of proving Conjecture 1.3: we are able to prove it for H containing at most three vertices of degree three. Note that this strictly generalizes the conclusions of Theorems 1.1 and 1.2 for H being a subdivided claw (with a $\mathcal{O}(|V(G)|^{40/41} \log |V(G)|)$ term in the exponent of the running time bound of the subexponential-time algorithm).

Furthermore, as a side result we obtain a QPTAS and a subexponential-time algorithm for graphs excluding a long hole.

Theorem 1.4. *For every $t \geq 4$ there exists an algorithm that, given a graph G that does not contain any cycle of length at least t as an induced subgraph, a weight function $\mathbf{w} : V(G) \rightarrow \mathbb{N}$, and an accuracy parameter $\varepsilon > 0$, computes a $(1 - \varepsilon)$ -approximation to MAXIMUM WEIGHT INDEPENDENT SET on (G, \mathbf{w}) in time exponential in a polynomial of $\log |V(G)|$ and ε^{-1} . Furthermore, in the same graph class MAXIMUM WEIGHT INDEPENDENT SET can be solved exactly in time exponential in $\mathcal{O}(|V(G)|^{1/2} \log |V(G)|)$.*

The techniques of Theorem 1.4 allow us also to state the following graph-theoretical corollary that generalizes an analogous result for P_t -free graphs [4, 23] and for graphs excluding any induced cycle of length at least 5 [10].

Theorem 1.5. *For every $t \geq 4$ there exists a constant c_t such that every graph G that does not contain any cycle of length at least t as an induced subgraph has treewidth bounded by $c_t \Delta$, where Δ is the maximum degree of G .*

Recent progress. This work, announced in 2019 [9,11], initiated a sequence of rapid developments.

In the case of P_t -free graphs, Gartland and Lokshantov showed in 2020 a quasipolynomial-time exact algorithm [19] for MWIS; the result was later simplified by Pilipczuk, Pilipczuk and Rzażewski [39]. The authors of the aforementioned two papers joined forces and generalized this result to the class of graphs that does not contain any cycle of length at least t as an induced subgraph and to a wide family of problems that include MWIS and FEEDBACK VERTEX SET [21].

Meanwhile, in the regime of polynomial-time algorithms, Abrishami et al [2] and Chudnovsky et al [8] improved the framework of [24,28], generalizing the algorithm for MWIS in P_6 -free graphs to the same set of problems as in [21] and some other related graph classes.

For the general case of H -free graphs where every component of H is a path or a subdivided claw, in 2021 Abrishami et al [1] showed a polynomial-time algorithm with the additional assumption of bounded maximum degree. Then, in 2022 Majewski et al [32] showed a structural result being the analog of the Gyárfás’ path argument, providing an arguably simpler proof of Theorems 1.1 and 1.2 with better running time bounds. Building on their work and the algorithms for P_t -free graphs [19,39], Gartland et al [20] very recently announced a quasipolynomial-time exact algorithm for MWIS. Furthermore, Abrishami et al [3] showed how to use the structural result of [32] to simplify the arguments of [1] and generalize the assumption of bounded maximum degree to the assumption of being K_s -free and $K_{s,s}$ -free for a constant s .

In all aforementioned works for P_t -free graphs, the Gyárfás path argument is essential, while in all aforementioned works for graphs excluding subdivided claws, the usage of extended strip decompositions and the three-in-a-tree theorem outlined in this work plays pivotal role. It remains at the moment highly unclear how to combine these tools with the framework of *potential maximal cliques*, being the base of the known polynomial-time algorithms in P_5 - and P_6 -free graphs [2,8,24,28].

Organization. After brief preliminaries in Section 2, we present our framework in Section 3. In Section 4 we treat heavy vertices. As a warm-up, the argument for P_t -free graphs is described in Section 5; this section also contains proofs of Theorems 1.4 and 1.5. Section 6, the main technical part of the paper, considers the case of H -free graphs where H is a subdivided claw, with Theorems 1.1 and 1.2 inferred in Section 6.2. Finally, in Section 7 we prove Conjecture 1.3 for H being a forest with at most three vertices of degree three.

2 Preliminaries

For an (undirected, simple) graph G and a vertex $v \in V(G)$, $N(v)$ denotes the (open) neighborhood of v , and $N[v] = \{v\} \cup N(v)$ is the closed neighborhood of v . We extend it to sets of vertices $X \subseteq V(G)$ by $N[X] = \bigcup_{v \in X} N[v]$ and $N(X) = N[X] \setminus X$. Whenever the graph G is not clear from the context, we clarify it by putting it in the subscript. For brevity, we sometimes identify subgraphs with their vertex set when this does not create any confusion: if D is a subgraph of G , then $G - D$, $N[D]$, and $N(D)$ are shorthands for $G - V(D)$, $N[V(D)]$, and $N(V(D))$, respectively. Two disjoint sets $A, B \subseteq V(G)$ are *fully adjacent* (*fully anti-adjacent*) if there are all possible edges (no edges, respectively) between A and B . By P_t we denote a path on t vertices. For a graph G , $cc(G)$ is the family of connected components of G .

2.1 Maximum Weight Independent Set

Let G be a graph and let $\mathbf{w} : V(G) \rightarrow \mathbb{N}$ be a weight function. For a set $X \subseteq V(G)$ we denote $\mathbf{w}(X) = \sum_{x \in X} \mathbf{w}(x)$. The MAXIMUM WEIGHT INDEPENDENT SET (MWIS) problem asks for an

independent set $I \subseteq V(G)$ maximizing $\mathbf{w}(I)$. We say that an independent set I is an α -*approximation* for MWIS in (G, \mathbf{w}) if for every independent set I' in G we have $\mathbf{w}(I) \geq \alpha \cdot \mathbf{w}(I')$. In this work, given G, \mathbf{w} , and an accuracy parameter $\varepsilon > 0$, we ask for an independent set I that is a $(1 - \varepsilon)$ -approximation. For simplicity, we will develop an algorithm that gives only a $(1 - c \cdot \varepsilon)$ -approximation for some universal constant c , as we can then use it with rescaled value of ε . We denote $n = |V(G)|$.

2.2 Extended strip decomposition and the three-in-a-tree theorem

Let G be a graph. An *extended strip decomposition* of G consists of the following:

1. a simple non-empty graph H ,
2. a vertex set $\eta(e) \subseteq V(G)$ for every $uv = e \in E(H)$ and subsets $\eta(e, u), \eta(e, v) \subseteq \eta(e)$,
3. a vertex set $\eta(v) \subseteq V(G)$ for every $v \in V(H)$, and
4. a vertex set $\eta(T) \subseteq V(G)$ for every triangle T in H ,

with the following properties:

1. the vertex sets of $\eta(e)$, $\eta(v)$, and $\eta(T)$ form a partition of $V(G)$;
2. for every $v \in V(H)$ and every two distinct edges $vu, vw \in E(H)$ incident with v , the set $\eta(vu, v)$ is fully adjacent to $\eta(vw, v)$ in G ;
3. every edge $xy \in E(G)$ is either contained in one of the graphs $G[\eta(e)]$, $G[\eta(v)]$, $G[\eta(T)]$, or is one of the following types:
 - $x \in \eta(e, v)$, $y \in \eta(e', v)$ for two distinct edges e, e' of H incident with a common vertex $v \in V(H)$;
 - $x \in \eta(v)$ and $y \in \eta(e, v)$ for some edge $e \in E(H)$ incident with a vertex $v \in V(H)$;
 - $x \in \eta(T)$ and $y \in \eta(e, v) \cap \eta(e, u)$ for some triangle T in H and an edge $e = uv$ of this triangle.

The main result of [13] is the following.

Theorem 2.1 ([13]). *Let G be a connected graph and let $Z \subseteq V(G)$ be a set of size at least two such that for every induced tree T of G , $|V(T) \cap Z| \leq 2$. Then there exists an extended strip decomposition (H, η) of G such that for every $z \in Z$ there exists a distinct vertex $w_z \in V(H)$ of degree one in H with $\eta(e_z, w_z) = \{z\}$ where e_z is the unique edge of H incident with w_z . Furthermore, given G and Z , such a decomposition can be computed in polynomial time.*

Given a graph G and an extended strip decomposition (H, η) of G , a vertex z satisfying the property expressed in Theorem 2.1 will be called *peripheral* in (H, η) . Concretely, z is peripheral in (H, η) if there exists a vertex w_z of H , said to be *occupied* by z , such that w_z has degree 1 in H and satisfies $\eta(e_z, w_z) = \{z\}$, where e_z is the unique edge incident to w_z in H .

We will also need the notion of a *trivial* extended strip decomposition. Given a graph G , a *trivial extended strip decomposition* (H, η) consists of an edgeless graph H that has a vertex x_C for every connected component C of G and $\eta(x_C) = C$.

3 Disperser yields a QPTAS

Let G be a graph and let (H, η) be an extended strip decomposition of G . For an edge $e \in E(H)$, let $\mathcal{T}(e)$ be the set of all triangles of H that contain e . We define a number of *atoms* as follows. For

every edge $e = uv \in E(H)$, we define the following atoms:

$$\begin{aligned} A_e^\perp &= \eta(e) \setminus (\eta(e, u) \cup \eta(e, v)), & A_e^u &= (\eta(u) \cup \eta(e)) \setminus \eta(e, v), \\ A_e^v &= (\eta(v) \cup \eta(e)) \setminus \eta(e, u), & A_e^{uv} &= \eta(u) \cup \eta(v) \cup \eta(e) \cup \bigcup_{T \in \mathcal{T}(e)} \eta(T). \end{aligned}$$

Furthermore, we define an atom $A_v = \eta(v)$ for every $v \in V(H)$ and an atom $A_T = \eta(T)$ for every triangle T in H . A *trivial atom* is an atom $A_v = \eta(v)$ for an isolated vertex v of H with A_v being a singleton containing an isolated vertex of G .

Let $\mathbf{w}: V(G) \rightarrow \mathbb{N}$ be a weight function and let $\gamma, \delta > 0$ be reals. Let $X \subseteq V(G)$ and let (H, η) be an extended strip decomposition of $G - X$. We say that $(X, (H, \eta))$ is

- δ -*shrinking* if for every nontrivial atom A of (H, η) we have $\mathbf{w}(A) \leq (1 - \delta)\mathbf{w}(V(G))$;
- γ -*safe* if $\mathbf{w}(X) \leq \gamma \cdot \mathbf{w}(V(G))$ and, furthermore, for every nontrivial atom A of (H, η) it holds that $\mathbf{w}(X) \leq \gamma \cdot \mathbf{w}(V(G) \setminus A)$;
- (γ, δ) -*good* if it is both δ -shrinking and γ -safe.

For a set $I \subseteq V(G)$, a weight function \mathbf{w}_I is defined as $\mathbf{w}_I(v) = \mathbf{w}(v)$ for every $v \in I$ and $\mathbf{w}_I(v) = 0$ for every $v \in V(G) \setminus I$.

For approximation schemes, we need the following notion.

Definition 3.1. For a graph G and a weight function $\mathbf{w}: V(G) \rightarrow \mathbb{N}$, a (γ, δ) -*dispenser* is a family \mathcal{D} such that:

- every member of \mathcal{D} is a pair of the form $(X, (H, \eta))$, where (H, η) is an extended strip decomposition of $G - X$; and
- for every independent set I in G with $\mathbf{w}(I) > 0$ there exists $(X, (H, \eta)) \in \mathcal{D}$ that is (γ, δ) -good for G and \mathbf{w}_I .

If one is interested in subexponential-time algorithms, it suffices to consider the following simpler notion that considers only uniform weights.

Definition 3.2. For a constant $\xi \in (0, 1)$ and a graph G , a ξ -*uniform dispenser* is a pair $(X, (H, \eta))$, where $X \subseteq V(G)$ and (H, η) is an extended strip decomposition of $G - X$ such that

$$|X| \leq |V(G)|^{-\xi} \cdot |V(G) \setminus A| \quad \text{and} \quad |A| \leq |V(G)| - |V(G)|^\xi \quad \text{for every atom } A \text{ of } (H, \eta).$$

3.1 Intuition

The main result of this section is that an algorithm producing dispensers with good parameters yields a QPTAS and, similarly, an algorithm producing uniform dispensers with good parameters yields an exact subexponential-time algorithm. Let us now give some intuition.

Let G be a graph and let (H, η) be an extended strip decomposition of G . Let A_1 and A_2 be two atoms of (H, η) . We say that A_1 and A_2 are *conflicting* if they are potentially not disjoint; that is, for every $e = uv \in E(H)$

- (i) A_e^\perp , A_e^u , A_e^v , and A_e^{uv} are pairwise in conflict;
- (ii) both A_e^u and A_e^{uv} conflict with A_u and both A_e^v and A_e^{uv} conflict with A_v ;
- (iii) A_e^{uv} and A_e^u conflicts with $A_{e'}^{uv'}$ and $A_{e'}^u$, for every edge $e' = uv' \in E(H)$ incident with u , and similarly for the v endpoint; and
- (iv) A_e^{uv} and A_T are in conflict for every $T \in \mathcal{T}(e)$.

Observe that if A_1 and A_2 are not conflicting then not only $A_1 \cap A_2 = \emptyset$ but also $E(A_1, A_2) = \emptyset$. Informally, two atoms A_1 and A_2 are not conflicting if and only if the definition of the extended strip decomposition ensures that they are disjoint and there is no edge of G between A_1 and A_2 . A family \mathcal{A} of atoms of (H, η) is *independent* if every two distinct elements of \mathcal{A} are not conflicting.

For an independent set I in G , we define the following family \mathcal{A}_I of atoms of (H, η) :

- A_e^{uv} for every $e = uv \in E(H)$ with $I \cap \eta(e, u) \neq \emptyset$ and $I \cap \eta(e, v) \neq \emptyset$,
- A_e^u for every $e = uv \in E(H)$ with $I \cap \eta(e, u) \neq \emptyset$ but $I \cap \eta(e, v) = \emptyset$,
- A_e^v for every $e = uv \in E(H)$ with $I \cap \eta(e, v) \neq \emptyset$ but $I \cap \eta(e, u) = \emptyset$,
- A_e^\perp for every $e = uv \in E(H)$ with $I \cap (\eta(e, u) \cup \eta(e, v)) = \emptyset$,
- A_v for every $v \in V(H)$ such that for every e incident with v we have $I \cap \eta(e, v) = \emptyset$,
- A_T for every triangle T in H such that for all edges $e = uv$ of T we have $I \cap \eta(e, u) = \emptyset$ or $I \cap \eta(e, v) = \emptyset$.

Observe that for every $v \in V(H)$, I may intersect at most one set $\eta(e, v)$ for e incident with v . From this, a direct check verifies the following crucial observation:

Claim 3.3. *For every independent set I in G , the family \mathcal{A}_I is independent and $I \subseteq \bigcup \mathcal{A}_I$.*

Proof. We consider the four cases of how the atoms can be conflicting one-by-one. For Case (i), observe that for every $e = uv \in E(H)$, the conditions for $A_e^\perp, A_e^u, A_e^v, A_e^{uv}$ are mutually exclusive and exactly one of these atoms is in \mathcal{A}_I . For Case (ii), by definition $A_v \in \mathcal{A}_I$ only if $A_e^v, A_e^{uv} \notin \mathcal{A}_I$ for every edge $e = uv$ incident with v .

Case (iii) is the most interesting: the definition of the extended strip decomposition ensures that $\eta(e, v)$ and $\eta(e', v)$ are fully adjacent for two different edges e, e' incident with v , and thus for every $v \in V(H)$ the independent set I can contain a vertex of at most one set $\eta(e, v)$ over all edges e incident with v . Consequently, \mathcal{A}_I contains at most one set A_e^{uv} or A_e^v over all edges $e = uv$ incident with v .

Finally, for Case (iv), A_T is conflicting only with atoms A_e^{uv} for edges $e = uv$ of T , but the condition for including A_T into \mathcal{A}_I is a negation of the condition for excluding any A_e^{uv} for edges $e = uv$ of T . ┘

In the other direction, if we are given an independent set $I(A) \subseteq A$ for every atom $A \in \mathcal{A}$ of an independent family \mathcal{A} of atoms, then $\bigcup_{A \in \mathcal{A}} I(A)$ is an independent set in G .

Thus, one can reduce finding a (good approximation of) maximum-weight independent set in G to finding such a (good approximation of) independent set in subgraphs $G[A]$ for atoms $A \in \mathcal{A}_I$, where I is the sought maximum-weight independent set. In the definition of a disperser, if one recurses in the above sense on $G - X$ and (H, η) for every $(X, (H, \eta))$ in the disperser, the notion of δ -shrinking ensures that such recursion is of small depth, while the notion of γ -safety ensures that by sacrificing the set X we lose only a small fraction of the optimum at every recursion step. In uniform dispersers, the bound on the size of X allows us to branch exhaustively on X in the recursion step; this cost is amortized by the decrease in the size of graphs considered in the branches.

However, there is one major obstacle to the above outline: we do not know the family \mathcal{A}_I . Instead, we can recurse on every atom of (H, η) .

Then, we need an observation that assembling results from the recursion in the best possible way reduces to a maximum-weight matching problem in an auxiliary graph, in a similar fashion that finding maximum-weight independent set in line graphs corresponds to finding maximum-weight matching in the preimage graph.

3.2 Formal statements

The following definition encompasses the idea that a graph class admits efficiently computable dispersers.

Definition 3.4. Let $\gamma \in (0, 1/2)$ be a real, $\delta: \mathbb{N} \rightarrow (0, 1/2)$ be a nonincreasing function, and $\mathbf{S}, \mathbf{T}: \mathbb{N} \rightarrow \mathbb{N}$ be nondecreasing functions. A hereditary graph class \mathcal{C} is called $(\gamma, \delta, \mathbf{S}, \mathbf{T})$ -dispersible if there exists an algorithm that, given an n -vertex graph $G \in \mathcal{C}$ and a weight function $\mathbf{w}: V(G) \rightarrow \mathbb{N}$, runs in time $\mathbf{T}(n)$ and computes a $(\gamma, \delta(n))$ -disperser for G and \mathbf{w} of size at most $\mathbf{S}(n)$.

The main theorem concerning approximation schemes is the following.

Theorem 3.5. *Let \mathcal{C} be a hereditary graph class with the following property: For every $\gamma \in (0, 1/2)$ there exist functions $\delta, \mathbf{S}, \mathbf{T}$ where*

$$(\delta(n))^{-1} \in \text{poly}(\log n, \gamma^{-1}) \quad \text{and} \quad \mathbf{S}(n), \mathbf{T}(n) \in 2^{\text{poly}(\log n, \gamma^{-1})}$$

and $\delta(n)$ is computable in polynomial time given γ and n , such that \mathcal{C} is $(\gamma, \delta, \mathbf{S}, \mathbf{T})$ -dispersible. Then MWIS restricted to graphs from \mathcal{C} admits a QPTAS.

From now on, hereditary classes \mathcal{C} satisfying the assumptions of Theorem 3.5 will be called *QP-dispersible*. Thus, Theorem 3.5 states that MWIS admits a QPTAS on every QP-dispersible class, while in the next sections we will prove that several classes are indeed QP-dispersible.

The above definitions are suited for all our results, but in some simpler cases we will construct dispersers that have a simpler form. More precisely, a disperser \mathcal{D} is *strong* if for each $(X, (H, \eta)) \in \mathcal{D}$, (H, η) is the trivial extended strip decomposition of $G - X$. Recall that this means that (H, η) simply decomposes $G - X$ into connected components: H is an edgeless graph with vertices mapped bijectively to connected components of $G - X$; then the atoms of (H, η) are exactly the connected components of $G - X$. As for strong dispersers the decomposition (H, η) is uniquely determined by X , we will somewhat abuse notation and regard strong dispersers as simply families of sets X , instead of pairs of the form $(X, (H, \eta))$. Intuitively, a strong disperser for G is simply a family of subsets of vertices such that for every possible weight function \mathbf{w} , some member of the family is a balanced separator for \mathbf{w} that has a small weight by itself. The notions of QP-dispersibility lifts to *strong QP-dispersibility* by considering strong dispersers instead of regular ones.

Similarly, uniform dispersers imply subexponential-time algorithms.

Theorem 3.6. *Let \mathcal{C} be a hereditary graph class with the following property: there exist constants $n_0 > 0$, $\tau > 0$, and $\xi \in (0, 1)$ and an algorithm that, given a connected graph $G \in \mathcal{C}$ with $n \geq n_0$ vertices and such that $|N_G[v]| \leq \tau n^\xi$ for every $v \in V(G)$, outputs in polynomial time a ξ -uniform disperser for G . Then, MWIS restricted to graphs from \mathcal{C} admits an algorithm with time complexity $2^{\mathcal{O}(n^{1-\xi} \log n)}$.*

In Theorem 3.6, the constant hidden in the big- \mathcal{O} notation may depend on n_0 , τ , and ξ . A hereditary graph class \mathcal{C} satisfying the assumptions of Theorem 3.6 for n_0 , τ , and ξ is called *ξ -uniformly dispersible*.

The rest of this section is devoted to the proofs of Theorems 3.5 and 3.6.

3.3 Using maximum-weight matching

Assume that a graph G is equipped with a weight function \mathbf{w} and an extended strip decomposition (H, η) . Furthermore, for every atom A of (H, η) we are given an independent set $I(A) \subseteq A$.

Construct a graph H' as follows: start with the graph H and then, for every edge $e = uv$ of H , add a new vertex x_e and edges $x_e u$ and $x_e v$. Furthermore, define weight function \mathbf{w}' on $E(H')$ as follows:

$$\begin{aligned}\mathbf{w}'(x_e u) &= \mathbf{w}(I(A_e^u)) - \mathbf{w}(I(A_u)) - \mathbf{w}(I(A_e^\perp)), \\ \mathbf{w}'(x_e v) &= \mathbf{w}(I(A_e^v)) - \mathbf{w}(I(A_v)) - \mathbf{w}(I(A_e^\perp)), \\ \mathbf{w}'(e) &= \mathbf{w}(I(A_e^{uv})) - \mathbf{w}(I(A_u)) - \mathbf{w}(I(A_v)) - \mathbf{w}(I(A_e^\perp)) - \sum_{T \in \mathcal{T}(e)} \mathbf{w}(I(A_T)).\end{aligned}$$

We claim that the problem of finding maximum-weight matching in (H', \mathbf{w}') is closely related to the problem of finding MWIS in (G, \mathbf{w}) . Let

$$a = \sum_{v \in V(H)} \mathbf{w}(I(A_v)) + \sum_{e \in E(H)} \mathbf{w}(I(A_e^\perp)) + \sum_{T \in \mathcal{T}(H)} \mathbf{w}(I(A_T)).$$

For a family \mathcal{A} of atoms of (H, η) , we define $M(\mathcal{A}) \subseteq E(H')$ as follows. For every $e = uv \in E(H)$, we insert into $M(\mathcal{A})$:

- the edge e if $A_e^{uv} \in \mathcal{A}$,
- the edge $x_e u$ if $A_e^u \in \mathcal{A}$, and
- the edge $x_e v$ if $A_e^v \in \mathcal{A}$.

A direct check shows the following.

Claim 3.7. *If \mathcal{A} is an independent family of atoms of (H, η) , then $M(\mathcal{A})$ is a matching in H' . Furthermore,*

$$\mathbf{w}'(M(\mathcal{A})) \geq -a + \sum_{A \in \mathcal{A}} \mathbf{w}(I(A)). \quad (1)$$

Proof. First we verify that $M(\mathcal{A})$ is a matching in H' . From the definition of independent set of atoms we infer that for every $e = uv \in E(H)$ at most one of the edges e , $x_e u$, or $x_e v$ belongs to $M(\mathcal{A})$. Furthermore, if $x_e u$ or e belongs to $M(\mathcal{A})$, we have A_e^{uv} or A_e^u belonging to \mathcal{A} , from which we infer that neither A_u nor $A_{e'v'}$ nor A_e^v belongs to \mathcal{A} for any other $e' = uv' \in E(H)$ incident with u in H . In particular, neither e' nor $x_{e'} u$ belongs to $M(\mathcal{A})$. Also, if $A_e^{uv} \in \mathcal{A}$ and $T \in \mathcal{T}(e)$, then $A_T \notin \mathcal{A}$ and $A_{e'}^{u'v'} \notin \mathcal{A}$ for every other edge $e' = u'v'$ of T .

For the weight bound, we consider their contribution to the left and right hand side of (1) one-by-one.

- for every atom A of the form A_e^{uv} , A_e^u , or A_e^v ,
 - if $A \in \mathcal{A}$, then the term $\mathbf{w}(I(A))$ appears once on the left hand side and once on the right hand side,
 - if $A \notin \mathcal{A}$, then the term $\mathbf{w}(I(A))$ does not appear at all in (1);
- for every $e = uv \in E(H)$,
 - if $A_e^\perp \in \mathcal{A}$, then the term $\mathbf{w}(I(A_e^\perp))$ does not appear on the left hand side (as then $A_e^u, A_e^v, A_e^{uv} \notin \mathcal{A}$) and its appearances on right hand side in a and $\sum_{A \in \mathcal{A}} \mathbf{w}(I(A))$ cancel out,
 - if $A_e^\perp \notin \mathcal{A}$, then the term $\mathbf{w}(I(A_e^\perp))$ appears with -1 coefficient on the right hand side (in the $-a$ term), while on the left hand side it appears with -1 coefficient if A_e^u, A_e^v , or A_e^{uv} belongs to \mathcal{A} , and does not appear at all otherwise.

- for every $v \in V(H)$,
 - if $A_v \in \mathcal{A}$, then the appearances of $\mathbf{w}(I(A_v))$ on the right hand side cancel out, while this term does not appear on the left hand side (the definition of independence ensures that no atom A_e^v nor A_e^{uv} is in \mathcal{A} for any edge $e = uv$ incident with v),
 - if $A_v \notin \mathcal{A}$, then $\mathbf{w}(I(A_v))$ appears with -1 coefficient on the right hand side, while the independence of \mathcal{A} implies that for at most one edge $e = uv$ incident with v the atom A_e^v or A_e^{uv} belongs to \mathcal{A} and, consequently, $\mathbf{w}(I(A_v))$ either does not appear on the left hand side or appears once with -1 coefficient;
- for every triangle T in H ,
 - if $A_T \in \mathcal{A}$, then the appearances of $\mathbf{w}(I(A_T))$ on the right hand side cancel out, while this term does not appear on the left hand side (the definition of independence ensures that no atom A_e^{uv} is in \mathcal{A} for any edge $e = uv$ of T),
 - if $A_T \notin \mathcal{A}$, then $\mathbf{w}(I(A_T))$ appears with -1 coefficient on the right hand side, while the independence of \mathcal{A} implies that for at most one edge $e = uv$ of T the atom A_e^{uv} belongs to \mathcal{A} and, consequently, $\mathbf{w}(I(A_T))$ either does not appear on the left hand side or appears once with -1 coefficient.

Thus, we have shown that for every atom A , the coefficient in front of $\mathbf{w}(I(A))$ on the left hand side of (1) is not smaller than the coefficient on the right hand side. This finishes the proof of the claim. \square

In the other direction, for $M \subseteq E(H')$ define a family $\mathcal{A}(M)$ of atoms of G as follows.

- For every edge $e = uv \in E(H) \cap M$, insert A_e^{uv} into $\mathcal{A}(M)$.
- For every edge $x_e u \in M \setminus E(H)$, insert A_e^u into $\mathcal{A}(M)$.
- For every edge $e = uv \in E(H)$ such that neither e , $x_e u$, nor $x_e v$ is in M , insert A_e^\perp into $\mathcal{A}(M)$.
- For every vertex $v \in V(H)$ such that none of the edges of M is incident with v , insert A_v into $\mathcal{A}(M)$.
- For every triangle T in H such that none of the edges of M is in M , insert A_T into $\mathcal{A}(M)$.

Again, a direct check shows the following.

Claim 3.8. *If M is a matching in H' , then $\mathcal{A}(M)$ is an independent family of atoms of (H, η) . Furthermore,*

$$\sum_{A \in \mathcal{A}(M)} \mathbf{w}(I(A)) = a + \mathbf{w}'(M). \quad (2)$$

Proof. To show that $\mathcal{A}(M)$ is independent, we consider the cases how two atoms can be conflicting one-by-one. For Case (i), since at most one edge e , $x_e u$, $x_e v$ for $e = uv \in E(H)$ belongs to M , we have that exactly one of the atoms A_e^\perp , A_e^u , A_e^v , A_e^{uv} belongs to $\mathcal{A}(M)$. For Case (ii), we insert A_v into $\mathcal{A}(M)$ only if none of the edges of M is incident with v , which in particular implies that neither A_e^v nor A_e^{uv} is in $\mathcal{A}(M)$ for any edge $e = uv \in E(H)$ incident with v . For Case (iii), since M is a matching, for every $u \in V(H)$ and two distinct edges $e = uv$ and $e' = uv'$ incident with u in H , at most one of the edges e , e' , $x_e u$, and $x_{e'} u$ belong to M , and thus at most one of the atoms A_e^u , $A_{e'}^{uv}$, $A_{e'}^u$, and $A_{e'}^{uv'}$ belong to $\mathcal{A}(M)$. Finally, for Case (iv), if $A_T \in \mathcal{A}(M)$, then none of the edges of T are in M and thus no atom A_e^{uv} for $e = uv$ of T is in $\mathcal{A}(M)$.

For the weight bound, we consider atoms and their contribution to (2) one-by-one.

- for every atom A_e^{uv} for $e = uv \in E(H)$,

- if $e \in M$, then the term $\mathbf{w}(I(A_e^{uv}))$ appears once on the left hand side of (2) (as $A_e^{uv} \in \mathcal{A}(M)$) and once on the right hand side (as a part of $\mathbf{w}'(e)$),
- if $e \notin M$, then the term $\mathbf{w}(I(A_e^{uv}))$ does not appear at all in (2);
- for every atom A_e^u for $e = uv \in E(H)$,
 - if $x_e u \in M$, then the term $\mathbf{w}(I(A_e^u))$ appears once on the left hand side of (2) (as $A_e^u \in \mathcal{A}(M)$) and once on the right hand side (as a part of $\mathbf{w}'(x_e u)$),
 - if $e \notin M$, then the term $\mathbf{w}(I(A_e^u))$ does not appear at all in (2);
- for every atom A_e^\perp for $e = uv \in E(H)$,
 - if neither of the edges $x_e u$, $x_e v$, or e belongs to M , then $A_e^\perp \in \mathcal{A}(M)$ and the term $\mathbf{w}(I(A_e^\perp))$ appears once on the left hand side of (2), while appearing once on the right hand side (once in a and not appearing in $\mathbf{w}'(M)$),
 - if one of the edges $x_e u$, $x_e v$, or e belongs to M , then the corresponding atom A being A_e^u , A_e^v , or A_e^{uv} , respectively, belongs to $\mathcal{A}(M)$, and the term $\mathbf{w}(I(A_e^\perp))$ does not appear on the left hand side while its appearances on the right hand side cancel out with the coefficient $+1$ in the term a and coefficient -1 in the term $\mathbf{w}'(x_e u)$, $\mathbf{w}'(x_e v)$, or $\mathbf{w}(e)$, respectively;
- for every atom A_v for $v \in V(H)$,
 - if there is an edge of M incident with v , say $x_e v$ or e for some $e = uv \in E(H)$, then $\mathbf{w}(I(A_v))$ does not appear on the left hand side of (2), while the appearances of $\mathbf{w}(I(A_v))$ on the right hand side cancel out with the coefficient $+1$ in the term a and coefficient -1 in the term $\mathbf{w}'(x_e v)$ or $\mathbf{w}'(e)$, respectively,
 - if there is no edge of M incident with v , then $A_v \in \mathcal{A}(M)$ and term $\mathbf{w}(I(A_v))$ appears once on the left hand side, while it appears once in a on the right hand side and does not appear in $\mathbf{w}'(M)$;
- for every atom A_T for a triangle T in H ,
 - if there is an edge e of T in M , then $\mathbf{w}(I(A_T))$ does not appear on the left hand side of (2), while the appearances of $\mathbf{w}(I(A_T))$ on the right hand side cancel out with the coefficient $+1$ in the term a and coefficient -1 in the term $\mathbf{w}'(e)$,
 - if no edges of T belong to M , then $A_T \in \mathcal{A}(M)$ and the term $\mathbf{w}(I(A_T))$ appears once on the left hand side, while on the right hand side it appears once in a and does not appear in $\mathbf{w}'(M)$.

Thus, we have shown that for every atom A , the coefficient in front of $\mathbf{w}(I(A))$ on the left hand side of (2) is equal to the one on the right hand side. This finishes the proof of the claim. \square

3.4 Proof of Theorem 3.5

The algorithm of Theorem 3.5 is a standard recursive divide-and-conquer procedure. Let $G \in \mathcal{C}$ be an input graph and \mathbf{w} be a weight function. Fix an accuracy constant $\varepsilon > 0$; w.l.o.g. assume that $1/\varepsilon$ is an integer.

Since we are aiming at an approximation algorithm, we can limit the stretch of the weights value. The problem is trivial if $\mathbf{w}(v) = 0$ for every $v \in V(G)$, so assume otherwise. First, rescale the weight function \mathbf{w} such that $\max_{v \in V(G)} \mathbf{w}(v) = n/\varepsilon$ (allowing rational values of weights). Second, round each weight down to the nearest integer value; since there exists an independent set in G

of weight at least n/ε (take the vertex with maximum weight), this decreases the weight of the maximum-weight independent set by a factor of at least $(1 - n \cdot \varepsilon/n) = (1 - \varepsilon)$. Third, discard all vertices of G of weight 0. Consequently, we can assume that on input the values of \mathbf{w} are integers within range $[1, n/\varepsilon]$.

Initially, we set up an upper bound $\mathbf{m} := n^2/\varepsilon$ on the weight of any independent set in G and \mathbf{w} and fix $\gamma := \varepsilon/(1 + \log(n^2/\varepsilon))$. In a recursive call, we are given an induced subgraph G' of G with the goal to output an independent set I' in G' (that, as we will prove, will be a good approximation). We also pass to a recursive call an upper bound \mathbf{m}' on the weight of the sought independent set.

In the base of the recursion, if G' is edgeless, then we return $I' = V(G')$. Also, if $\mathbf{m}' < 1$, then we return $I' = \emptyset$. In the recursive step, we use the fact that \mathcal{C} is QP-dispersible: for the parameter γ fixed above, there are functions $\delta, \mathbf{S}, \mathbf{T}$ with

$$(\delta(x))^{-1} \in \text{poly}(\log x, \varepsilon^{-1}) \quad \text{and} \quad \mathbf{S}(x), \mathbf{T}(x) \in 2^{\text{poly}(\log x, \varepsilon^{-1})}$$

such that \mathcal{C} is $(\gamma, \delta, \mathbf{S}, \mathbf{T})$ -dispersible. We compute a $(\gamma, \delta(|V(G')|))$ -disperser \mathcal{D} for $(G', \mathbf{w}|_{V(G')})$.

For every $(X, (H, \eta)) \in \mathcal{D}$, we recurse on every atom A of (H, η) , passing an upper bound of $\mathbf{m}' \cdot (1 - \delta(|V(G')|))$, obtaining an independent set $I(A)$. As explained in Section 3.3, we construct the graph H' from H and weight function \mathbf{w}' on $E(H')$ using independent sets $I(A)$. We find a matching M in H' with maximum weight with respect to \mathbf{w}' . We define $I_{(X, (H, \eta))} = \bigcup_{A \in \mathcal{A}(M)} I(A)$. Finally, we return the produced independent set $I_{(X, (H, \eta))}$ of maximum weight among all elements $(X, (H, \eta)) \in \mathcal{D}$.

Running time bound. Since δ is a nonincreasing function, \mathbf{m}' drops below 1 at recursion depth $\mathcal{O}((\delta(n))^{-1} \log(n^2/\varepsilon))$. Since the sets $\eta(e)$, $\eta(v)$, and $\eta(T)$ are pairwise disjoint, there are at most $5n$ nonempty atoms in every (H, η) for $(X, (H, \eta)) \in \mathcal{D}$. Consequently, the recursion tree has size bounded by

$$(\mathbf{S}(n) \cdot 5n)^{\mathcal{O}((\delta(n))^{-1} \log(n^2/\varepsilon))}.$$

At every step, we spend $\mathbf{T}(n)$ time to compute \mathcal{D} , polynomial in n time to compute $\delta(n)$, and $\mathbf{S}(n) \cdot n^{\mathcal{O}(1)}$ time to handle simple manipulations of \mathcal{D} and the find maximum-weight matchings in H' . Hence, the algorithm runs in time bounded by an exponential function of a polynomial in $\log n$ and ε^{-1} .

Approximation guarantee. Let I_0 be an independent set in G of maximum weight. We mark some recursion calls. Initially we mark the initial root call for G . Consider a marked step of the recursion with subgraph G' . Let $(X_0, (H_0, \eta_0))$ be an element of the computed disperser \mathcal{D} that is $(\gamma, \delta(|V(G')|))$ -good for G' and $\mathbf{w}_{I_0 \cap V(G')}$; we henceforth call $(X_0, (H_0, \eta_0))$ the *correct element* of the considered recursive call. Consider the family of atoms $\mathcal{A}_{I_0 \cap V(G' - X_0)}$ for the extended strip decomposition (H_0, η_0) of $G' - X_0$ and the independent set $I_0 \cap V(G' - X_0)$. Claim 3.3 ensures that $\mathcal{A}_{I_0 \cap V(G' - X_0)}$ is independent and its union contains $I_0 \cap V(G' - X_0)$. We mark all recursive calls (being children of the recursive call for G') for atoms $A \in \mathcal{A}_{I_0 \cap V(G' - X_0)}$.

Due to our weight rescaling and rounding, initially $\mathbf{w}(I_0) \leq n^2/\varepsilon$. By a straightforward top-to-bottom induction on the recursion tree, using the definition of being δ -shrinking, we show that at every marked recursive call, if G' is the graph considered in the call and \mathbf{m}' is the passed upper bound, then $\mathbf{w}(I_0 \cap V(G')) \leq \mathbf{m}'$.

In particular, whenever $\mathbf{m}' < 1$, then $I_0 \cap V(G') = \emptyset$ as \mathbf{w} has range contained in $[1, n/\varepsilon]$. Also, if G' is edgeless, then the algorithm returns a maximum-weight independent set in G' . Consequently, at every marked leaf of the recursion with graph G' the returned independent set in G' is of weight at least $\mathbf{w}(I_0 \cap V(G'))$.

Consider a nonleaf marked recursive call and let G' be the graph considered in this call. Let $(X_0, (H_0, \eta_0))$ be the correct element for this recursive call. Furthermore, let $I(A)$ be the independent set output by every recursive call invoked by the considered call for atom A of (H_0, η_0) . Claims 3.7 and 3.8 ensure that the computed independent set for $(X_0, (H_0, \eta_0))$ satisfy

$$\mathbf{w}(I_{(X_0, (H_0, \eta_0))}) \geq \sum_{A \in \mathcal{A}_{I_0 \cap V(G' - X_0)}} \mathbf{w}(I(A)).$$

In particular, the independent set output by the considered recursive call for G' is of weight at least the right hand side of the above inequality.

Let \mathcal{X} be the family of all correct elements over all nonleaf marked recursive calls. We infer that the weight of the independent set output by the root of the recursion is at least

$$\mathbf{w}(I_0) - \sum_{(X_0, (H_0, \eta_0)) \in \mathcal{X}} \mathbf{w}(I_0 \cap X_0).$$

Thus, it remains to estimate the sum of $\mathbf{w}(I_0 \cap X_0)$ over all $(X_0, (H_0, \eta_0)) \in \mathcal{X}$.

Let T be the subtree of the recursion tree induced by all marked calls. We call a nonleaf marked call z *strange* if every marked child of z corresponds to a trivial atom of the correct element $(X_0, (H_0, \eta_0))$ at z , and *normal* otherwise.

For every normal marked call z , denote by $f(z)$ the marked child call for a nontrivial atom A with maximum $\mathbf{w}(I_0 \cap A)$ (breaking ties arbitrarily) and mark the edge $zf(z)$ of T . Let $F \subseteq E(T)$ be the set of marked edges. Clearly, $(V(T), F)$ is a set of upward paths in T . Let Z be the set of top endpoints of these paths, that is, Z consists of the root of T and all recursive calls such that the edge of T between the call and its parent is not marked. For every $z \in V(T)$, let G'_z be the subgraph of G considered in the call z . Note that all marked leaves of T that correspond to trivial atoms are in Z . Let S be the family of strange marked nodes.

As at every marked recursive call, the marked children of the call consider disjoint atoms, we infer that every $v \in I_0$ is contained in at most $1 + \log(\mathbf{w}(I_0))$ graphs G'_z for $z \in Z$ (in at most one leaf corresponding to a trivial atom and, for every other $z \in Z$ with $v \in V(G'_z)$, the weight of the vertices of I_0 in G'_z is at most half of the weight of the vertices of I_0 in the graph G' at the parent of z).

Furthermore, for every normal marked call z , from γ -safeness of the correct element $(X_0, (H_0, \eta_0))$ for $\mathbf{w}_{I_0 \cap V(G'_z)}$ we infer that

$$\mathbf{w}(X_0 \cap I_0) \leq \gamma \cdot \left(\mathbf{w}(I_0 \cap V(G'_z)) - \mathbf{w}(I_0 \cap V(G'_{f(z)})) \right).$$

Summing over all nonleaf marked calls z we infer that

$$\begin{aligned} \sum_{(X_0, (H_0, \eta_0)) \in \mathcal{X}} \mathbf{w}(X_0 \cap I_0) &\leq \gamma \cdot \sum_{z \in Z} \mathbf{w}(I_0 \cap V(G'_z)) + \gamma \cdot \sum_{s \in S} \mathbf{w}(I_0 \cap V(G'_s)) \\ &\leq \frac{\varepsilon}{1 + \log(n^2/\varepsilon)} \cdot (\log(\mathbf{w}(I_0)) \cdot \mathbf{w}(I_0) + \mathbf{w}(I_0)) \\ &\leq \varepsilon \cdot \mathbf{w}(I_0). \end{aligned}$$

Consequently, the returned independent set at the root recursive call is of weight at least $(1 - \varepsilon)\mathbf{w}(I_0)$. This finishes the proof of Theorem 3.5.

3.5 Proof of Theorem 3.6

The algorithm of Theorem 3.6 is again a standard divide-and-conquer procedure, simpler than in the case of Theorem 3.5. By choosing n_0 appropriately we may assume that $n_0 > e^{1/\xi}$, i.e., $n_0^\xi > e$.

Let G be the input graph with a weight function \mathbf{w} . If $n := |V(G)| \leq n_0$, we solve the problem by brute-force in constant time. If G is disconnected, we recurse on every connected component. Otherwise, if there exists $v \in V(G)$ with $|N[v]| > \tau n^\xi$, branch exhaustively on v : in one branch, delete v from G and recurse (consider v not included in the sought solution) and in the second branch, delete $N[v]$ from G , recurse, and add v to the independent set obtained from the recursive call (consider v included in the sought solution). Finally, output the one of the two obtained solutions that has larger weight.

In the remaining case we have $n = |V(G)| > n_0$ and $|N[v]| \leq \tau n^\xi$ for every $v \in V(G)$. Invoke the assumed algorithm that outputs a pair $(X, (H, \eta))$ where $X \subseteq V(G)$ and (H, η) is an extended strip decomposition of $G - X$ such that:

$$|X| \leq n^{-\xi}(n - |A|) \quad \text{and} \quad |A| \leq n - n^\xi \quad \text{for every atom } A \text{ of } (H, \eta).$$

For every independent set $Y \subseteq X$, we proceed as follows. Let (H, η_Y) be (H, η) restricted to $G - (X \cup N[Y])$. That is, (H, η_Y) is obtained from (H, η) by removing all vertices of $N[Y]$ from all the sets in the image of η ; it is straightforward to see that then (H, η_Y) is an extended strip decomposition of $G - (X \cup N[Y])$. We recurse on every atom A of (H, η_Y) , obtaining an independent set $I_Y(A)$. As explained in Section 3.3, we construct the graph H' from H and weight function \mathbf{w}' on $E(H')$ using independent sets $I_Y(A)$. We find a matching M in H' with maximum weight with respect to \mathbf{w}' . We define $I_Y = Y \cup \bigcup_{A \in \mathcal{A}(M)} I_Y(A)$. Finally, we return the independent set I_Y that has the maximum weight among all produced for independent sets $Y \subseteq X$.

Correctness. It is straightforward to verify that every recursive call is invoked on some induced subgraph of G and the set returned by any recursive call is an independent set. By induction on $|V(G)|$, we prove that an application of the algorithm to a graph G returns a maximum-weight independent set in G .

This is obvious for the cases when we apply a brute-force search and when G is disconnected and we recurse on the connected components of G . If we branch on a vertex v with $|N[v]| > \tau n^\xi$, then the correctness is again straightforward as we consider exhaustively cases of v being and not being included in the sought solution. Otherwise, we are in the case where we obtained a ξ -uniform disperser $(X, (H, \eta))$. Let I_0 be a maximum-weight independent set in G' and consider the case $Y = I_0 \cap X$. Then, $\mathcal{A}_{I_0 \setminus Y}$ is an independent family of atoms of (H, η_Y) and by Claim 3.7, $M(\mathcal{A}_{I_0 \setminus Y})$ is a matching in H' . Furthermore, by the inductive assumption, for every atom A of (H, η_Y) , $I_Y(A)$ is an independent set of maximum weight in $G'[A]$. Note that we may apply the inductive assumption here due to $|A| \leq n - n^\xi < n$. In particular, for every $A \in \mathcal{A}_{I_0 \setminus Y}$, $\mathbf{w}(I_Y(A)) \geq \mathbf{w}(I_0 \cap A)$. Therefore by Claims 3.7 and 3.8, the weight of I_Y is at least

$$\mathbf{w}(Y) + \sum_{A \in \mathcal{A}_{I_0 \setminus Y}} \mathbf{w}(I_Y(A)) \geq \mathbf{w}(Y) + \sum_{A \in \mathcal{A}_{I_0 \setminus Y}} \mathbf{w}(I_0 \cap A) = \mathbf{w}(I_0).$$

This concludes the inductive proof that the algorithm returns a maximum-weight independent set in G .

Running time bound. We prove by induction on n that when the algorithm is applied on an n -vertex graph G , the number of leaves of the recursion tree is bounded by $e^{Cn^{1-\xi}(1+\ln n)}$ for some

constant C depending on ξ and n_0 . Since the time spent at internal computation in each recursive call is polynomial in n , the claimed running time bound will follow.

The claim is straightforward for the leaves of the recursion and for non-leaf recursive calls when G is disconnected. In a non-leaf recursive call, if the algorithm branches on a vertex $v \in V(G)$ with $|N[v]| > \tau n^\xi$, in one child recursive call the number of vertices drops by 1, in the second drops by at least τn^ξ . Then the inductive step follows by standard calculations, which we omit here.²

In the remaining case, we have obtained a ξ -uniform disperser $(X, (H, \eta))$. Let k be the number of vertices in the largest atom of (H, η) . By the properties of $(X, (H, \eta))$, we have

$$|X| \leq n^{-\xi} \cdot (n - k) \quad \text{and} \quad n - k \geq n^\xi. \quad (3)$$

For a fixed independent set $Y \subseteq X$, the algorithm recurses on at most $5n$ atoms, each of size at most k , which is strictly smaller than n . Hence, by the inductive hypothesis, for a sufficiently large constant C we have that the total number of leaf nodes of the recursion in descendants of the considered node is bounded by

$$2^{|X|} \cdot 5n \cdot 2^{C \cdot k^{1-\xi}(1+\log k)} \quad (4)$$

If $k \leq e^{1/\xi} \leq n_0$, then all the recursive calls are leaves in the recursion tree, so by (3) their number is bounded by

$$2^{|X|} \cdot 5n \leq \exp\left(\ln 2 \cdot n^{1-\xi} + \ln n + \ln 5\right).$$

This value can be bounded as desired by taking $C \geq \ln 2 + 1 + \ln 5$. Hence, we assume

$$k > e^{1/\xi}. \quad (5)$$

We need the following inequality:

$$\begin{aligned} n^{1-\xi}(1 + \ln n) - k^{1-\xi}(1 + \ln k) &\geq (n - k) \cdot \min_{k \leq x_0 \leq n} \left(\frac{d}{dx} \left(x^{1-\xi}(1 + \ln x) \right) \Big|_{x=x_0} \right) \\ &= (n - k) \cdot \min_{k \leq x_0 \leq n} \left((1 - \xi)x_0^{-\xi}(1 + \ln x_0) + x_0^{-\xi} \right) \\ &= (n - k)n^{-\xi} \left((1 - \xi)(1 + \ln n) + 1 \right). \end{aligned} \quad (6)$$

Here, in the last equality we have used (5), as $x \mapsto x^{-\xi}$ is decreasing for $x > 0$ and $x \mapsto x^{-\xi} \ln x$ is decreasing for $x \geq e^{1/\xi}$.

By applying $n - k \geq n^\xi$, from (6) we obtain that:

$$n^{1-\xi}(1 + \ln n) - k^{1-\xi}(1 + \ln k) \geq (1 - \xi)(n - k)n^{-\xi} + (1 - \xi)(1 + \ln n). \quad (7)$$

With (7) in hand, we are now ready to given an upper bound on (4):

$$\begin{aligned} &2^{|X|} \cdot 5n \cdot \exp\left(C \cdot k^{1-\xi}(1 + \ln k)\right) \\ &\leq \exp\left(\ln 2 \cdot n^{-\xi} \cdot (n - k) + \ln n + \ln 5 + Ck^{1-\xi}(1 + \ln k)\right) \\ &\leq \exp\left(\ln 2 \cdot n^{-\xi} \cdot (n - k) + \ln n + \ln 5 + Cn^{1-\xi}(1 + \ln n) \right. \\ &\quad \left. - C(1 - \xi)(n - k)n^{-\xi} - C(1 - \xi)(1 + \ln n)\right) \\ &\leq \exp\left(Cn^{1-\xi}(1 + \ln n)\right). \end{aligned}$$

²The crucial calculation is $n^{1-\xi} - (n - \tau n^\xi)^{1-\xi} \sim (1 - \xi)\tau$. Alternatively, one can observe that in the recursion tree on every root-to-leaf path there are at most $\mathcal{O}(n^{1-\xi} \ln n / \tau)$ edges that correspond to branching on a vertex v with $|N[v]| > \tau n^\xi$ and deleting the whole $N[v]$.

In the last inequality we have used (7) and $C \geq \frac{\ln 5}{1-\xi}$.

This finishes the proof of the time complexity and of Theorem 3.6.

4 Heavy vertices and strong dispersers

Let G be a graph, $\mathbf{w}: V(G) \rightarrow \mathbb{N}$ be a weight function, and $I \subseteq V(G)$ be an independent set. For a real $\beta \in [0, 1]$, a vertex $w \in V(G)$ is β -heavy (with respect to I) if $\mathbf{w}(N[w] \cap I) \geq \beta \cdot \mathbf{w}(I)$. A simple coupon-collector argument shows the following.

Lemma 4.1. *Let G be an n -vertex graph for $n \geq 2$, $\mathbf{w}: V(G) \rightarrow \mathbb{N}$ 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 $J \subseteq I$ of size at most $\lceil \beta^{-1} \log n \rceil$ such that $N[J]$ contains all β -heavy vertices with respect to I .*

Proof. Let Z be the set of β -heavy vertices. We consider a probability distribution on I where a vertex $v \in I$ is chosen with probability $\mathbf{w}(v)/\mathbf{w}(I)$. For every $z \in Z$, a vertex $v \in I$ chosen at random according to this distribution satisfies $z \in N[v]$ with probability at least β . Consequently, if J is a set of $\lceil \beta \log n \rceil$ vertices of I each chosen independently at random according to this distribution, then for every $z \in Z$ the probability that $z \notin N[J]$ is less than $(1 - \beta)^{\beta \log n} < 1/n$ (here we used that $\beta \leq 1/2$ and $n \geq 2$). By the union bound, the probability that $Z \subseteq N[J]$ is positive. \square

Next we prove a general-usage lemma that reduces the task of finding small dispersers to connected graphs where the neighborhood of every vertex is not β -heavy with regards to some fixed maximum-weight independent set we are looking for. This is done essentially as follows: we first guess the set J of β -heavy vertices of size $\text{poly}(\gamma^{-1}, \log n)$ using Lemma 4.1, focus on the heaviest connected component of $G - N[J]$, and construct a suitable disperser for this component. This idea can be used to prove the following statement.

Lemma 4.2. *Let \mathcal{C} be a hereditary graph class. Suppose there is a polynomial $p(\cdot)$ such that given any $\sigma > 0$ and n -vertex connected graph $G \in \mathcal{C}$ one can in polynomial time compute a family \mathcal{N} with $|\mathcal{N}| \leq \text{poly}(n)$ consisting of pairs of the form $(X, (H, \eta))$, where $X \subseteq V(G)$ and (H, η) is an extended strip decomposition of $G - X$, such that the following holds: For every weight function $\mathbf{w}: V(G) \rightarrow \mathbb{N}$ satisfying $\mathbf{w}(N[v]) \leq p(\sigma)\mathbf{w}(V(G))$ for each $v \in V(G)$ there exists $(X, (H, \eta)) \in \mathcal{N}$ such that*

$$\mathbf{w}(A) \leq (1 - p(\sigma)) \cdot \mathbf{w}(G) \quad \text{and} \quad \mathbf{w}(X) \leq \sigma \cdot \mathbf{w}(G - A) \quad \text{for every atom } A \text{ of } (H, \eta).$$

Then the class \mathcal{C} is QP-dispersible. Moreover, if it is always the case that all the extended strip decompositions appearing in the family \mathcal{N} are trivial (i.e. corresponding to the partition into connected components), then \mathcal{C} is strongly QP-dispersible.

Proof. Suppose without loss of generality that $p(x) \geq x$ for all positive x . Fix $\gamma \in (0, 1/2)$. Fix $G \in \mathcal{C}$ on n vertices supplied with a weight function $\mathbf{w}: V(G) \rightarrow \mathbb{N}$.

We present the construction of a disperser for G as a nondeterministic procedure that, for a given independent set I with $\mathbf{w}(I) > 0$, produces a pair $(X, (H, \eta))$, where $X \subseteq V(G)$ and (H, η) is an extended strip decomposition of $G - X$, that is $(\gamma, p(\gamma))$ -good for \mathbf{w}_I , i.e. we shall have $\delta(n) = p(\gamma)$. We argue that this nondeterministic procedure has $\mathbf{S}(n)$ possible runs that can be enumerated in time $\mathbf{S}(n) \cdot \text{poly}(n)$ without the knowledge of I , where the function $\mathbf{S}(n)$ will be chosen later. Then the constructed disperser \mathcal{D} comprises of all sets X constructed by all possible runs, and thus has size at most $\mathbf{S}(n)$. As each run has polynomial length, the running time of the construction of \mathcal{D} is $\mathbf{T}(n) \leq \mathbf{S}(n) \cdot \text{poly}(n)$.

Therefore, fix an independent set I in G with $\mathbf{w}(I) > 0$. Recall that \mathbf{w}_I is a weight function on G obtained from \mathbf{w} by changing the weight of vertices outside of I to 0.

First, apply Lemma 4.1 to G , \mathbf{w}_I , I , and constant $\beta = p(\gamma)/2$. This yields a set $J \subseteq I$ of size at most $2p(\gamma)^{-1} \log n + 1 = \text{poly}(\gamma^{-1}, \log n)$ such that $N[J]$ contains all vertices that are $p(\gamma)/2$ -heavy w.r.t. \mathbf{w}_I . The procedure nondeterministically guesses the set J ; note that there are $2^{\text{poly}(\gamma^{-1}, \log n)}$ choices for J . Then

$$\mathbf{w}_I(N[v]) \leq p(\gamma)/2 \cdot \mathbf{w}_I(G) \quad \text{for every vertex } v \in V(G) \setminus N[J]. \quad (8)$$

Let G' be the heaviest (w.r.t. \mathbf{w}_I) connected component of $G - N[J]$. Our nondeterministic procedure guesses G' (n options) and whether $\mathbf{w}_I(G') \leq \mathbf{w}_I(G)/2$ or not (2 options).

Suppose first that $\mathbf{w}_I(G') \leq \mathbf{w}_I(G)/2$. Then observe that putting $X = N(J)$ and (H, η) as the trivial extended strip decomposition of $G - X$, we find that $(X, (H, \eta))$ is $(0, 1/2)$ -good for G . Indeed, in $G - X$ every vertex of J is isolated, so it corresponds to a trivial atom of (H, η) , while every other atom of (H, η) corresponds to a connected component of $G - N[J]$ and hence it has weight at most $\mathbf{w}_I(G)/2$. On the other hand, $\mathbf{w}_I(X) = 0$, because $X = N(J)$ is disjoint with I .

Therefore, from now on we focus on the second case when

$$\mathbf{w}_I(G') > \mathbf{w}_I(G)/2. \quad (9)$$

Since \mathcal{C} is hereditary, we have $G' \in \mathcal{C}$. Hence, we may apply the assumed algorithm to G' for $\sigma = \gamma$, yielding in polynomial time a family \mathcal{N} of size $\text{poly}(n)$ consisting of pairs of the form $(X', (H', \eta'))$, where (H', η') is an extended strip decomposition of $G' - X'$. As by (8) and (9) we have

$$\mathbf{w}_I(N_{G'}[v]) \leq \mathbf{w}_I(N_G[v]) \leq p(\gamma)/2 \cdot \mathbf{w}_I(G) \leq p(\gamma) \cdot \mathbf{w}_I(G') \quad \text{for every } v \in V(G'),$$

by assumption there exists $(X', (H', \eta')) \in \mathcal{N}$ satisfying the following:

$$\mathbf{w}_I(A) \leq (1 - p(\gamma)) \cdot \mathbf{w}_I(G') \quad \text{and} \quad \mathbf{w}_I(X') \leq \gamma \cdot \mathbf{w}(V(G') \setminus A) \quad \text{for every atom } A \text{ of } (H', \eta').$$

By choosing among $|\mathcal{N}| = \text{poly}(n)$ options, our nondeterministic procedure guesses $(X', (H', \eta'))$ satisfying the above.

Consider $X = X' \cup N(J)$. Observe that since G' is a connected component of $G - N(J)$, we have $\text{cc}(G' - X') \subseteq \text{cc}(G - X)$. Let now (H, η) be the extended strip decomposition of $G - X$ obtained from (H', η') by adding every connected component $C \in \text{cc}(G - X) \setminus \text{cc}(G' - X')$ as a separate piece of the decomposition: we add a new node w_C that is isolated in H and set $\eta(w_C) = V(C)$.

Claim 4.3. *The pair $(X, (H, \eta))$ is $(\gamma, p(\gamma))$ -good for G and \mathbf{w}_I .*

Proof. First, observe that since $N(J) \cap I = \emptyset$, we have

$$\mathbf{w}_I(X) = \mathbf{w}_I(X') \leq \gamma \cdot \mathbf{w}_I(G' - B) \leq \gamma \cdot \mathbf{w}_I(G'),$$

where B is any nontrivial atom of (H', η') .

Consider any nontrivial atom A of (H, η) . Since vertices of J form trivial atoms in (H, η) , we have that either A is a connected component of $G - N[J]$ that is different from G' , or A is an atom of (H', η') .

In the first case, by (9) we infer that $\mathbf{w}_I(A) < \mathbf{w}_I(G)/2$. Moreover, since G' and A are disjoint, we have $\mathbf{w}_I(G - A) \geq \mathbf{w}_I(G')$. The latter assertion together with $\mathbf{w}_I(X) \leq \gamma \cdot \mathbf{w}_I(G')$ implies that $\mathbf{w}_I(X) \leq \gamma \cdot \mathbf{w}_I(G - A)$, as required.

Consider now the second case. First, by assumption we have $\mathbf{w}_I(A) < (1 - p(\gamma))\mathbf{w}_I(G') \leq (1 - p(\gamma))\mathbf{w}_I(G)$. Second, again by assumption we have $\mathbf{w}_I(X) = \mathbf{w}_I(X') \leq \gamma \cdot \mathbf{w}_I(G' - A) \leq \gamma \cdot \mathbf{w}_I(G - A)$.

Thus, in both cases we conclude that $(X, (H, \eta))$ is $(\gamma, p(\gamma))$ -good for G and \mathbf{w}_I . \square

Therefore, in all cases the nondeterministic procedure produced a pair $(X, (H, \eta))$ that is $(\gamma, p(\gamma))$ -good for G and I .

We conclude by observing that the nondeterminism used by the procedure comes from:

- choosing J , for which there are $2^{\text{poly}(\gamma^{-1}, \log n)}$ choices;
- choosing G' and whether $\mathbf{w}_I(G') \leq \mathbf{w}_I(G)/2$, for which there are at most $2n$ choices; and
- choosing $(X', (H', \eta')) \in \mathcal{N}$, for which there are $\text{poly}(n)$ choices.

Hence, we can set $\mathbf{S}(n) \in 2^{\text{poly}(\gamma^{-1}, \log n)}$ for the size of the computed strong disperser and, consequently, also the construction running time is $\mathbf{T}(n) = \mathbf{S}(n) \cdot \text{poly}(n) = 2^{\text{poly}(\gamma^{-1}, \log n)}$. We conclude that \mathcal{C} is $(\gamma, p(\gamma), 2^{\text{poly}(\gamma^{-1}, \log n)}, 2^{\text{poly}(\gamma^{-1}, \log n)})$ -dispersible for every $\gamma \in (0, 1/2)$, hence it is QP-dispersible. Moreover, it can be easily seen that if the assumed algorithm only returns trivial extended strip decompositions, then also all the constructed extended strip decompositions are trivial and, consequently, \mathcal{C} is strongly QP-dispersible. \square

5 Dispersers in P_t -free graphs and graphs without a long hole

In this section we focus on the class of P_t -free graphs and graphs excluding a long hole.

As a warm-up, to show how our framework works, we prove the following statement.

Theorem 5.1. *For every $t \in \mathbb{N}$, the class of P_t -free graphs is strongly QP-dispersible and $\frac{1}{2}$ -uniformly dispersible.*

The proof of Theorem 5.1 relies on a classical construction used by Gyárfás [26] to prove that P_t -free graphs are χ -bounded, which is usually called the *Gyárfás path*. In Section 5.1 we encapsulate this concept in a versatile claim, as we will reuse it later on.

For $t \in \mathbb{N}$, a graph G is $C_{\geq t}$ -free if G excludes every cycle C_ℓ for $\ell \geq t$ as an induced subgraph. For instance, the *long-hole-free* graphs considered in [10] are exactly $C_{\geq 5}$ -free graphs. In Section 5.2, we prove the following strengthening of Theorem 5.1 that implies Theorem 1.4.

Theorem 5.2. *For every $t \in \mathbb{N}$, the class of $C_{\geq t}$ -free graphs is strongly QP-dispersible and $\frac{1}{2}$ -uniformly dispersible.*

The structural results obtained in Section 5.2 also directly imply Theorem 1.5.

5.1 Gyárfás' path

The following lemma encapsulates a classical construction of Gyárfás [26].

Lemma 5.3. *Let $\alpha \in (0, 1/2)$ be a real. Let G be a connected graph endowed with a weight function $\mathbf{w}: V(G) \rightarrow \mathbb{N}$, and let u be any vertex of G . Then there is an induced path $Q = (v_0, v_1, \dots, v_k)$ in G (possibly with $k = -1$ and Q being empty) such that, denoting $G_0 = G - v_0$ and $G_i = G - N[v_0, \dots, v_{i-1}]$ for $i \in \{1, \dots, k+1\}$, the following holds:*

- (P1) $u = v_0$ unless $k = -1$ (where we put $G_0 = G$);
- (P2) for every $C \in \text{cc}(G_{k+1})$, we have $\mathbf{w}(C) \leq (1 - \alpha)\mathbf{w}(G)$; and
- (P3) for every $i \in \{0, 1, \dots, k\}$, there is a connected component D of G_i such that $\mathbf{w}(D) > (1 - \alpha)\mathbf{w}(G)$ and D contains a neighbor of v_i .

Moreover, given G and u one can compute in polynomial time a family \mathcal{Q} consisting of $\mathcal{O}(|V(G)|^2)$ induced paths in G , each starting at u , so that for every $\alpha \in (0, 1/2)$ and weight function $\mathbf{w}: V(G) \rightarrow \mathbb{N}$ there exists $Q \in \mathcal{Q}$ satisfying the above properties for α and \mathbf{w} .

Proof. We first prove the existential statement and then argue how the reasoning can be turned into a suitable algorithm.

Call an induced subgraph H of G *heavy* if $\mathbf{w}(H) > (1 - \alpha)\mathbf{w}(G)$ and *light* otherwise. We construct P inductively so that after constructing v_0, \dots, v_ℓ , these vertices induce a path (v_0, \dots, v_ℓ) in G and property (P3) is satisfied for all $i \in \{0, 1, \dots, \ell\}$. If no component of G is heavy, we may finish the construction immediately by setting $k = -1$ and Q as the empty path. Otherwise, we start by setting $v_0 = u$. Since $G_0 = G - v_0$ and G is connected, the unique (due to $\alpha < 1/2$) heavy component of G_0 is adjacent to v_0 and (P3) is satisfied for $i = 0$.

For $\ell \geq 0$, the construction of $v_{\ell+1}$ is implemented as follows. By (P3) for $i = \ell$, there is a connected component D of G_ℓ that is heavy and adjacent to v_ℓ . As $\alpha < 1/2$, no other connected component of G_ℓ can be heavy. Since $G_{\ell+1}$ is an induced subgraph of G_ℓ , either every connected component of $G_{\ell+1}$ is light, or there is exactly one heavy connected component D' of $G_{\ell+1}$ that is moreover an induced subgraph of D . In the former case, we may finish the construction by setting $k = \ell$, as then (P2) is satisfied. Otherwise, observe that $G_{\ell+1}$ is obtained from G_ℓ by removing vertices of $N[v_\ell] \setminus N[v_0, \dots, v_{\ell-1}]$, hence D' is a connected component of $D - (N[v_\ell] \cap V(D))$. Here observe that $N[v_\ell] \cap V(D)$ is non-empty, because D is adjacent to v_ℓ . Consequently, there exists a vertex $v_{\ell+1} \in V(D)$ that is simultaneously adjacent to v_ℓ and to D' . Since $v_{\ell+1} \in V(D)$, $v_{\ell+1}$ is not adjacent to any of the vertices $v_0, \dots, v_{\ell-1}$. We conclude that the induced path (v_0, \dots, v_ℓ) can be extended by $v_{\ell+1}$ so that (P3) is satisfied for $i = \ell + 1$.

Since G is finite, the construction eventually finishes yielding a path Q satisfying both (P2) and (P3). We are left with arguing the algorithmic statement.

Observe that in the above reasoning, we used the constant α and the function \mathbf{w} only in order to verify whether the construction should be finished, or to identify the heavy connected component D' of $D - (N[v_\ell] \cap V(D))$. Having identified D' , $v_{\ell+1}$ can be chosen freely among the common neighbors of D' and v_ℓ . Fix beforehand a total order of $V(G)$ and assume that $v_{\ell+1}$ is always chosen as the smallest eligible vertex. Consider any run of the algorithm for G, α, \mathbf{w} and for $i \in \{0, \dots, k-1\}$ let D_i be the unique heavy connected component of G_i . Since $\alpha < 1/2$, subgraphs D_i pairwise intersect. Since $G_0, G_1, G_2, \dots, G_{k-1}$ is a descending chain in the induced subgraph order and each D_i is a connected component of G_i , we conclude that $D_0, D_1, D_2, \dots, D_{k-1}$ is also a descending chain in the induced subgraph order. Consequently, there exists a vertex z that is contained in each of D_0, D_1, \dots, D_{k-1} . Now comes the main observation: knowing z and having constructed G_i , we may identify D_i as the unique connected component of G_i that contains z . Thus, a path Q suitable for α, \mathbf{w} can be constructed knowing only k and z (given the total order fixed beforehand). Constructing such a path Q for every choice of k and z , of which there are at most $\mathcal{O}(|V(G)|^2)$ many, yields the desired family \mathcal{Q} . \square

Note that in the statement of Theorem 5.3, graph G_{k+1} is equal to $G - N[Q]$ unless Q is empty, when it is equal to $G - u$.

Now Theorem 5.1 follows from a straightforward combination of Lemmas 4.2 and Lemma 5.3.

Proof of Theorem 5.1. We first argue the $\frac{1}{2}$ -uniform dispersibility. Set $\tau = \frac{1/4}{t-1}$ and assume G is an n -vertex connected P_t -free graph and $|N[v]| \leq \tau\sqrt{n}$ for every $v \in V(G)$. Apply Lemma 5.3 to G , arbitrary $u \in V(G)$, $\alpha = 1/4$, and uniform weight function \mathbf{w} , obtaining a path Q . Note that Q cannot be empty, as G is connected and the weight function is uniform. Since G is P_t -free, Q has at most $t-1$ vertices, so $X := N[V(Q)]$ has size at most $\frac{1}{4}\sqrt{n}$. On the other hand, every connected component C of $G - X = G_{k+1}$ has at most $(1 - \alpha)|V(G)| = \frac{3}{4}n$ vertices, so $|X| \leq n^{-\frac{1}{2}}(n - |C|)$. Hence, we can return X and a trivial extended strip decomposition of $G - X$ as the desired uniform disperser.

For QP-dispersibility, the argument is only slightly longer. Without loss of generality assume $t \geq 4$. We argue that the class of P_t -free graphs satisfies the prerequisites of Lemma 4.2. Thus we assume we are given a connected P_t -free graph G and a parameter $\sigma > 0$. Consider applying Lemma 5.3 to G and any vertex $u \in V(G)$. We infer that in polynomial time we can construct a polynomial-size family \mathcal{Q} of induced paths in G satisfying in particular the following: for each weight function $\mathbf{w}: V(G) \rightarrow \mathbb{N}$ there exists $Q \in \mathcal{Q}$ such that $\mathbf{w}(C) \leq \frac{3}{4}\mathbf{w}(G)$ for every $C \in \text{cc}(G - X)$, where $X = N[Q]$ if Q is non-empty and $X = \{u\}$ otherwise. Since G is P_t -free, every path in \mathcal{Q} has less than t vertices. Consequently, supposing $\mathbf{w}(N[v]) \leq \frac{\sigma}{4t} \cdot \mathbf{w}(V(G))$ for every vertex v , we have $\mathbf{w}(X) \leq \sigma/4 \cdot \mathbf{w}(V(G))$ for every $Q \in \mathcal{Q}$, and in particular $\mathbf{w}(X) \leq \sigma \cdot \mathbf{w}(G - C)$ for every $C \in \text{cc}(G - X)$.

From \mathcal{Q} construct a family \mathcal{N} by including, for every $Q \in \mathcal{Q}$, a pair $(X, (H, \eta))$ where X is as above and (H, η) is the trivial extended strip decomposition of $G - X$. The reasoning of the previous paragraph shows that the assumptions of Lemma 4.2 are satisfied for $p(\sigma) = \frac{\sigma}{4t}$. Therefore, from Lemma 4.2 we conclude that the class of P_t -free graphs is strongly QP-dispersible. \square

5.2 Graphs without long holes

The proof of Theorem 5.2 follows from applying exactly the same reasoning as in the proof of Theorem 5.1, except that in order to obtain a suitable path family \mathcal{Q} we use the following Lemma 5.4, instead of Lemma 5.3. Furthermore, the lemma below also directly implies Theorem 1.5 via standard arguments (see e.g. Corollary 1 of [4]).

Lemma 5.4. *Let G be a connected $C_{\geq t}$ -free graph supplied with a weight function $\mathbf{w}: V(G) \rightarrow \mathbb{N}$. Then in G there is an induced path Q on less than t vertices such that*

$$\mathbf{w}(C) \leq \frac{3}{4}\mathbf{w}(G) \quad \text{for every } C \in \text{cc}(G - N[Q]).$$

Moreover, given G alone, one can enumerate in polynomial time a family \mathcal{Q} of $\mathcal{O}(|V(G)|^2)$ induced paths on less than t vertices with a guarantee that for every weight function \mathbf{w} there exists $Q \in \mathcal{Q}$ satisfying the above for \mathbf{w} .

Proof. Without loss of generality assume $t \geq 4$. We first focus on proving the existential statement. At the end we will argue how the enumeration statement can be derived from the enumeration statement of Lemma 5.3.

Fix any vertex u in G and apply the existential statement of Lemma 5.3 to G , vertex u , weight function \mathbf{w} , and $\alpha = \frac{1}{4}$. This yields an induced path $R = (v_0, v_1, \dots, v_k)$ satisfying properties (P2) and (P3), where $v_0 = u$. If $k + 1 < t$ then, by (P2), we may simply take $Q = R$, or $Q = (u)$ in case R is the empty path. Hence, from now on assume that $k \geq t - 1$.

Let R' and R'' be the subpaths of R defined as

$$R' = (v_{k-t+1}, \dots, v_{k-1}) \quad \text{and} \quad R'' = (v_{k-t+2}, \dots, v_k).$$

Note that each of R', R'' has $t - 1$ vertices. In the rest of the proof we argue the following claim: one of paths R', R'' satisfies the condition required of Q .

Suppose otherwise: there are components $D' \in \text{cc}(G - N[R'])$ and $D'' \in \text{cc}(G - N[R''])$ with $\mathbf{w}(D') > \frac{3}{4}\mathbf{w}(G)$ and $\mathbf{w}(D'') > \frac{3}{4}\mathbf{w}(G)$. Note that then D' and D'' are unique. We observe the following.

Claim 5.5. *D' is adjacent to v_k .*

Proof. By property (P3) of Lemma 5.3, $G - N[v_0, \dots, v_{k-1}]$ contains a (unique) connected component C of weight more than $\frac{3}{4}\mathbf{w}(G)$ that is moreover adjacent to v_k . As $G - N[v_0, \dots, v_{k-1}]$ is an induced subgraph of $G - N[R']$ and $\mathbf{w}(D') > \frac{3}{4}\mathbf{w}(G)$, it follows that C is contained in D' . Hence D' is adjacent to v_k . \lrcorner

Claim 5.6. D'' is adjacent to v_{k-t+1} .

Proof. Graph $G - N[v_0, \dots, v_k]$ can be obtained from $G - N[R'']$ by removing vertices v_0, \dots, v_{k-t} and all the neighbors of v_0, \dots, v_{k-t+1} that do not belong to $N[R'']$; denote the set of those vertices by Z . Thus, every connected component of $G - N[R'']$ that is not a connected component of $G - N[v_0, \dots, v_k]$ necessarily contains a vertex of Z . Since by property (P2) of Lemma 5.3, no connected component of $G - N[v_0, \dots, v_k]$ has weight more than $\frac{3}{4}\mathbf{w}(G)$, while this is the case for D'' , we conclude that $V(D'') \cap Z \neq \emptyset$. Now observe that $G[Z \cup \{v_{k-t+1}\}]$ is connected and all vertices of Z are present in $G - N[R'']$. Hence some vertex of $V(D'') \cap Z$ is adjacent to v_{k-t+1} , implying the claim. \lrcorner

Claim 5.7. $V(D') \cap V(D'') \neq \emptyset$.

Proof. Follows immediately from $\mathbf{w}(D') > \frac{3}{4}\mathbf{w}(G)$ and $\mathbf{w}(D'') > \frac{3}{4}\mathbf{w}(G)$. \lrcorner

By Claims 5.5, 5.6, 5.7 it follows that there exists an induced path P with endpoints v_{k-t+1} and v_k whose all internal vertices belong to $V(D') \cup V(D'')$. As vertices of $V(D') \cup V(D'')$ are non-adjacent to $v_{k-t+2}, \dots, v_{k-1}$ by definition, path P together with the subpath of R from v_{k-t+2} to v_{k-1} induce a cycle of length at least t , a contradiction.

For the enumeration statement, it suffices to compute the family \mathcal{R} provided by Lemma 5.3 and, for every $R \in \mathcal{R}$, include in \mathcal{Q} either R , if its number of vertices is less than t , or both R' and R'' as defined above for R . \square

6 Rooted subdivided claw

In this section we will focus on the classes of graphs excluding a claw subdivided a fixed number of times. We try to construct such subdivided claws with the use of Theorem 2.1. This provides us with extended strip decompositions of the considered graphs.

We introduce a useful lemma that encapsulates the way we will use Theorem 2.1. We first need a definition.

Definition 6.1. Let G be a graph and let $Z \subseteq V(G)$ be such that $|Z| = 3$. An extended strip decomposition (H, η) *shatters* Z if the following condition holds: whenever P_1, P_2, P_3 is a triple of induced paths in G that are pairwise disjoint and non-adjacent, and each of them has one endpoint in Z , then there is no atom in (H, η) that intersects or is adjacent to each of P_1, P_2, P_3 .

Lemma 6.2. *Let G be a graph and let $Z \subseteq V(G)$ be such that $|Z| = 3$. Then one can in polynomial time find either an induced tree in G containing all vertices of Z , or an extended strip decomposition (H, η) of G that shatters Z .*

The proof of Lemma 6.2 is postponed to Section 6.1. Note that contrary to Theorem 2.1, Lemma 6.2 does not assume that the graph is connected.

We move to the main point of this section, which concerns classes excluding subdivided claws.

Definition 6.3. A *subdivided claw* is a graph obtained from the claw $K_{1,3}$ and subdividing each of its edges an arbitrary number of times. The degree-1 vertices are then called the *tips* of the claw, while the unique vertex of degree 3 is the *center*. A subdivided claw is a $(\geq t)$ -*claw* if all its tips are at distance at least t from its center. A graph G is $Y_{\geq t}$ -*free* if it does not contain any $(\geq t)$ -claw as an induced subgraph.

Theorem 6.4. *For every $t \in \mathbb{N}$, the class of $Y_{\geq t}$ -free graphs is QP-dispersible and $\frac{1}{9}$ -uniformly dispersible.*

Theorem 6.4 is a consequence of Lemma 6.5 below. Indeed, to obtain QP-dispersibility it suffices to combine Lemma 6.5 with Lemma 4.2, while to obtain $\frac{1}{9}$ -uniformly dispersibility, apply Lemma 6.5 for $\sigma = n^{-1/9}$ (setting n_0 sufficiently large such that $\sigma < \frac{1}{100t}$), uniform weight function, and any u , and observe that a pair $(X, (H, \eta))$ satisfying (C2) is a $\frac{1}{9}$ -uniform disperser; we can find such a pair in polynomial time by inspecting all the members of \mathcal{N} .

Lemma 6.5. *Fix an integer $t \geq 4$. Let G be a connected graph supplied with a weight function $\mathbf{w}: V(G) \rightarrow \mathbb{N}$ and let $\sigma \in (0, \frac{1}{100t})$ be such that*

$$\mathbf{w}(N[v]) \leq \sigma^8 \cdot \mathbf{w}(G) \text{ for every } v \in V(G). \quad (10)$$

Let u be any vertex of G . Then there is either

- (C1) an induced $(\geq t)$ -claw in G with one of the tips being u , or*
- (C2) a subset of vertices $X \subseteq V(G)$ and an extended strip decomposition (H, η) of $G - X$ such that*

$$\mathbf{w}(A) \leq (1 - \sigma^7) \cdot \mathbf{w}(G) \quad \text{and} \quad \mathbf{w}(X) \leq \sigma \cdot \mathbf{w}(G - A) \quad \text{for every atom } A \text{ of } (H, \eta).$$

Moreover, given G and u one can in polynomial time either find conclusion (C1), or enumerate a family \mathcal{N} of $\mathcal{O}(|V(G)|^4)$ pairs $(X, (H, \eta))$ such that for every $\sigma \in (0, \frac{1}{100t})$ and every weight function $\mathbf{w}: V(G) \rightarrow \mathbb{N}$ satisfying (10) there exists $(X, (H, \eta)) \in \mathcal{N}$ satisfying (C2) for \mathbf{w} and σ .

Proof. We first focus on proving the existential statement. At the end we will argue how the enumeration statement can be derived using the enumeration statement of Lemma 5.3.

Apply Lemma 5.3 to G , u , \mathbf{w} , and $\alpha = \sigma$, yielding a suitable path $Q = (v_0, \dots, v_k)$, where $v_0 = u$ (unless $k = -1$ and Q is empty). As in Lemma 5.3, denote $G_0 = G - u$ and $G_i = G - N[v_0, \dots, v_{i-1}]$ for $i \in \{1, \dots, k+1\}$. For $i \in \{0, \dots, k+1\}$, let D_i be the heaviest (w.r.t. \mathbf{w}) connected component of G_i . Then by (P3) and (P2) we have

$$\mathbf{w}(D_i) > (1 - \sigma) \cdot \mathbf{w}(G) \text{ for } i \leq k \quad \text{and} \quad \mathbf{w}(D_{k+1}) \leq (1 - \sigma) \cdot \mathbf{w}(G). \quad (11)$$

Also, as argued in the proof of Lemma 5.3, D_j is an induced subgraph of D_i for each $i, j \in \{0, \dots, k\}$ with $i \leq j$.

If $\mathbf{w}(D_0) \leq (1 - \sigma^5) \cdot \mathbf{w}(G)$, then conclusion (C2) can be obtained by taking $X = \{v_0\}$ and (H, η) to be the trivial extended strip decomposition of $G - X$. This is because $\mathbf{w}(X) = \mathbf{w}(v_0) \leq \sigma^8 \cdot \mathbf{w}(G)$ due to (10), while $\mathbf{w}(G - D) \geq \sigma^5 \cdot \mathbf{w}(G)$ for every connected component D of $G - X$. Note that if $k = -1$, then in particular $\mathbf{w}(D_0) \leq (1 - \sigma) \cdot \mathbf{w}(G) \leq (1 - \sigma^5) \cdot \mathbf{w}(G)$, so the above analysis can be applied as well. Hence, from now on assume that $k \geq 0$ and $\mathbf{w}(D_0) > (1 - \sigma^5) \cdot \mathbf{w}(G)$.

Define p and q as the largest indices satisfying the following:

$$\mathbf{w}(D_p) > (1 - \sigma^5) \cdot \mathbf{w}(G) \quad \text{and} \quad \mathbf{w}(D_q) > (1 - \sigma^3) \cdot \mathbf{w}(G).$$

By (11) and the discussion of the previous paragraph we have that p and q are well-defined and satisfy $0 \leq p \leq q \leq k$.

We now observe that indices $0, p, q, k$ have to be well-separated from each other, or otherwise we are done. For this, consider the following paths in G :

$$R_1 = (v_0, v_1, \dots, v_{p-2}), \quad R_2 = (v_p, v_{p+1}, \dots, v_{q-2}), \quad R_3 = (v_q, v_{q+1}, \dots, v_{k-1}).$$

Note that the above path formally may be empty in case the index of the second endpoint is smaller than that of the first endpoint; in a moment we will see that this is actually never the case. We now verify that the neighborhood of each of these paths has to have a significant weight, or otherwise we are done.

Claim 6.6. *If we have*

$$\mathbf{w}(N[R_1]) \leq (\sigma^6/2) \cdot \mathbf{w}(G) \quad \text{or} \quad \mathbf{w}(N[R_2]) \leq (\sigma^4/2) \cdot \mathbf{w}(G) \quad \text{or} \quad \mathbf{w}(N[R_3]) \leq (\sigma^2/2) \cdot \mathbf{w}(G),$$

then conclusion (C2) can be obtained.

Proof. We first consider the case when $\mathbf{w}(N[R_1]) \leq (\sigma^6/2) \cdot \mathbf{w}(G)$, which is slightly simpler. By assumption we have $\mathbf{w}(D_{p+1}) \leq (1 - \sigma^5) \cdot \mathbf{w}(G)$ where D_{p+1} is the heaviest connected component of $G - N[v_0, v_1, \dots, v_p]$. On the other hand, we have

$$\begin{aligned} \mathbf{w}(N[v_0, v_1, \dots, v_p]) &\leq \mathbf{w}(N[R_1]) + \mathbf{w}(N[v_{p-1}]) + \mathbf{w}(N[v_p]) \\ &\leq (\sigma^6/2 + 2\sigma^8) \cdot \mathbf{w}(G) \\ &\leq \sigma^6 \cdot \mathbf{w}(G). \end{aligned}$$

Hence, we can obtain conclusion (C2) by taking $X = N[v_0, v_1, \dots, v_p]$ and the trivial extended strip decomposition of $G - X$. Indeed, for every connected component D of $G - X$ we have $\mathbf{w}(D) \leq \mathbf{w}(D_{p+1}) \leq (1 - \sigma^5) \cdot \mathbf{w}(G)$, implying also that $\mathbf{w}(X) \leq \sigma^6 \cdot \mathbf{w}(G) \leq \sigma \cdot \mathbf{w}(G - D)$.

Now, consider the case when $\mathbf{w}(N[R_2]) \leq (\sigma^4/2) \cdot \mathbf{w}(G)$. Observe that we also have $\mathbf{w}(N[R_1]) \leq \mathbf{w}(G) - \mathbf{w}(D_p) < \sigma^5 \cdot \mathbf{w}(G)$, because D_p and $N[R_1]$ are disjoint. By assumption we have $\mathbf{w}(D_{q+1}) \leq (1 - \sigma^3) \cdot \mathbf{w}(G)$ where D_{q+1} is the heaviest connected component of $G - N[v_0, v_1, \dots, v_q]$. On the other hand, we have

$$\begin{aligned} \mathbf{w}(N[v_0, v_1, \dots, v_q]) &\leq \mathbf{w}(N[R_1]) + \mathbf{w}(N[R_2]) + \mathbf{w}(N[v_{p-1}]) + \mathbf{w}(N[v_{q-1}]) \\ &\leq (\sigma^5 + \sigma^4/2 + 2\sigma^8) \cdot \mathbf{w}(G) \\ &\leq \sigma^4 \cdot \mathbf{w}(G). \end{aligned}$$

Hence, we can obtain conclusion (C2) by taking $X = N[v_0, v_1, \dots, v_q]$ and the trivial extended strip decomposition of $G - X$. Indeed, for every connected component D of $G - X$ we have $\mathbf{w}(D) \leq \mathbf{w}(D_{q+1}) \leq (1 - \sigma^3) \cdot \mathbf{w}(G)$, implying also that $\mathbf{w}(X) \leq \sigma^4 \cdot \mathbf{w}(G) \leq \sigma \cdot \mathbf{w}(G - D)$.

Finally, consider the case when $\mathbf{w}(N[R_3]) \leq (\sigma^2/2) \cdot \mathbf{w}(G)$. As in the previous case, we have $\mathbf{w}(N[R_1]) < \sigma^5 \cdot \mathbf{w}(G)$ and $\mathbf{w}(N[R_2]) < \sigma^3 \cdot \mathbf{w}(G)$. By the construction of Q we have $\mathbf{w}(D_{k+1}) \leq (1 - \sigma) \cdot \mathbf{w}(G)$ where D_{k+1} is the heaviest connected component of $G - N[v_0, v_1, \dots, v_k]$. On the other hand, we have

$$\begin{aligned} \mathbf{w}(N[v_0, v_1, \dots, v_k]) &\leq \mathbf{w}(N[R_1]) + \mathbf{w}(N[R_2]) + \mathbf{w}(N[R_3]) + \mathbf{w}(N[v_{p-1}]) + \mathbf{w}(N[v_{q-1}]) + \mathbf{w}(N[v_k]) \\ &\leq (\sigma^5 + \sigma^3 + \sigma^2/2 + 3\sigma^8) \cdot \mathbf{w}(G) \\ &\leq \sigma^2 \cdot \mathbf{w}(G). \end{aligned}$$

Hence, we can obtain conclusion (C2) by taking $X = N[v_0, v_1, \dots, v_k]$ and the trivial extended strip decomposition of $G - X$. Indeed, for every connected component D of $G - X$ we have $\mathbf{w}(D) \leq \mathbf{w}(D_{k+1}) \leq (1 - \sigma) \cdot \mathbf{w}(G)$, implying also that $\mathbf{w}(X) \leq \sigma^2 \cdot \mathbf{w}(G) \leq \sigma \cdot \mathbf{w}(G - D)$. \square

We proceed under the assumption that the prerequisite of Claim 6.6 does not hold, that is,

$$\mathbf{w}(N[R_1]) > (\sigma^6/2) \cdot \mathbf{w}(G) \text{ and } \mathbf{w}(N[R_2]) > (\sigma^4/2) \cdot \mathbf{w}(G) \text{ and } \mathbf{w}(N[R_3]) > (\sigma^2/2) \cdot \mathbf{w}(G). \quad (12)$$

From this we argue that $0, p, q, k$ have to be well-separated from each other.

Claim 6.7. *It holds that*

$$p - 0 > t + 1 \quad \text{and} \quad q - p > t + 1 \quad \text{and} \quad k - q > t + 1.$$

Proof. Observe that if $p - 0 \leq t + 1$, then

$$\mathbf{w}(N[R_1]) \leq \sum_{i=0}^{p-1} \mathbf{w}(N[v_i]) \leq (t+1)\sigma^8 \cdot \mathbf{w}(G) < (\sigma^6/2) \cdot \mathbf{w}(G),$$

contradicting the assumption (12). The proof for the other two inequalities is analogous. \square

We will also consider the following subpaths of Q :

$$Q_1 = (v_0, v_1, \dots, v_{t-1}), \quad Q_2 = (v_p, v_{p+1}, \dots, v_{p+t-1}), \quad Q_3 = (v_q, v_{q+1}, \dots, v_{q+t-1}).$$

Note that by Claim 6.7, paths Q_1, Q_2, Q_3 are pairwise disjoint and non-adjacent, and they are prefixes of R_1, R_2, R_3 , respectively. Also, each of them consists of t vertices.

Now, let

$$G' = G - ((N(Q_1) \cup N(Q_2) \cup N(Q_3)) \setminus \{v_t, v_{p+t}, v_{q+t}\}).$$

Note that in G' , paths Q_1, Q_2, Q_3 are preserved, but they become *detached* in the following sense: only one endpoint ($v_{t-1}, v_{p+t-1}, v_{q+t-1}$, respectively) is adjacent to one vertex from the rest of the graph (v_t, v_{p+t}, v_{q+t} , respectively). Also, paths R_1, R_2, R_3 are also preserved in G' .

We now apply Lemma 6.2 to graph G' with

$$Z = \{v_0, v_p, v_q\}.$$

This either yields an induced tree T in G' containing v_0, v_p, v_q , or an extended strip decomposition (H', η') of G' which shatters v_0, v_p, v_q . In the first case, by the construction of G' it follows that T has to contain an induced ($\geq t$)-claw T' with tips v_0, v_p, v_q . As $v_0 = u$, then T' witnesses that conclusion (C1) holds. Hence, from now on we assume the second case.

Observe that

$$\mathbf{w}(N[Q_1] \cup N[Q_2] \cup N[Q_3]) \leq 3t \cdot \sigma^8 \cdot \mathbf{w}(G) \leq (\sigma^7/2) \cdot \mathbf{w}(G).$$

Hence, it now suffices to prove the following:

$$\mathbf{w}(A) \leq (1 - \sigma^6/2) \cdot \mathbf{w}(G) \text{ for every atom } A \text{ of } (H, \eta). \quad (13)$$

Indeed, if (13) holds, then we can obtain conclusion (C2) by taking $X = N[Q_1] \cup N[Q_2] \cup N[Q_3]$ and (H, η) to be (H', η') with all the vertices of $V(Q_1) \cup V(Q_2) \cup V(Q_3) \cup \{v_t, v_{p+t}, v_{q+t}\}$ removed, because then

$$\mathbf{w}(X) \leq (\sigma^7/2) \cdot \mathbf{w}(G) \leq \sigma \cdot \mathbf{w}(G - A) \text{ for every atom } A \text{ of } (H, \eta).$$

Suppose that, contrary to (13), there exists an atom A in (H', η') such that $\mathbf{w}(A) > (1 - \sigma^6/2) \cdot \mathbf{w}(G)$. Note that since Q is an induced path in G , we have that R_1, R_2, R_3 are induced paths in G'

that are disjoint and pairwise non-adjacent. Since (H', η') shatters $\{v_0, v_p, v_q\}$, we conclude that the atom A is disjoint with $N[R_t]$ for at least one $t \in \{1, 2, 3\}$. However, this combined with (12) and the assumption that $\mathbf{w}(A) > (1 - \sigma^6/2) \cdot \mathbf{w}(G)$ yields that $\mathbf{w}(A \cup N[R_t]) > \mathbf{w}(G)$, a contradiction. This concludes the proof of the existential statement.

For the enumeration statement, it suffices to enumerate the family \mathcal{Q} provided by Lemma 5.3, and for every $Q = (v_0, \dots, v_k)$ and $0 \leq p \leq q \leq k$ include in \mathcal{N} the following pairs:

- $X = N[v_0, \dots, v_p]$, and the trivial extended strip decomposition of $G - X$;
- $X = N[v_0, \dots, v_q]$, and the trivial extended strip decomposition of $G - X$;
- $X = N[v_0, \dots, v_k]$, and the trivial extended strip decomposition of $G - X$;
- $X = N[v_0, \dots, v_{t-1}] \cup N[v_p, \dots, v_{p+t-1}] \cup N[v_q, \dots, v_{q+t-1}]$, and the extended strip decomposition obtained by applying Theorem 2.1 to G' (in the notation from the proof above) and $Z = \{v_0, v_p, v_q\}$.

In the last point, if for any choice of Q, p, q we obtain an induced $(\geq t)$ -claw with u as one of the tips, then it can be reported by the algorithm. Otherwise from the above proof it is clear that the enumerated family \mathcal{N} consists of $\mathcal{O}(|V(G)|^4)$ pairs and satisfies the required property. \square

6.1 Proof of Lemma 6.2

The following technical lemma describes how triples of disjoint, non-adjacent paths starting at peripheral vertices behave in an extended strip decomposition of a graph.

Lemma 6.8. *Let (H, η) be an extended strip decomposition of a graph G . Suppose P_1, P_2, P_3 are three induced paths in G that are pairwise disjoint and non-adjacent, and moreover each of P_1, P_2, P_3 has an endpoint that is peripheral in (H, η) . Then in (H, η) there is no atom that would intersect or be adjacent to each of P_1, P_2, P_3 .*

Proof. A feature of (H, η) is a vertex, an edge, or a triangle of H . We introduce the following incidence relation between features: two edges are incident if they share a vertex, a vertex of H is incident to all edges of H it is an endpoint of, and a triangle of H is incident to all edges of H that it contains. Thus, vertices and triangles are considered to be non-incident. Note that every edge of G connects either vertices from $\eta(f)$ for the same feature f , or from $\eta(f)$ and $\eta(f')$ for two incident features f, f' .

Consider an induced path Q in G . A *visit* of a feature f by Q is a maximal subpath of Q consisting of vertices belonging to $\eta(f)$. The order of vertices on Q naturally gives rise to an order of visits of features by Q . We now establish a few basic properties of how induced paths in G behave w.r.t. the decomposition (H, η) in order to get an understanding of the interaction between P_1, P_2, P_3 in (H, η) .

Claim 6.9. *Suppose Q is an induced path in G . Consider some visit W of a feature f by Q , where f is either a vertex or a triangle. Let W_1 be the visit on Q directly before W and W_2 be the visit on Q directly after W ; possibly W_1 or W_2 does not exist when W is the first, respectively last visit of a feature on Q . Then W_1 and W_2 , if existent, are visits of an edge in H that is incident to f , and if they are both existent, then this is the same edge of H .*

Proof. Let f_1, f_2 be the features visited by Q in W_1, W_2 , respectively. The fact that f_1, f_2 are both edges incident to f follows directly from the definition of an extended strip decomposition, in particular the conditions on edges of G . We are left with proving that if both W_1, W_2 exist (i.e., visit W appears neither at the front nor at the end of Q), then $f_1 = f_2$.

Consider first the case when f is a vertex. Then f_1 and f_2 are both edges incident to f . Moreover, then the last vertex of the visit W_1 belongs to $\eta(f_1, f)$, while the first vertex of the visit of W_2 belongs to $\eta(f_2, f)$. But if $f_1 \neq f_2$, then $\eta(f_1, f)$ and $\eta(f_2, f)$ would be complete to each other, which would contradict the assumption that P is induced. Therefore we conclude that $f_1 = f_2$.

Consider now the case when f is a triangle; then f_1 and f_2 are both edges contained in f . Supposing $f_1 \neq f_2$, we may denote $f = uvw$, $f_1 = uv$, $f_2 = uw$. Then the last vertex of the visit W_1 belongs to $\eta(uv, u) \cap \eta(uv, v)$, while the first vertex of the visit of W_2 belongs to $\eta(uw, u) \cap \eta(uw, w)$. This means that these two vertices are adjacent, because they belong to $\eta(uv, u)$ and $\eta(uw, u)$, respectively. This is a contradiction with the assumption that P is an induced path. \square

Claim 6.10. *Suppose Q_1 and Q_2 are two induced paths in G that do not intersect and are non-adjacent. Suppose further that Q_1 has endpoint z_1 and Q_2 has endpoint z_2 such that z_1, z_2 are peripheral. Then there does not exist an edge uv of H such that both Q_1 and Q_2 intersect $\eta(uv, u)$.*

Proof. Orient Q_1, Q_2 so that z_1, z_2 are their first vertices, respectively. Suppose the claim does not hold and let (uv, u) be such that both Q_1 and Q_2 intersect $\eta(uv, u)$; among such pairs, choose (uv, u) so that the distance from z_1 to the first vertex of $\eta(uv, u)$ on Q_1 plus the distance from z_2 to the first vertex of $\eta(uv, u)$ on Q_2 is as small as possible. Let y_1, y_2 be the first vertices on Q_1, Q_2 that belong to $\eta(uv, u)$, respectively.

Consider first the corner case when $z_1 = y_1$ and $z_2 = y_2$. Since both z_1, z_2 are peripheral and $z_1, z_2 \in \eta(uv)$, it must be that $\eta(uv, u) = \{z_1\}$ and $\eta(uv, v) = \{z_2\}$, or vice versa. But then $z_2 \notin \eta(uv, u)$, a contradiction.

Hence, either $y_1 \neq z_1$ or $y_2 \neq z_2$. Assume without loss of generality the former and let x_1 be the vertex directly preceding y_1 on Q_1 ; clearly, $x_1 \notin \eta(uv, u)$ by the choice of y_1 .

First observe that x_1 cannot belong to $(\eta(\cdot)$ of) any vertex or triangle of H . Indeed, if this was the case, then by Claim 6.9 we would conclude that Q_1 would already intersect $\eta(uv, u)$ before x_1 , so y_1 would not be the first vertex of $\eta(uv, u)$ on Q_1 . Hence, either $x_1 \in \eta(uw, u)$ for some $w \neq v$, or $x_1 \in \eta(uv) \setminus \eta(uv, u)$. In the former case we infer that x_1 and y_2 would be adjacent, a contradiction with the assumption that Q_1 and Q_2 are non-adjacent. Hence, we have $x_1 \in \eta(uv) \setminus \eta(uv, u)$. Since Q_1 starts in a peripheral vertex z_1 , we conclude that on Q_1 there is a vertex $t_1 \in \eta(uv, v)$ that appears no later than x_1 (possibly $t_1 = x_1$).

Consider now the corner case when $z_2 = y_2$. Let $ww' \in E(H)$ be such that w has degree 1 in H and $\eta(ww', w) = \{z_2\}$. Then $(uv, u) = (ww', w)$ or $(uv, u) = (ww', w')$. In the former case we would have $y_1 \in \eta(ww', w)$ and $y_1 \neq y_2 = z_2$, a contradiction to $|\eta(ww', w)| = 1$. In the latter case, however, we would have $t_1 \in \eta(ww', w)$, again a contradiction to $|\eta(ww', w)| = 1$, because $t_1 \neq z_2$.

Hence, from now on assume that $z_2 \neq y_2$. By applying the same reasoning to Q_2 as we did for Q_1 we infer that on Q_2 there is a vertex $t_2 \in \eta(uv, v)$ that appears earlier than y_2 . However, now the existence of $t_1, t_2 \in \eta(uv, v)$ is a contradiction with the choice of the pair (uv, u) . \square

We proceed to the proof of the lemma statement. It suffices to prove the statement for atoms of the form A_e^{uv} for some edge $e = uv \in E(H)$, as every atom of (H, η) is contained in an atom of this form, apart from atoms corresponding to isolated vertices of H for which the statement holds trivially. Recall that then

$$A_e^{uv} = \eta(u) \cup \eta(v) \cup \eta(uv) \cup \bigcup_{T \supseteq uv} \eta(T).$$

We first note the following.

Claim 6.11. *Among paths P_1, P_2, P_3 , at most one can intersect the set $\eta(u) \cup \bigcup_{w: uw \in E(H)} \eta(uw, u)$.*

Proof. As each of the paths P_1, P_2, P_3 starts in a peripheral vertex, intersecting $\eta(u)$ entails intersecting $\bigcup_{w: uw \in E(H)} \eta(uw, u)$. By Claim 6.10, no two of the paths P_1, P_2, P_3 intersect the same set $\eta(uw, u)$, for some w with $uw \in E(H)$. However, if, say, P_1 intersected $\eta(uw_1, u)$ and P_2 intersected $\eta(uw_2, u)$ for some $uw_1, uw_2 \in E(H)$, $w_1 \neq w_2$, then P_1 and P_2 would contain adjacent vertices, a contradiction. \square

Denote

$$\begin{aligned} K_u &= \eta(u) \cup \bigcup_{w: uw \in E(H)} \eta(uw, u), \\ K_v &= \eta(v) \cup \bigcup_{w: vw \in E(H)} \eta(vw, v), \\ L &= \bigcup_{T \supseteq uv} \eta(T) \cup (\eta(uv) \setminus (\eta(uv, u) \cup \eta(uv, v))), \end{aligned}$$

and observe that

$$N[A_e^{uv}] = K_u \cup K_v \cup L.$$

By Claim 6.11, K_u above can be intersected by at most one of the paths P_1, P_2, P_3 , and similarly for K_v . Hence, if $N[A_e^{uv}]$ is intersected by all three paths P_1, P_2, P_3 , then one of them, say P_3 , intersects L while not intersecting $K_u \cup K_v$. Note that

$$N(L) \subseteq \eta(uv, u) \cup \eta(uv, v) \subseteq K_u \cup K_v,$$

hence we conclude that P_3 is entirely contained in L . This is a contradiction with the assumption that one of the endpoints of P_3 is peripheral in (H, η) . \square

The proof of Lemma 6.2 is now an easy combination of Theorem 2.1 and Lemma 6.8.

Proof of Lemma 6.2. Consider first the case when vertices of Z are not all in the same connected component of G . Then we can output the trivial extended strip decomposition of G , as it clearly shatters Z .

Suppose now that all vertices of Z are in the same connected component C of G . Apply Theorem 2.1 to Z in C . Then, in polynomial time we can either find an induced tree T in C that contains all vertices of Z (which can be directly returned), or an extended strip decomposition (H_C, η_C) of C such that all vertices of Z are peripheral in (H_C, η_C) . In the latter case, by Lemma 6.8 we conclude that Z is shattered by (H_C, η_C) in C . We augment (H_C, η_C) to an extended strip decomposition (H, η) of G by adding for every component $C' \in \text{cc}(G)$, $C' \neq C$, a new isolated vertex $v_{C'}$ with $\eta(v_{C'}) = V(C')$. Then it is easy to see that (H, η) shatters Z in G . \square

6.2 Proof of Theorems 1.1 and 1.2

With Theorem 6.4, the proofs of Theorems 1.1 and 1.2 are straightforward.

Proof of Theorem 1.1. Let H be such that every connected component of H is a path or a subdivided claw. Let Y be a subdivided claw such that every connected component of H is an induced subgraph of Y .

Let G be H -free, let $\mathbf{w}: V(G) \rightarrow \mathbb{N}$ be a weight function, and let $\varepsilon > 0$ be an accuracy parameter. Set $\beta := \varepsilon / (2|V(H)|)$. Let I be an independent set in (G, \mathbf{w}) of maximum-weight. By Lemma 4.1, there exists a set $J \subseteq I$ of size at most $\lceil \beta \log n \rceil = \mathcal{O}(\varepsilon^{-1} \log n)$ such that all β -heavy vertices w.r.t. I are contained in $N[J]$. By branching into $n^{\mathcal{O}(\varepsilon^{-1} \log n)}$ subcases, we guess the set J .

Let $G' = G - N[J]$. Let \mathcal{C} be a maximal family of connected components of H such that $H[\bigcup \mathcal{C}]$ is an induced subgraph of G' . Let $H' = H[\bigcup \mathcal{C}]$ and note that H' is a proper induced subgraph of H . Let $X \subseteq V(G')$ be such that $G'[X]$ is isomorphic to H' . Note that $|X| < |V(H)|$.

Observe that $G'' := G' - N[X]$ is Y -free. Indeed, if G'' contains Y as an induced subgraph, then, by the choice of Y , it contains some connected component C of $H - V(H')$ as an induced subgraph. Together with $G'[X]$ isomorphic to H' , this contradicts the choice of \mathcal{C} .

Apply the algorithm of Theorem 6.4 to find an independent set I'' in G'' that is a $(1 - \varepsilon/2)$ -approximation to a maximum weight independent set problem on G'' and $\mathbf{w}|_{V(G'')}$. This takes time $2^{\text{poly}(\varepsilon^{-1}, \log n)}$ and we have $\mathbf{w}(I'') \geq (1 - \varepsilon/2)\mathbf{w}(I \cap V(G''))$. Finally, we return $I' := I'' \cup J$.

Consider the branch where J is guessed correctly. We have $\mathbf{w}(I \cap N[X]) \leq \beta |X| \mathbf{w}(I) < \varepsilon/2 \cdot \mathbf{w}(I)$. Furthermore,

$$\mathbf{w}(I) - \mathbf{w}(I'') \leq (\varepsilon/2) \cdot \mathbf{w}(I \cap V(G'')) + \mathbf{w}(I \cap N[X]) \leq \varepsilon \mathbf{w}(I).$$

This finishes the proof of Theorem 1.1. \square

Proof of Theorem 1.2. Again, let H be such that every connected component of H is a path or a subdivided claw and let Y be a subdivided claw such that every connected component of H is an induced subgraph of Y .

Let G be H -free and let $\mathbf{w}: V(G) \rightarrow \mathbb{N}$ be a weight function. If G is disconnected, recurse on every connected component separately. If G contains a vertex v such that $|N[v]| > |V(G)|^{1/9}$, branch exhaustively on v : in one branch, delete v and recurse, in the other branch, delete $N[v]$, recurse, and add v to the independent set returned by the recursive call; finally, output the one of the two obtained independent sets of higher weight.

Otherwise, let \mathcal{C} be a maximal family of connected components of H such that $H[\bigcup \mathcal{C}]$ is an induced subgraph of G . Let $H' = H[\bigcup \mathcal{C}]$ and note that H' is a proper induced subgraph of H . Let $X \subseteq V(G)$ be such that $G[X]$ is isomorphic to H' . Note that $|X| < |V(H)|$.

Observe that $G - N[X]$ is Y -free. Indeed, if $G - N[X]$ contains Y as an induced subgraph, then, by the choice of Y , it contains some connected component C of $H - V(H')$ as an induced subgraph. Together with $G[X]$ isomorphic to H' , this contradicts the choice of \mathcal{C} .

For every independent set $Z \subseteq N[X]$, invoke the algorithm of Theorem 6.4 on the Y -free graph $G - (N[X] \cup N[Z])$, obtaining an independent set $I(Z)$, and observe that $I_Z := Z \cup I(Z)$ is an independent set in G . Out of all independent sets I_Z for $Z \subseteq N[X]$, return the one of maximum weight.

Since we consider every independent set $Z \subseteq N[X]$, the returned solution is indeed an independent set in G of maximum possible weight.

For the running time bound, note that $|N[X]| < |V(G)|^{1/9} \cdot |V(H)|$, hence we invoke less than $2^{|V(G)|^{1/9} \cdot |V(H)|}$ calls to the algorithm of Theorem 6.4, each taking $2^{\mathcal{O}(|V(G)|^{8/9} \log |V(G)|)}$ time. In a recursive step, the analysis is straightforward if G is disconnected and follows by standard arguments if G contains a vertex v with $|N[v]| > |V(G)|^{1/9}$. This finishes the proof of Theorem 1.2. \square

7 A small generalization

In this section we generalize Theorem 1.1 by proving that Conjecture 1.3 holds for all subcubic forests H that have at most three vertices of degree three. Let L be the *lobster graph* depicted in Figure 1. For $t \in \mathbb{N}$, an $(\geq t)$ -lobster is any graph obtained from L by subdividing every edge at least $t - 1$ times. Then a graph is $L_{\geq t}$ -free if it does not contain any $(\geq t)$ -lobster as an induced subgraph.

By Theorem 3.5, to prove Conjecture 1.3 for all subcubic forests H that have at most three vertices of degree three it suffices to prove the following.




Figure 1: The lobster graph L .

Theorem 7.1. *For every $t \in \mathbb{N}$, the class of $L_{\geq t}$ -free graphs is QP-dispersible and $\frac{1}{41}$ -uniformly dispersible.*

Similarly as was the case with Theorem 6.4 and Lemma 6.5, to prove Theorem 7.1 it suffices to show the following.

Lemma 7.2. *Fix an integer $t \geq 4$. Let G be a connected graph supplied with a weight function $\mathbf{w}: V(G) \rightarrow \mathbb{N}$ and let $\sigma \in (0, \frac{1}{100t})$ be such that*

$$\mathbf{w}(N[v]) \leq \sigma^{40} \cdot \mathbf{w}(G) \text{ for every } v \in V(G). \quad (14)$$

Then there is either

- (L1) an induced $(\geq t)$ -lobster in G , or*
- (L2) a subset of vertices $X \subseteq V(G)$ and an extended strip decomposition (H, η) of $G - X$ such that*

$$\mathbf{w}(A) \leq (1 - \sigma^{39}) \cdot \mathbf{w}(G) \quad \text{and} \quad \mathbf{w}(X) \leq \sigma \cdot \mathbf{w}(G - A) \quad \text{for every atom } A \text{ of } (H, \eta).$$

Moreover, given G one can in polynomial time either find conclusion (L1), or enumerate a family \mathcal{N} of $\mathcal{O}(|V(G)|^{12})$ pairs $(X, (H, \eta))$ such that for every $\sigma \in (0, \frac{1}{100t})$ and every weight function $\mathbf{w}: V(G) \rightarrow \mathbb{N}$ satisfying (14) there exists $(X, (H, \eta)) \in \mathcal{N}$ satisfying (L2) for \mathbf{w} and σ .

The proof of Lemma 7.2 uses the same set of ideas as that of Lemma 6.5, but the number of steps in the construction of a lobster is larger and one needs to tend to more technical details. Essentially, the overall strategy can be summarized as follows. We try to construct an induced $(\geq t)$ -lobster in G ; each step of the construction may fail and produce conclusion (L2) as a result. We start by building the right claw T of the lobster using Lemma 6.5, however we make sure that one of the tips of this claw, call it w , is adjacent to a connected component of $G - N[T - w]$ that contains almost the whole weight of the graph. This is done by applying Lemma 5.3 to construct a long Gyárfás path Q , and then applying Lemma 6.5 not to any initial vertex, but to a vertex v_i of the Gyárfás path such that $\mathbf{w}(G_i)$ is significantly separated from $\mathbf{w}(G)$. Having constructed T and w , we forget about the first Gyárfás path Q and construct, using Lemma 5.3, a second Gyárfás path P , this time starting from w . We construct the left claw S of the lobster, but again we start this construction at later sections of P so that we can ensure the following: there is a tip v of S so that in the graph $G - N[S - v] - N[T - w]$ there is a connected component containing v, w , and a long prefix of P . Then we construct the “tail” (that is, the middle pendant edge) of the lobster from the saved prefix of P , by applying Lemma 6.2 in this component in a manner similar to how we did it in the proof of Lemma 6.5.

We now proceed to the formal details.

Proof of Lemma 7.2. As usual, we first focus on proving the existential statement, and at the end we argue how the proof can be turned into an enumeration algorithm.

Let a t -claw be a subdivided claw in which all the tips are at distance exactly t from the center. Note that a t -claw has exactly $3t + 1$ vertices. The first step is to use Lemmas 5.3 and 6.5 to find an induced t -claw in G that is placed robustly with respect to further constructions.

Claim 7.3. *We can either reach conclusion (L2), or find an induced t -claw T in G with a tip w satisfying the following property: there is a connected component D of the graph $G - N[V(T) \setminus \{w\}]$ that is adjacent to w and satisfies $\mathbf{w}(D) \geq (1 - \sigma^{35}) \cdot \mathbf{w}(G)$.*

Proof. Pick any vertex u of G and apply Lemma 5.3 to G , u , and $\alpha = \sigma$. This yields a suitable induced path $Q = (v_0, v_1, \dots, v_k)$ in G , where $v_0 = u$. We adopt the notation from Lemma 5.3 and let D_i be the heaviest connected component of G_i , for $i \in \{0, 1, \dots, k+1\}$. As in the proof of Lemma 6.5, we have that $\mathbf{w}(D_i) > (1 - \sigma) \cdot \mathbf{w}(G)$ for all $i \leq k$ and $\mathbf{w}(D_{k+1}) \leq (1 - \sigma) \cdot \mathbf{w}(G)$. In the same manner as in the proof of Lemma 6.5, we may assume that $\mathbf{w}(D_0) > (1 - \sigma^{39}) \cdot \mathbf{w}(G)$, which in particular entails $k \geq 0$, for otherwise conclusion (L2) can be immediately reached by taking $X = \{u\}$ and the trivial extended strip decomposition of $G_0 = G - u$.

We now define p as the largest index satisfying the following:

$$\mathbf{w}(D_p) > (1 - \sigma^{35}) \cdot \mathbf{w}(G).$$

Since $\mathbf{w}(D_0) > (1 - \sigma^{39}) \cdot \mathbf{w}(G)$ and $\mathbf{w}(D_{k+1}) \leq (1 - \sigma) \cdot \mathbf{w}(G)$, we have that p is well-defined and satisfies $0 \leq p \leq k$.

Consider now the connected graph $G' = G[\{v_p\} \cup V(D_p)]$ and the vertex $u' := v_p$ in it. Since $\mathbf{w}(G') \geq \mathbf{w}(D_p) > \mathbf{w}(G)/2$, we have $\mathbf{w}(N_{G'}[v]) \leq \sigma^{40} \cdot \mathbf{w}(G) \leq \sigma^{16} \cdot \mathbf{w}(G')$ for each vertex v of G' . Hence, we can apply Lemma 6.5 to G' (with the weight function $\mathbf{w}(\cdot)$), vertex u' , and parameters t and σ^2 . This either yields

- (C'1) an induced ($\geq t$)-claw T' in G' with u' being one of its tips; or
- (C'2) a vertex subset $X' \subseteq V(G')$ and an extended strip decomposition (H', η') of $G' - X'$ such that

$$\mathbf{w}(A) \leq (1 - \sigma^{14}) \cdot \mathbf{w}(G') \quad \text{and} \quad \mathbf{w}(X') \leq \sigma^2 \cdot \mathbf{w}(G' - A) \quad \text{for every atom } A \text{ of } (H', \eta').$$

We now argue that in the second case, when conclusion (C'2) is drawn, we can immediately reach conclusion (L2).

Claim 7.4. *If the above application of Lemma 6.5 leads to conclusion (C'2), then conclusion (L2) can be reached.*

Proof. Let us set

$$X = N[v_0, v_1, \dots, v_{p-1}] \cup X'.$$

Then the graph $G - X$ is the disjoint union of $G' - X' - u'$ and all the connected components of G_p different from D_p . Consequently, we can obtain an extended strip decomposition (H, η) of $G - X$ by taking (H', η') , removing u' from it if $u' \notin X'$, and adding, for each component $C \in \text{cc}(G_p)$ different from D_p , a new isolated vertex x_C with $\eta(x_C) = V(C)$. We claim that $(X, (H, \eta))$ satisfies all the properties required by conclusion (L2).

Recall that $\mathbf{w}(G') \geq \mathbf{w}(D_p) > (1 - \sigma^{35}) \cdot \mathbf{w}(G)$. Take any atom A of (H, η) . If A is the vertex set of a connected component C of G_p different from D_p , then we have

$$\mathbf{w}(A) = \mathbf{w}(C) \leq \mathbf{w}(G) - \mathbf{w}(D_p) < \sigma^{35} \cdot \mathbf{w}(G) < (1 - \sigma^{39}) \cdot \mathbf{w}(G), \quad (15)$$

as required. Now assume that A is an atom (H, η) that is also an atom of (H', η') (possibly with u' removed). Then by condition (C'2), we have

$$\mathbf{w}(A) \leq (1 - \sigma^{14}) \cdot \mathbf{w}(G') \leq (1 - \sigma^{14}) \cdot \mathbf{w}(G) < (1 - \sigma^{39}) \cdot \mathbf{w}(G), \quad (16)$$

again as required.

Finally, let us estimate the weight of X . By condition (C'2), for every atom A of (H, η) that is also an atom of (H', η') (possibly with u' removed) we have

$$\begin{aligned} \mathbf{w}(X) &\leq \mathbf{w}(N[v_0, v_1, \dots, v_{p-1}]) + \mathbf{w}(X') \\ &\leq (\mathbf{w}(G) - \mathbf{w}(D_p)) + \sigma^2 \cdot \mathbf{w}(G' - A) \\ &\leq \sigma^{35} \cdot \mathbf{w}(G) + \sigma^2 \cdot \mathbf{w}(G - A). \end{aligned} \tag{17}$$

On the other hand, by (16) we have

$$\mathbf{w}(G - A) = \mathbf{w}(G) - \mathbf{w}(A) \geq \sigma^{14} \cdot \mathbf{w}(G).$$

The above two inequalities together imply that

$$\mathbf{w}(X) \leq \sigma^{21} \cdot \mathbf{w}(G - A) + \sigma^2 \cdot \mathbf{w}(G - A) \leq \sigma \cdot \mathbf{w}(G - A).$$

This establishes the property required in conclusion (L2) for atoms A of (H, η) that are actually atoms of (H', η') , possibly with u' removed. It remains to verify this property for the other atoms, that is, for connected components of G_p different from D_p . Let then C be such a component; then by (15) we have $\mathbf{w}(C) \leq \sigma^{35} \cdot \mathbf{w}(G)$. Hence, by (17) we have

$$\mathbf{w}(X) \leq \sigma^{35} \cdot \mathbf{w}(G) + \sigma^2 \cdot \mathbf{w}(G') \leq 2\sigma^2 \cdot \mathbf{w}(G) \leq \sigma \cdot \mathbf{w}(G - C),$$

and we are done. \square

We continue the proof of Claim 7.3: we are left with considering what happens in case conclusion (C'1) is drawn as a consequence of applying Lemma 6.5. Let c be the center of the constructed $(\geq t)$ -claw T' and let T be the induced t -claw in T' , that is, T the subgraph of T' induced by all the vertices at distance at most t from the center c . We define w as the tip of T that lies on the path connecting u' and c in T' , and we let R be the subpath of this path with endpoints u' and w . See Figure 2 for an illustration. We claim that either we can again reach conclusion (L2), or T and w satisfy the properties from the statement of the claim.




Figure 2: Situation in the proof of Claim 7.3

Let D be the heaviest connected component of $G - N[V(T) \setminus \{w\}]$. Taking $X = N[V(T) \setminus \{w\}]$, we have $\mathbf{w}(X) \leq 3t\sigma^{40} \cdot \mathbf{w}(G) \leq \sigma^{39} \cdot \mathbf{w}(G)$. Therefore, if we had $\mathbf{w}(D) \leq (1 - \sigma^{35}) \cdot \mathbf{w}(G)$, then X together with the trivial extended strip decomposition of $G - X$ would satisfy the requirements of conclusion (L2). Indeed, for every connected component D' of $G - X$ we would have $\mathbf{w}(D') \leq \mathbf{w}(D) \leq (1 - \sigma^{35}) \cdot \mathbf{w}(G)$, which entails $\mathbf{w}(X) \leq \sigma^{39} \cdot \mathbf{w}(G) \leq \sigma^4 \cdot \mathbf{w}(G - D')$. Hence, from now on assume that $\mathbf{w}(D) > (1 - \sigma^{35}) \cdot \mathbf{w}(G)$.

It remains to argue that D is adjacent to w . Let \hat{R} be the path obtained by concatenating the prefix of Q from u to v_p with the path R , and removing w (note that in case $w = v_p$, we also remove

it from \widehat{R}). Observe that \widehat{R} is adjacent to w and is entirely contained in $G - N[V(T) \setminus \{w\}]$, because T' is an induced subdivided claw in $G' = G - (N[v_0, \dots, v_{p-1}] \setminus \{v_p\})$. Therefore, it suffices to argue that $N[\widehat{R}]$ and D intersect.

By the choice of p , every connected component of $G_{p+1} = G - N[v_0, v_1, \dots, v_p]$ has weight at most $(1 - \sigma^{35}) \cdot \mathbf{w}(G)$. On the other hand, we know that D is connected in G and $\mathbf{w}(D) > (1 - \sigma^{35}) \cdot \mathbf{w}(G)$. Therefore, D has to intersect $N[v_0, v_1, \dots, v_p]$. We now have either $w \neq v_p$ and $N[v_0, v_1, \dots, v_p] \subseteq N[\widehat{R}]$ and we are done, or $w = v_p$. In the latter case, either D actually intersects $N[v_0, v_1, \dots, v_{p-1}] \subseteq N[\widehat{R}]$, or D intersects $N[w]$, directly implying that D is adjacent to w . \square

We continue the proof of Lemma 7.2. By applying Claim 7.3, we may assume that we have constructed a suitable t -claw T and its tip w . Let us denote by D the heaviest connected component of $G - N[V(T) \setminus \{w\}]$; then Claim 7.3 ensures us that

$$\mathbf{w}(D) > (1 - \sigma^{35}) \cdot \mathbf{w}(G) \quad \text{and} \quad D \text{ is adjacent to } w.$$

Let us define

$$G'' = G[V(D) \cup \{w\}].$$

Note that G'' is connected.

We first verify that achieving an appropriate variant of conclusion (L2) for G'' is sufficient for our needs.

Claim 7.5. *Suppose we construct a set $X'' \subseteq V(G'')$ and an extended strip decomposition (H'', η'') of $G'' - X''$ with the following property:*

$$\mathbf{w}(A) \leq (1 - \sigma^{35}) \cdot \mathbf{w}(G'') \quad \text{and} \quad \mathbf{w}(X'') \leq \sigma^2 \cdot \mathbf{w}(G'' - A) \quad \text{for every atom } A \text{ of } (H'', \eta'').$$

Then we can reach conclusion (L2).

Proof. Set $X = X'' \cup N[V(T) \setminus \{w\}]$ and observe that the graph $G - X$ can be obtained by taking a disjoint union of the graph $G'' - X'' - w$ and adding all the connected components of $J := G - N[V(T) \setminus \{w\}]$ that are different from D . Hence, we can construct an extended strip decomposition (H, η) of $G - X$ by taking (H'', η'') , removing w if necessary, and adding, for each component $C \in \text{cc}(J)$ different from D , a new isolated vertex x_C with $\eta(x_C) = V(C)$. We claim that $(X, (H, \eta))$ satisfies all the properties required by conclusion (L2).

Recall that $\mathbf{w}(J) \geq \mathbf{w}(D) > (1 - \sigma^{35}) \cdot \mathbf{w}(G)$. Take any atom A of (H, η) . If A is the vertex set of a connected component C of J different from D , then we have

$$\mathbf{w}(A) = \mathbf{w}(C) \leq \mathbf{w}(J) - \mathbf{w}(D) < \sigma^{35} \cdot \mathbf{w}(G) < (1 - \sigma^{39}) \cdot \mathbf{w}(G), \quad (18)$$

as required. Now assume that A is an atom of (H, η) that is also an atom of (H'', η'') (possibly with w removed). Then by the assumption of the claim we have

$$\mathbf{w}(A) \leq (1 - \sigma^{35}) \cdot \mathbf{w}(G'') \leq (1 - \sigma^{39}) \cdot \mathbf{w}(G), \quad (19)$$

again as required.

Finally, let us estimate the weight of X . By the assumption, for every atom A of (H, η) that is also an atom of (H'', η'') (possibly with w removed) we have

$$\begin{aligned} \mathbf{w}(X) &\leq \mathbf{w}(N[V(T) \setminus \{w\}]) + \mathbf{w}(X'') \\ &\leq 3t\sigma^{40} \cdot \mathbf{w}(G) + \sigma^2 \cdot \mathbf{w}(G'' - A) \\ &\leq \sigma^{39} \cdot \mathbf{w}(G) + \sigma^2 \cdot \mathbf{w}(G - A). \end{aligned} \quad (20)$$

On the other hand, by (19) we have

$$\mathbf{w}(G - A) = \mathbf{w}(G) - \mathbf{w}(A) \geq \sigma^{35} \cdot \mathbf{w}(G).$$

The above two inequalities together imply that

$$\mathbf{w}(X) \leq \sigma^4 \cdot \mathbf{w}(G - A) + \sigma^2 \cdot \mathbf{w}(G - A) \leq \sigma \cdot \mathbf{w}(G - A).$$

This establishes the property required in conclusion (L2) for atoms A of (H, η) that are actually atoms of (H'', η'') , possibly with w removed. It remains to verify this property for the other atoms, that is, for connected components of J different from D . Let C be such a component; then by (18) we have $\mathbf{w}(C) \leq \sigma^{35} \cdot \mathbf{w}(G)$. Hence, by (20) we have

$$\mathbf{w}(X) \leq \sigma^{35} \cdot \mathbf{w}(G) + \sigma^2 \cdot \mathbf{w}(G) \leq 2\sigma^2 \cdot \mathbf{w}(G) \leq \sigma \cdot \mathbf{w}(G - C),$$

and we are done. \square

Therefore, from now on we may focus on the graph G'' . The intuition is that T is already one claw of the lobster, and in G'' we try to first construct the second claw, and finally the ‘‘tail’’.

Apply Lemma 5.3 to the graph G'' , vertex w , and $\alpha = \sigma$. This yields a suitable path $P = (y_0, y_1, y_2, \dots, y_\ell)$, where $y_0 = w$. We adopt the notation from the statement of Lemma 5.3 in the following form: $G''_0 = G'' - w$ and $G''_i = G'' - N[y_0, \dots, y_{i-1}]$ for $i \in \{1, \dots, \ell + 1\}$. Moreover, for $i \in \{0, 1, \dots, \ell + 1\}$, let D''_i be the heaviest connected component of G''_i ; then $\mathbf{w}(D''_i) > (1 - \sigma) \cdot \mathbf{w}(G'')$ for $i \leq \ell$ and $\mathbf{w}(D''_{\ell+1}) \leq (1 - \sigma) \mathbf{w}(G'')$. Again, we may assume that $\mathbf{w}(D''_0) > (1 - \sigma^{35}) \cdot \mathbf{w}(G'')$, which in particular entails $\ell \geq 0$: otherwise, the prerequisites of Claim 7.5 can be achieved by taking $X'' = \{w\}$ and the trivial extended strip decomposition of $G'' - X''$, so we can reach conclusion (L2).

Let us define p, q, r as the largest indices satisfying the following:

$$\mathbf{w}(D''_p) > (1 - \sigma^{30}) \cdot \mathbf{w}(G) \quad \text{and} \quad \mathbf{w}(D''_q) > (1 - \sigma^{25}) \cdot \mathbf{w}(G) \quad \text{and} \quad \mathbf{w}(D''_r) > (1 - \sigma^{20}) \cdot \mathbf{w}(G).$$

Since $\mathbf{w}(D''_0) > (1 - \sigma^{35}) \cdot \mathbf{w}(G)$ and $\mathbf{w}(D''_{\ell+1}) \leq (1 - \sigma) \cdot \mathbf{w}(G)$, the indices p, q, r are well-defined and satisfy $0 \leq p \leq q \leq r \leq \ell$.

Similarly as in the proof of Lemma 6.5, let us define the following subpaths of P :

$$R_1 = (y_0, y_1, \dots, y_{p-2}), \quad R_2 = (y_p, y_{p+1}, \dots, y_{q-2}), \quad R_3 = (y_q, y_{q+1}, \dots, y_{r-1}).$$

Observe that paths R_1, R_2, R_3 are pairwise disjoint and non-adjacent. Moreover, the same reasoning as in Claims 6.6 and 6.7 in the proof of Lemma 6.5 easily yields the following; we note that we verify the condition provided to Claim 7.5 in order to reach conclusion (L2).

Claim 7.6. *If we have*

$$\mathbf{w}(N[R_1]) \leq \sigma^{33} \cdot \mathbf{w}(G'') \quad \text{or} \quad \mathbf{w}(N[R_2]) \leq \sigma^{28} \cdot \mathbf{w}(G'') \quad \text{or} \quad \mathbf{w}(N[R_3]) \leq \sigma^{23} \cdot \mathbf{w}(G''),$$

then conclusion (L2) can be obtained. In particular, if the above condition does not hold, then

$$p - 0 > t + 1 \quad \text{and} \quad q - p > t + 1 \quad \text{and} \quad r - q > t + 1.$$

Hence, from now on we assume that the condition stated in Claim 7.6 does not hold, that is:

$$\mathbf{w}(N[R_1]) > \sigma^{33} \cdot \mathbf{w}(G'') \quad \text{and} \quad \mathbf{w}(N[R_2]) > \sigma^{28} \cdot \mathbf{w}(G'') \quad \text{and} \quad \mathbf{w}(N[R_3]) > \sigma^{23} \cdot \mathbf{w}(G''), \quad (21)$$

which in particular implies that $p > t + 1$, $q > p + t + 1$, and $r > q + t + 1$. Since $\mathbf{w}(G'') > \mathbf{w}(G)/2$, assertion (21) in particular implies that

$$\mathbf{w}(N[R_1]) > \sigma^{34} \cdot \mathbf{w}(G) \quad \text{and} \quad \mathbf{w}(N[R_2]) > \sigma^{29} \cdot \mathbf{w}(G) \quad \text{and} \quad \mathbf{w}(N[R_3]) > \sigma^{24} \cdot \mathbf{w}(G). \quad (22)$$

Consider now the connected graph $G''' = G''[\{y_r\} \cup V(D_r'')]$ and the vertex $u''' := y_r$ in it. Since $\mathbf{w}(G''') \geq \mathbf{w}(D_r'') > \mathbf{w}(G'')/2 > \mathbf{w}(G)/4$, we have $\mathbf{w}(N_{G'''}[v]) \leq \sigma^{40} \cdot \mathbf{w}(G) \leq \sigma^{24} \cdot \mathbf{w}(G''')$ for each vertex v of G''' . Hence, we can apply Lemma 6.5 to G''' (with the weight function $\mathbf{w}(\cdot)$), vertex u''' , and parameters t and σ^3 . This either yields

- (C''1) an induced ($\geq t$)-claw S' in G''' with u''' being one of its tips; or
- (C''2) a vertex subset $X''' \subseteq V(G''')$ and an extended strip decomposition (H''', η''') of $G''' - X'''$ such that

$$\mathbf{w}(A) \leq (1 - \sigma^{21}) \cdot \mathbf{w}(G''') \quad \text{and} \quad \mathbf{w}(X''') \leq \sigma^3 \cdot \mathbf{w}(G''' - A) \quad \text{for every atom } A \text{ of } (H''', \eta''').$$

We now argue that in the second case, when conclusion (C''2) is drawn, we can immediately reach conclusion (L2).

Claim 7.7. *If the above application of Lemma 6.5 leads to conclusion (C''2), then we can reach conclusion (L2).*

Proof. Let us define

$$X'' = N_{G''}[y_0, y_1, \dots, y_{r-1}] \cup X'''.$$

Then the graph $G'' - X''$ is the disjoint union of $G''' - X''' - u'''$ and all the connected components of G_r'' different from D_r'' . Consequently, we can obtain an extended strip decomposition (H'', η'') of $G'' - X''$ by taking (H''', η''') , removing u''' from it if $u''' \notin X'''$, and adding, for each component $C \in \text{cc}(G_r'')$ different from D_r'' , a new isolated vertex x_C with $\eta(x_C) = V(C)$. We claim that $(X'', (H'', \eta''))$ satisfies the prerequisites of Claim 7.5, which then entails conclusion (L2)

Recall that $\mathbf{w}(G''') \geq \mathbf{w}(D_r'') > (1 - \sigma^{20}) \cdot \mathbf{w}(G'')$. Take any atom A of (H'', η'') . If A is the vertex set of a connected component C of G_r'' different from D_r'' , then we have

$$\mathbf{w}(A) = \mathbf{w}(C) \leq \mathbf{w}(G'') - \mathbf{w}(D_r'') < \sigma^{20} \cdot \mathbf{w}(G'') < (1 - \sigma^{35}) \cdot \mathbf{w}(G''), \quad (23)$$

as required. Now assume that A is an atom (H'', η'') that is also an atom of (H''', η''') (possibly with u''' removed). Then by condition (C''2), we have

$$\mathbf{w}(A) \leq (1 - \sigma^{21}) \cdot \mathbf{w}(G''') \leq (1 - \sigma^{21}) \cdot \mathbf{w}(G'') \leq (1 - \sigma^{35}) \cdot \mathbf{w}(G''), \quad (24)$$

again as required.

Finally, let us estimate the weight of X'' . By condition (C''2), for every atom A of (H'', η'') that is also an atom of (H''', η''') (possibly with u''' removed), we have

$$\begin{aligned} \mathbf{w}(X'') &\leq \mathbf{w}(N_{G''}[y_0, y_1, \dots, y_{r-1}]) + \mathbf{w}(X''') \\ &\leq (\mathbf{w}(G'') - \mathbf{w}(D_r'')) + \sigma^3 \cdot \mathbf{w}(G''' - A) \\ &\leq \sigma^{20} \cdot \mathbf{w}(G'') + \sigma^3 \cdot \mathbf{w}(G'' - A). \end{aligned} \quad (25)$$

On the other hand, by (24) we have

$$\mathbf{w}(G'' - A) = \mathbf{w}(G'') - \mathbf{w}(A) \geq \sigma^{21} \cdot \mathbf{w}(G'').$$

The above two inequalities together imply that

$$\mathbf{w}(X'') \leq \sigma^6 \cdot \mathbf{w}(G'' - A) + \sigma^3 \cdot \mathbf{w}(G'' - A) \leq \sigma^2 \cdot \mathbf{w}(G'' - A).$$

This establishes the property required in conclusion (L2) for atoms A of (H'', η'') that are actually atoms of (H''', η''') , possibly with u''' removed. It remains to verify this property for the other atoms, that is, for connected components of G_r'' different from D_r'' . Let then C be such a component; then by (23) we have $\mathbf{w}(C) \leq \sigma^{20} \cdot \mathbf{w}(G'')$. Hence, by (25) we have

$$\mathbf{w}(X'') \leq \sigma^{20} \cdot \mathbf{w}(G'') + \sigma^2/2 \cdot \mathbf{w}(G''') \leq 3\sigma^2/4 \cdot \mathbf{w}(G'') \leq \sigma^2 \cdot \mathbf{w}(G'' - C),$$

and we are done. \lrcorner

Hence, from now on we may assume that the application of Lemma 6.5 leads to conclusion (C'1). That is, we constructed an induced $(\geq t)$ -claw S' in G''' with u''' being one of the tips.

Let S be the induced t -claw in S' , that is, S is induced in S' by all vertices at distance at most t from the center of S' . Let v be the tip of S that is the closest in S' to u''' . We now define R'_3 as the path obtained by concatenating: the path R_3 (leading from y_q to y_{r-1}) and the path within S from $u''' = y_r$ to v . Since $S - y_r$ is by construction contained in $G_r = G - N[y_0, \dots, y_{r-1}]$, and P is an induced path in G'' , we infer that paths R_1, R_2, R'_3 are pairwise disjoint and non-adjacent. Moreover, since R_3 is a subpath of R'_3 , by (22) we infer that $\mathbf{w}(N[R'_3]) > \sigma^{24} \cdot \mathbf{w}(G'')$.

Define the following prefix of R_2 :

$$P_2 = (y_p, y_{p+1}, \dots, y_{p+t-1}).$$

We now define the graph

$$G^{(4)} = G - (N[V(S) \setminus \{v\}] \cup N[V(T) \setminus \{w\}] \cup (N(P_2) \setminus y_{p+t-1})).$$

Note that in $G^{(4)}$, the path P_2 is preserved but becomes detached in the following sense: only the endpoint y_{p+t-1} is adjacent to one vertex from the rest of the graph, namely y_{p+t} . Observe that the paths R_1, R_2, R'_3 are also preserved in $G^{(4)}$, and of course they are still disjoint and pairwise non-adjacent.

We now apply Lemma 6.2 to graph $G^{(4)}$ with

$$Z = \{v, w, y_p\}.$$

This either yields an induced tree U in $G^{(4)}$ that contains v, w, y_p , or an extended strip decomposition $(H^{(4)}, \eta^{(4)})$ of $G^{(4)}$ that shatters Z .

In the first case, letting U be inclusion-wise minimal subject to being connected and containing v, w, y_p , we observe that the set

$$V(T) \cup V(U) \cup V(S)$$

induces an $(\geq t)$ -lobster in G . Thus, we reach conclusion (L1); see Figure 3.

We now argue that in the second case we may reach conclusion (L2) by taking

$$X = N[S] \cup N[T] \cup N[P_2],$$

and an extended strip decomposition (H, η) of $G - X$ defined as follows: take $(H^{(4)}, \eta^{(4)})$ and, keeping $H = H^{(4)}$, remove all vertices that belong to X from all the sets in the image of $\eta^{(4)}(\cdot)$.

Since $(H^{(4)}, \eta^{(4)})$ shatters Z in $G^{(4)}$, while R_1, R_2, R'_3 are pairwise disjoint and non-adjacent paths in $G^{(4)}$, each having an endpoint in Z , we infer that every atom A of $(H^{(4)}, \eta^{(4)})$ is disjoint




Figure 3: Final construction of the lobster

with either $N[R_1]$, or $N[R_2]$, or $N[R_3]$. By (22) we infer that $\mathbf{w}(A) \leq (1 - \sigma^{34}) \cdot \mathbf{w}(G)$ for every atom A of $(H^{(4)}, \eta^{(4)})$. Since atoms of (H, η) are subsets of atoms of $(H^{(4)}, \eta^{(4)})$, we also have $\mathbf{w}(A) \leq (1 - \sigma^{34}) \cdot \mathbf{w}(G)$ for every atom A of (H, η) .

Now, observe that since $|X| \leq 7t + 2$, we have

$$\mathbf{w}(X) \leq (7t + 2)\sigma^{40} \cdot \mathbf{w}(G) \leq \sigma^{39} \cdot \mathbf{w}(G).$$

As $\mathbf{w}(A) \leq (1 - \sigma^{34}) \cdot \mathbf{w}(G)$ for every atom A of (H, η) , we also have $\mathbf{w}(G - A) \geq \sigma^{34} \cdot \mathbf{w}(G)$, which in conjunction with the above yields that

$$\mathbf{w}(X) \leq \sigma^5 \cdot \mathbf{w}(G - A) \quad \text{for every atom } A \text{ of } (H, \eta).$$

This means that we have indeed reached conclusion (L2).

For the enumeration statement, it suffices to examine the consecutive steps of the reasoning and replace all steps where we invoke the existential statements of Lemmas 5.3 and 6.5 with iteration over the families obtained by respective enumeration statements. The final family \mathcal{N} consists of all the pairs $(X, (H, \eta))$ that we might have obtained at any point in the reasoning as witnesses for conclusion (L2), for all possible choices of objects from the families provided by Lemmas 5.3 and 6.5. To be more precise, we first invoked Lemma 5.3 followed by Lemma 6.5 in the proof of Claim 7.3, which results in either finding an induced t -claw T or a suitable family \mathcal{N} of size $\mathcal{O}(|V(G)|^6)$. Then we again invoked Lemma 5.3 followed by Lemma 6.5 in the remainder of the proof, which again results in either finding an induced ($\geq t$)-lobster or a suitable family \mathcal{N} of size $\mathcal{O}(|V(G)|^6)$. \square

References

- [1] Tara Abrishami, Maria Chudnovsky, Cemil Dibek, and Pawel Rzazewski. Polynomial-time algorithm for maximum independent set in bounded-degree graphs with no long induced claws. In Joseph (Seffi) Naor and Niv Buchbinder, editors, *Proceedings of the 2022 ACM-SIAM Symposium on Discrete Algorithms, SODA 2022, Virtual Conference / Alexandria, VA, USA, January 9 - 12, 2022*, pages 1448–1470. SIAM, 2022.
- [2] Tara Abrishami, Maria Chudnovsky, Marcin Pilipczuk, Pawel Rzazewski, 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.

- [3] Tara Abrishami, Maria Chudnovsky, Marcin Pilipczuk, and Paweł Rzażewski. Max weight independent set in sparse graphs with no long claws, 2023.
- [4] Gábor Bacsó, Daniel Lokshtanov, Dániel Marx, Marcin Pilipczuk, Zsolt 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.
- [5] Brenda S. Baker. Approximation algorithms for NP-complete problems on planar graphs. *J. ACM*, 41(1):153–180, 1994.
- [6] Andreas Brandstädt and Raffaele Mosca. Maximum weight independent sets for $(P_7, \text{triangle})$ -free graphs in polynomial time. *Discret. Appl. Math.*, 236:57–65, 2018.
- [7] Christoph Brause. A subexponential-time algorithm for the Maximum Independent Set problem in p_t -free graphs. *Discrete Applied Mathematics*, 231:113–118, 2017.
- [8] Maria Chudnovsky, Rose McCarty, Marcin Pilipczuk, Michal Pilipczuk, and Paweł Rzażewski. Sparse induced subgraphs in P_6 -free graphs. *CoRR*, abs/2307.07330, 2023.
- [9] Maria Chudnovsky, Marcin Pilipczuk, Michal 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.
- [10] Maria Chudnovsky, Marcin Pilipczuk, Michał Pilipczuk, and Stéphan Thomassé. On the maximum weight independent set problem in graphs without induced cycles of length at least five. *SIAM J. Discret. Math.*, 34(2):1472–1483, 2020.
- [11] Maria Chudnovsky, Marcin Pilipczuk, Michal 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.
- [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.
- [13] Maria Chudnovsky and Paul D. Seymour. The three-in-a-tree problem. *Combinatorica*, 30(4):387–417, 2010.
- [14] Derek G. Corneil, H. Lerchs, and L. Stewart Burlingham. Complement reducible graphs. *Discrete Applied Mathematics*, 3(3):163–174, 1981.
- [15] Yuri Faenza, Gianpaolo Oriolo, and Gautier Stauffer. An algorithmic decomposition of claw-free graphs leading to an $O(n^3)$ -algorithm for the weighted stable set problem. In Dana Randall, editor, *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2011, San Francisco, California, USA, January 23-25, 2011*, pages 630–646. SIAM, 2011.
- [16] Yuri Faenza, Gianpaolo Oriolo, and Gautier Stauffer. Solving the weighted stable set problem in claw-free graphs via decomposition. *J. ACM*, 61(4):20:1–20:41, 2014.
- [17] M. R. Garey and David S. Johnson. The rectilinear steiner tree problem in NP complete. *SIAM Journal of Applied Mathematics*, 32:826–834, 1977.
- [18] Michael R. Garey and David S. Johnson. *Computers and intractability*, volume 29. WH Freeman, 2002.
- [19] 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.
- [20] Peter Gartland, Daniel Lokshtanov, Tomáš Masarík, Marcin Pilipczuk, Michal Pilipczuk, and Paweł Rzażewski. Maximum weight independent set in graphs with no long claws in quasi-polynomial time. *CoRR*, abs/2305.15738, 2023.
- [21] Peter Gartland, Daniel Lokshtanov, Marcin Pilipczuk, Michal 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.
- [22] Michael U. Gerber, Alain Hertz, and Vadim V. Lozin. Stable sets in two subclasses of banner-free graphs. *Discrete Applied Mathematics*, 132(1-3):121–136, 2003.

- [23] Carla Groenland, Karolina Okrasa, Pawel Rzazewski, Alex D. Scott, Paul D. Seymour, and Sophie Spirkl. H -colouring P_t -free graphs in subexponential time. *Discret. Appl. Math.*, 267:184–189, 2019.
- [24] Andrzej Grzesik, Tereza Klimosová, Marcin Pilipczuk, and Michal Pilipczuk. Polynomial-time algorithm for maximum weight independent set on P_6 -free graphs. *ACM Trans. Algorithms*, 18(1):4:1–4:57, 2022.
- [25] Venkatesan Guruswami and Ali Kemal Sinop. The complexity of finding independent sets in bounded degree (hyper)graphs of low chromatic number. In Dana Randall, editor, *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2011, San Francisco, California, USA, January 23-25, 2011*, pages 1615–1626. SIAM, 2011.
- [26] András Gyárfás. Problems from the world surrounding perfect graphs. *Applicationes Mathematicae*, 19:413–441, 1987.
- [27] Johan Håstad. Clique is hard to approximate within $n^{1-\epsilon}$. *Acta Math.*, 182(1):105–142, 1999.
- [28] Daniel Lokshtanov, 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.
- [29] Vadim V. Lozin and Martin Milanic. A polynomial algorithm to find an independent set of maximum weight in a fork-free graph. *J. Discrete Algorithms*, 6(4):595–604, 2008.
- [30] Vadim V. Lozin and Dieter Rautenbach. Some results on graphs without long induced paths. *Inf. Process. Lett.*, 88(4):167–171, 2003.
- [31] Frédéric Maffray and Lucas Pastor. Maximum weight stable set in (P_7, bull) -free graphs and $(S_{1,2,3}, \text{bull})$ -free graphs. *Discret. Math.*, 341(5):1449–1458, 2018.
- [32] Konrad Majewski, Tomáš Masarík, Jana Novotná, Karolina Okrasa, Marcin Pilipczuk, Pawel Rzazewski, and Marek Sokolowski. Max weight independent set in graphs with no long claws: An analog of the gyárfás’ path argument. In Mikolaj Bojanczyk, Emanuela Merelli, and David P. Woodruff, editors, *49th International Colloquium on Automata, Languages, and Programming, ICALP 2022, July 4-8, 2022, Paris, France*, volume 229 of *LIPICs*, pages 93:1–93:19. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022.
- [33] George J. Minty. On maximal independent sets of vertices in claw-free graphs. *J. Comb. Theory, Ser. B*, 28(3):284–304, 1980.
- [34] Raffaele Mosca. Stable sets of maximum weight in (P_7, banner) -free graphs. *Discrete Mathematics*, 308(1):20–33, 2008.
- [35] Daishin Nakamura and Akihisa Tamura. The generalized stable set problem for claw-free bidirected graphs. In Robert E. Bixby, E. Andrew Boyd, and Roger Z. Ríos-Mercado, editors, *Integer Programming and Combinatorial Optimization, 6th International IPCO Conference, Houston, Texas, USA, June 22-24, 1998, Proceedings*, volume 1412 of *Lecture Notes in Computer Science*, pages 69–83. Springer, 1998.
- [36] Paolo Nobili and Antonio Sassano. An $O(n^2 \log n)$ algorithm for the weighted stable set problem in claw-free graphs. *CoRR*, abs/1501.05775, 2015.
- [37] Paolo Nobili and Antonio Sassano. An $O(m \log n)$ algorithm for the weighted stable set problem in claw-free graphs with $\alpha(g) \leq 3$. *Math. Program.*, 164(1-2):157–165, 2017.
- [38] Gianpaolo Oriolo, Ugo Pietropaoli, and Gautier Stauffer. A new algorithm for the maximum weighted stable set problem in claw-free graphs. In Andrea Lodi, Alessandro Panconesi, and Giovanni Rinaldi, editors, *Integer Programming and Combinatorial Optimization, 13th International Conference, IPCO 2008, Bertinoro, Italy, May 26-28, 2008, Proceedings*, volume 5035 of *Lecture Notes in Computer Science*, pages 77–96. Springer, 2008.
- [39] Marcin Pilipczuk, Michal Pilipczuk, and Pawel Rzazewski. 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.
- [40] Svatopluk Poljak. A note on stable sets and colorings of graphs. *Commentationes Mathematicae Universitatis Carolinae*, 2(15):307–309, 1974.
- [41] Najiba Sbihi. Algorithme de recherche d’un stable de cardinalité maximum dans un graphe sans étoile. *Discrete Mathematics*, 29(1):53–76, 1980. (in French).

- [42] David Zuckerman. Linear degree extractors and the inapproximability of Max Clique and Chromatic Number. *Theory of Computing*, 3(1):103–128, 2007.