

Graphs with No Long Claws: An Improved Bound for the Analog of the Gyárfás' Path Argument

Romain Bourneuf 

Univ. Bordeaux, CNRS, Bordeaux INP, LaBRI, UMR 5800, F-33400, Talence, France

Jana Masaříková 

University of Warsaw, Poland

Wojciech Nadara 

University of Warsaw, Poland

Technical University of Denmark, Kongens Lyngby, Denmark

Marcin Pilipczuk 

University of Warsaw, Poland

Abstract

For a fixed integer $t \geq 1$, a $(t-)$ long claw, denoted $S_{t,t,t}$, is the unique tree with three leaves, each at distance exactly t from the vertex of degree three. Majewski et al. [ICALP 2022, ACM ToCT 2024] proved an analog of the Gyárfás' path argument for $S_{t,t,t}$ -free graphs: given an n -vertex $S_{t,t,t}$ -free graph, one can delete neighborhoods of $\mathcal{O}(\log n)$ vertices so that the remainder admits an extended strip decomposition (an appropriate generalization of partition into connected components) into particles of multiplicatively smaller size. In this work, we refine the argument of Majewski et al. to its arguably final form: we show that a constant number of neighborhoods suffice.

The statement of Majewski et al. is one of the two pillars of a recent quasi-polynomial time algorithm for MAXIMUM WEIGHT INDEPENDENT SET in $S_{t,t,t}$ -free graphs [Gartland et al., STOC 2024]; our work immediately improves the quasi-polynomial function in the running time bound. Furthermore, our result significantly simplifies known polynomial-time algorithms for MAXIMUM WEIGHT INDEPENDENT SET in $S_{t,t,t}$ -free graphs with an additional sparsity assumption such as bounded degree or excluding a fixed biclique as a subgraph.

2012 ACM Subject Classification Mathematics of computing \rightarrow Graph algorithms

Keywords and phrases long-claw-free graphs, extended strip decomposition, maximum weight independent set, Gyárfás' path, three in a tree

Digital Object Identifier 10.4230/LIPIcs.MFCS.2025.28

Related Version Full Version: <https://arxiv.org/abs/2501.13907>

Funding Jana Masaříková: Supported by Polish National Science Centre PRELUDIUM 21 grant number 2022/45/N/ST6/04232.

Wojciech Nadara: Supported by Independent Research Fund Denmark grant 2020-2023 (9131-00044B) "Dynamic Network Analysis" (while being employed in Denmark) and European Union's Horizon 2020 research and innovation programme, grant agreement No. 948057 – BOBR (while being employed in Poland).

Marcin Pilipczuk: Supported by Polish National Science Centre SONATA BIS-12 grant number 2022/46/E/ST6/00143.

Acknowledgements The majority of this work was done during the Sparse Graphs Coalition workshop "Algebraic, extremal, and structural methods and problems in graph colouring" that happened online in February 2024.



© Romain Bourneuf, Jana Masaříková, Wojciech Nadara, and Marcin Pilipczuk; licensed under Creative Commons License CC-BY 4.0

50th International Symposium on Mathematical Foundations of Computer Science (MFCS 2025).

Editors: Paweł Gawrychowski, Filip Mazowiecki, and Michał Skrzypczak; Article No. 28; pp. 28:1–28:16

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1 Introduction

The MAXIMUM WEIGHT INDEPENDENT SET (MWIS) problem is a fundamental problem in combinatorial optimization, where the objective is to find an independent set of vertices in a graph such that the sum of their weights is maximized. The complexity of MWIS significantly varies depending on the structure of the input graph. Some graph classes allow for polynomial-time solutions (e.g., chordal graphs), while in general graphs the problem is NP-hard and hard to approximate within an $n^{1-\varepsilon}$ factor [27, 20]. A crucial research direction involves identifying graph classes where the absence of specific substructures simplifies the MWIS problem. This leads to the study of hereditary graph classes – those closed under vertex deletion. Equivalently, they can be defined by a list of forbidden induced subgraphs. In particular, H -free graphs – graphs that do not contain a specific graph H as an induced subgraph – are of significant interest.

As observed by Alekseev in the 1980s [5, 6], the problem of MWIS remains NP-hard in H -free graphs for most graphs H : MWIS remains NP-hard unless every connected component of H is either a path or a subdivided claw. Moreover, it cannot even be solved in subexponential time, unless the Exponential-Time Hypothesis fails. Thus, significant attention has been devoted to P_t -free graphs (P_t is a path on t vertices) and to $S_{t,t,t}$ -free graphs ($S_{t,t,t}$ is a tree with three leaves at distance t from the vertex of degree 3).

The class of P_4 -free graphs (also known as cographs) is well-structured and many problems, including MWIS, are polynomial-time solvable in it. Similarly, MWIS is polynomial-time solvable in $S_{1,1,1}$ -free graphs [24, 26] (also known as claw-free graphs). For the latter, claw-free graphs are closely related to line graphs, and finding a maximum-weight independent set in a line graph corresponds to finding a maximum-weight matching. The next smallest subdivided claw is $S_{2,1,1}$ (also known as the fork), for which a polynomial algorithm was shown by Alekseev [7] in 2004.

However, moving further is far from being trivial. Currently, $S_{1,2,2}$ is the smallest open case where the polynomiality of MWIS has not been confirmed. In 2014, Lokshantov, Vatschelle, and Villanger [22] showed that MWIS is solvable in polynomial time on P_5 -free graphs using a framework of potential maximal cliques. The result was extended to P_6 -free graphs by Grzesik, Klimošová, Pilipczuk, and Pilipczuk in 2019 [17]. The case of P_7 -free graphs remains open. Several related partial results have been obtained, see for instance [2, 4]. In 2020, in a breakthrough result, Gartland and Lokshantov [13] obtained a quasipolynomial-time algorithm for MWIS in P_t -free graphs. The result later was significantly simplified by Pilipczuk, Pilipczuk, and Rzażewski [25] (but still quasipolynomial time) and generalized to a larger class of problems and to $C_{>t}$ -free graphs (graphs without induced cycle of length more than t) [16].

In $S_{t,t,t}$ -free graphs several partial results were obtained. Abrishami, Chudnovsky, Dibek, and Rzażewski announced a polynomial-time algorithm for MWIS in $S_{t,t,t}$ -free graphs of bounded degree [1] that was later improved in [3] where the same set of authors and Pilipczuk also provide a polynomial-time algorithm for MWIS in graphs excluding any (arbitrary but fixed) graph whose every component is a subdivided claw as an induced subgraph, and any (arbitrary but fixed) biclique as a subgraph, not necessarily induced. However, much less was known for the general cases (even for a specific small t) up until recently. In 2023, another breakthrough, a quasipolynomial algorithm for $S_{t,t,t}$ -free graphs, was presented by Gartland, Lokshantov, Masařík, Pilipczuk, Pilipczuk, and Rzażewski [15]. This progress corroborates the conjecture that MWIS is polynomial-time solvable in H -free graphs for all the open cases left by the original hardness reductions, that is, whenever H is a forest whose every connected component has at most three leaves.

The work [15] is built upon two pillars. The first one is the branching strategy similar to the first P_t -free paper [13]. The second pillar is a work [23] that gives a structural result for $S_{t,t,t}$ -free graphs that is an analog of the Gyárfás' path argument from P_t -free graphs. In this work, we improve the main result of [23] to a form with near optimal parameter dependency. Before we state the result formally, let us discuss the relevant results and techniques.

Firstly, we return to P_t -free graphs. In the 1980s, Gyárfás showed that for every fixed t the class of P_t -free graphs is χ -bounded [18, 19]. Bacsó, Lokshantov, Marx, Pilipczuk, Tuza, and van Leeuwen [8] gave a subexponential-time algorithm for MWIS in P_t -free graphs, using the following important corollary of the Gyárfás' path argument.

► **Theorem 1.** *Given an n -vertex graph G with nonnegative vertex weights, one can in polynomial time find an induced path Q in G such that every connected component of $G - N[V(Q)]$ has weight at most half of the total weight of $V(G)$.*

Observe that in P_t -free graphs, Q contains at most $t - 1$ vertices. In particular, having connected components of multiplicatively smaller size after removing the neighborhood of a constant number of vertices became very useful in the recursive approach.

For the progress in general $S_{t,t,t}$ -free graphs, the paper [23] shows that the notion of an *extended strip decomposition*, developed by Chudnovsky and Seymour in their project to understand claw-free graphs [11], became useful as an analog of the Gyárfás' path argument. For a formal definition of extended strip decomposition, we refer to Section 2. For an intuition to develop, we remark that in an extended strip decomposition of a graph, we can distinguish *particles* (specific induced subgraphs of the graph). The crucial property is that a maximum weight independent set of the whole graph can be combined from solutions obtained independently in individual particles ([10]). Thus, “particles” of an extended strip decomposition can be viewed as analogs of connected components. Particles of multiplicatively smaller size in $S_{t,t,t}$ -free graphs have been obtained in [23].

► **Theorem 2 ([23]).** *Given an n -vertex graph G with nonnegative vertex weights and an integer $t \geq 1$, one can in polynomial time either:*

- *output an induced copy of $S_{t,t,t}$ in G , or*
- *output a set \mathcal{P} consisting of at most $11 \log n + 6$ induced paths in G , each of length at most $t + 1$, and a rigid extended strip decomposition of $G - N[\bigcup_{P \in \mathcal{P}} V(P)]$ whose every particle has weight at most half of the total weight of $V(G)$.*

Using this structural result, the authors provided a subexponential-time algorithm with running time $2^{\mathcal{O}(\sqrt{n} \log n)}$ and a quasipolynomial-time approximation scheme with running time $2^{\mathcal{O}(\varepsilon^{-1} \log^5 n)}$. Both results improved and simplified the state-of-the-art knowledge by Chudnovsky, Pilipczuk, Pilipczuk, and Thomassé [10] from 2020.

This result has already found several applications besides simplifying the analysis of the quasipolynomial-time algorithm [3]. Most importantly, it became as a black box one of the pillars of the quasipolynomial-time algorithm for MWIS in $S_{t,t,t}$ -free graphs [15]. There, the key structural ingredient for the algorithm is the following. For a fixed integer $t \geq 1$, given a graph G , one can in polynomial time find either an induced $S_{t,t,t}$ in G , or a balanced separator consisting of $\mathcal{O}(\log |V(G)|)$ vertex neighborhoods in G , or an extended strip decomposition of G with each particle of weight multiplicatively smaller than the weight of G . We remark that the third output gives an extended strip decomposition of the whole graph, in contrast with Theorem 2.

1.1 Our Contribution

We address an open question from [23] that asked to eliminate the logarithmic factor in output 2 of Theorem 2. Indeed, it turns out that removing neighborhoods of a constant number of *short* paths is sufficient to obtain an extended strip decomposition with each particle of weight multiplicatively smaller than the total weight of the graph. Our main contribution is the following theorem.

► **Theorem 3.** *Given a graph $G = (V, E)$ with nonnegative vertex weights $\mathfrak{w} : V \rightarrow \mathbb{N}$ and an integer $t \geq 1$, one can in polynomial time either:*

- *output an induced copy of $S_{t,t,t}$ in G , or*
- *output a set \mathcal{S} of size at most $3t+11$, and a rigid extended strip decomposition of $G - N[\mathcal{S}]$ whose every particle has weight at most $\mathfrak{w}(V)/2$.*

Observe that removing some (unbounded) number of vertices is sometimes needed to get a nontrivial extended strip decomposition: adding to any $S_{t,t,t}$ -free graph a large number of universal vertices keeps the graph $S_{t,t,t}$ -free (for $t \geq 2$), but destroys any useful extended strip decomposition of the graph. Hence, in other words, Theorem 3 improves one of the two pillars of the quasipolynomial algorithm for MWIS in $S_{t,t,t}$ -free graphs of [15] to an arguably optimal bound of removing only a constant number of neighborhoods.

1.2 Technical overview

Similarly, as in [23], our starting point is the Gyárfás’ path, refining the selection of splitting vertices. If its length is within a small multiple of t , we even have a stronger result of components of half of the weight of G . Otherwise, we identify specific vertices of the Gyárfás’ path and query the three-in-a-tree theorem, originally developed by Chudnovsky and Seymour [12] with improved running time by Lai, Lu, and Thorup [21], on them. The output of the three-in-a-tree theorem is either an induced tree connecting the given vertices or an extended strip decomposition of G . By the choice of the vertices we make sure the second output applies, otherwise an induced copy of $S_{t,t,t}$ is found. The main difference from [23] lies in a more involved choice of the vertices with a better understanding of the structure of the extended strip decomposition. Among others, we use the obvious but powerful property of the last two vertices of the Gyárfás’ path. After removing the neighborhood of the Gyárfás’ path up to its second-to-last vertex, the largest connected component is still “big” (has weight larger than half of the weight of G). However, after removing the neighborhood of the last vertex, its weight suddenly drops below this threshold. This observation lets us look more closely at the interaction of the “big” particle and the “big” connected component.

1.3 Algorithmic implications

Theorem 3 improves the quasi-polynomial running time bound of the algorithm for MWIS in $S_{t,t,t}$ -free graphs of [15], but does not yield a polynomial-time algorithm. The running time bound of the algorithm of [15] is $n^{\mathcal{O}_t(\log^{16} n)}$, where \mathcal{O}_t hides factors depending on t . Few, but not all, of the $\log n$ factors in the exponent come from Theorem 2; using Theorem 3 instead removes them from the exponent. We refrain from re-doing here a detailed analysis of the branching algorithm of [15] to specify exactly what the improvement is; this analysis will appear in the journal version of [15].

In the specific case of sparse $S_{t,t,t}$ -free graphs considered by [3], Theorem 3 significantly simplifies the situation. For the bounded degree case, the polynomial-time algorithm is now immediate: apply Theorem 3, exhaustively branch on $N[\mathcal{S}]$, and recurse on the particles of

the extended strip decomposition. The algorithm for the more general case of excluding any fixed graph whose every component is a subdivided claw as an induced subgraph, and any fixed biclique as a subgraph can also be significantly simplified (but not to half a sentence as the previous case).

We remark that a very similar result to Theorem 3 has been independently obtained by Chudnovsky, Codsi, Milanič, Lokshtanov, and Sivashankar [9], via a slightly different proof.

2 Preliminaries

2.1 Notation

All graphs are simple. Let G be a graph. For $X \subseteq V(G)$, by $G[X]$ we denote the subgraph of G induced by X , i.e., $(X, \{uv \in E(G) : u, v \in X\})$. If the graph G is clear from the context, we will often identify induced subgraphs with their vertex sets.

For a vertex v , by $N_G(v)$ we denote the set of neighbors of v , and by $N_G[v]$ we denote the set $N_G(v) \cup \{v\}$. For a set $X \subseteq V(G)$, we also define $N_G(X) := \bigcup_{v \in X} N_G(v) - X$, and $N_G[X] = N_G(X) \cup X$. If it does not lead to confusion, we omit the subscript and write simply $N(\cdot)$ and $N[\cdot]$. By $\binom{V(G)}{3}$ we mean the set of all subsets of $V(G)$ of size 3. By $T(G)$, we denote the set of all triangles in G . Note that $T(G) \subseteq \binom{V(G)}{3}$. Similarly to writing $xy \in E(G)$, we will write $xyz \in T(G)$ to indicate that $G[\{x, y, z\}]$ is a triangle. We say that two sets $X, Y \subseteq V(G)$ *touch* if $X \cap Y \neq \emptyset$ or there is an edge with one end in X and another in Y . For a family \mathcal{Q} of sets, by $\bigcup \mathcal{Q}$ we denote $\bigcup_{Q \in \mathcal{Q}} Q$. For a function $\mathfrak{w} : V \rightarrow \mathbb{Z}_{\geq 0}$ and subset $V' \subseteq V$, we denote $\mathfrak{w}(V') := \sum_{v \in V'} \mathfrak{w}(v)$.

A graph G is said to be F -free for some graph F if no induced subgraph of G is isomorphic to F . For integers $t, a, b, c > 0$, by P_t we denote the path on t vertices, and by $S_{a,b,c}$, we denote the tree with three leaves at respective distance a, b , and c from the unique vertex of degree 3 of the tree.

2.2 Extended strip decompositions

See [23, Figure 1] or [14, Figure 3] for an illustration of the definition below.

► **Definition 4** (Extended strip decomposition, [12]). *An extended strip decomposition of a graph G is a pair (H, η) that consists of:*

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

which satisfy the following properties:

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

For an edge $xy \in E(H)$, the sets $\eta(xy, x)$ and $\eta(xy, y)$ are called the *interfaces* of the edge xy . An extended strip decomposition (H, η) is *rigid* if (i) for every $xy \in E(H)$ it holds that $\eta(xy, x) \neq \emptyset$, and (ii) for every $x \in V(H)$ such that x is an isolated vertex it holds that $\eta(x) \neq \emptyset$. Observe that if we *restrict* η to $V' \subset V(G)$, i.e. we keep in η only vertices of V' , then (H, η) after the restriction remains an extended strip decomposition, but it might not be rigid anymore.

We say that a vertex $v \in V(G)$ is *peripheral* in (H, η) if there is a degree-one vertex x of H , such that $\eta(xy, x) = \{v\}$, where y is the (unique) neighbor of x in H . Note that we can also assume $\eta(x) = \emptyset$, as otherwise all vertices from $\eta(x)$ can be moved to $\eta(xy)$. For a set $Z \subseteq V(G)$, we say that (H, η) is an *extended strip decomposition of (G, Z)* if H has $|Z|$ degree-one vertices and each vertex of Z is peripheral in (H, η) . The following theorem by Chudnovsky and Seymour [12] is a slight strengthening of their celebrated solution of the famous *three-in-a-tree* problem.

► **Theorem 5** (Chudnovsky, Seymour [12, Section 6]). *Let G be an n -vertex graph and $Z \subseteq V(G)$ with $|Z| \geq 2$. There is an algorithm that runs in time $\mathcal{O}(n^4)$ and returns one of the following:*

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

The running time was improved to $\mathcal{O}(m \log^2 n)$ (where $m = |E|$) in [21].

► **Definition 6.** *Let (H, η) be an extended strip decomposition of a graph G . We distinguish the following special subsets of $V(G)$:*

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

We refer to these sets as *particles* and discern between their respective types.

A vertex particle A_x is *trivial* if x is an isolated vertex in H . Similarly, an extended strip decomposition (H, η) is *trivial* if H is an edgeless graph. Given a weight function $\mathfrak{w} : V(G) \rightarrow \mathbb{N}$, we say that a particle is *small* if its weight is at most $\frac{\mathfrak{w}(V(G))}{2}$, and an extended strip decomposition (H, η) of G is *refined* if all its particles are small.

3 Properties of extended strip decompositions and the proof of Theorem 3

3.1 Overview of the proof

Before the actual proof, we show several properties of extended strip decompositions. For the purposes of this overview, the most important is Lemma 11, which states that if (H, η) is an extended strip decomposition of a graph G then any induced path P (in G) between peripheral vertices is entirely contained in edge particles. Moreover, in each edge particle, P has at most one vertex in each interface.

Let us now describe the key ideas of the proof of Theorem 3. We start by considering a Gyárfás' path Q in G . If Q is short, we can set $\mathcal{S} = V(Q)$ and return a trivial extended strip decomposition of $G - N[\mathcal{S}]$ satisfying the requirements of Theorem 3, so assume that

Q is long. First, we mark specific vertices of Q on which we will query the three-in-a-tree theorem: The first vertex of Q (call it x) and the vertices at distance $t + 3$ (call it y) and $t + 1$ (call it z) from the end of Q other than x . Before querying the three-in-a-tree theorem, we remove the neighborhood outside Q of three subpaths of Q : the subpaths of Q starting at x and z and taking the next t vertices in Q , respectively, and the subpath taking the t vertices of Q preceding y . We also remove the vertex of Q between y and z . See Figure 1 for an illustration. Observe that this selection ensures that the three-in-a-tree theorem either outputs an induced copy of $S_{t,t,t}$ (which is a desired output of Theorem 3) or an extended strip decomposition of the graph minus the mentioned neighborhoods.

We take a closer look at the returned extended strip decomposition. If all its particles are small, we can simply return it, so assume it has a big particle A . Let us assume that A is a full edge particle, i.e., $A = A_{pq}^{pq}$, the other cases being easy to handle. We distinguish two cases depending on whether the subpath Q_1 of Q from x to y intersects A or not. If Q_1 is disjoint from A , the only vertices of Q_1 that can have a neighbor in A are vertices in an interface of a neighboring edge of pq in H . Lemma 11 implies that there can be at most four of them. Finally, using that Q is a Gyárfás' path, we observe that after removing the neighborhoods of the three previously mentioned subpaths, of these four vertices and of four other well-chosen vertices, every remaining connected component has weight at most half of the total weight of G . If Q_1 intersects A , denote by Q_2 be the subpath of Q from z to the end of Q . The fact that Q is induced and Lemma 11 together imply that Q_2 does not have any neighbor in A . From the property of Gyárfás' path, letting ℓ denote the last vertex of Q , we have that $G - N[V(Q) - \ell]$ has a big connected component C but $G - N[V(Q)]$ does not. Therefore, ℓ has a neighbor in $V(C)$. Two big parts must intersect, so $V(C)$ and A intersect. It then follows from Lemma 11 that $V(C) \subseteq A$. This contradicts the fact that Q_2 does not have any neighbor in A .

3.2 Properties of extended strip decompositions

We start the proof of Theorem 3 with several properties of extended strip decompositions.

► **Lemma 7.** *Let (H, η) be an extended strip decomposition of a graph G . Let $A = A_{xy}^{xy}$ be a full edge particle. For $z \in \{x, y\}$, let F_z be the set of edges of H incident with z . Let*

$$B = \bigcup_{z \in \{x, y\}} \bigcup_{f \in F_z - \{xy\}} \eta(f, z).$$

Then, $N_G(A) \subseteq B$. Furthermore, if (H, η) is rigid, then $N_G(A) = B$.

Proof. The claim $N_G(A) \subseteq B$ follows from the third property of the definition of an extended strip decomposition: A_{xy}^{xy} consists of $\eta(xy)$, $\eta(x)$, $\eta(y)$, and $\eta(xyz)$ for all triangles xyz , and vertices in those sets are allowed only to have neighbors outside A_{xy}^{xy} in sets $\eta(f, x)$ for $f \in F_x - \{xy\}$ and $\eta(f, y)$ for $f \in F_y - \{xy\}$.

If (H, η) is rigid, then $\eta(xy, x)$ is nonempty and any vertex of $\eta(xy, x)$ is complete to $\bigcup_{f \in F_x - \{xy\}} \eta(f, x)$; a symmetric claim holds for $\eta(xy, y)$. This proves that $N_G(A) = B$ in this case. ◀

► **Lemma 8.** *Let (H, η) be a rigid extended strip decomposition of a graph G . Let $A = A_{xy}^{xy}$ be a full edge particle. There exists a set $X_A \subseteq A$ of size at most 2 such that $N_G(A) \subseteq N_G(X_A)$.*

Proof. Since (H, η) is rigid, there exist $v_x \in \eta(xy, x)$ and $v_y \in \eta(xy, y)$. Set $X_A = \{v_x, v_y\}$. Let $u \in N_G(A)$. Since $u \notin A$ is adjacent to a vertex in $A = A_{xy}^{xy}$, by Lemma 7, there exists $z \in \{x, y\}$ and an edge $f \neq xy$ incident to z such that $u \in \eta(f, z)$. Then, u is adjacent to $v_z \in \{v_x, v_y\}$ so $u \in N_G(X_A)$. ◀

We note that the above observation was mentioned in [23, Observation 8], we include it here for completeness.

As discussed above, if (H, η) is an extended strip decomposition of a graph G , if G' is an induced subgraph of G and η' is the restriction of η to $V(G')$ then (H, η') is an extended strip decomposition of G' , but it is not necessarily rigid anymore. The following lemma (whose proof we postpone to Section 3.4 in order not to disturb the flow of arguments here) shows how to modify it to obtain a rigid extended strip decomposition, while preserving the same bounds for the maximum weight of a particle.

► **Lemma 9.** *Let $G = (V, E)$ be a graph with nonnegative vertex weights $\mathfrak{w} : V \rightarrow \mathbb{N}$, $w \geq \mathfrak{w}(V)/2$ and (H, η) be an extended strip decomposition of G whose every particle has weight at most w . Then, one can in polynomial time compute either a rigid extended strip decomposition of G whose every particle has weight at most w , or a set $X \subseteq V(G)$ of size at most 2 such that every connected component of $G - N[X]$ has weight at most w .*

We now look at induced paths in G whose endpoints are peripheral in (H, η) , with particular emphasis on where their vertices appear in (H, η) . The following two lemmata characterize how an induced path (in G) between peripheral vertices can traverse an extended strip decomposition.

► **Lemma 10.** *Let (H, η) be an extended strip decomposition of a graph G and $z \in V(G)$ be a peripheral vertex in (H, η) . Let $Q = (z = x_0, x_1, \dots, x_{k-1}, x_k)$ be an induced path in G with endpoint z . Let $A = A_{pq}^{pq}$ be a full edge particle such that $z \notin \eta(pq)$. If $V(Q) \cap N[A] \neq \emptyset$, there exists an edge $f \neq pq$ and $r \in \{p, q\}$, $r \in f$ such that $\eta(f, r) \cap V(Q) \neq \emptyset$.*

Proof. Since $z \notin \eta(pq)$ and z is peripheral in (H, η) then $z \notin A$. Suppose $V(Q) \cap N[A] \neq \emptyset$ and let i be minimum such that $x_i \in N[A]$.

If $i = 0$, then $z \in N[A]$. Let f be the unique edge of H such that $z \in \eta(f)$. Then, $f \neq pq$ by assumption. Furthermore, since $z \in N[A] - A = N(A)$, by Lemma 7, there exists $r \in \{p, q\}$ such that $z \in \eta(f, r)$.

Otherwise, by minimality of i , we have $x_{i-1} \notin N[A]$, which implies $x_i \notin A$. Thus, $x_i \in N[A] - A = N(A)$ so there exists an edge f and $r \in \{p, q\}$ such that $x_i \in \eta(f, r)$. Since $x_i \notin A$ then $f \neq pq$, which concludes the proof. ◀

► **Lemma 11.** *Let (H, η) be an extended strip decomposition of a graph G and $x, y \in V(G)$ be two distinct peripheral vertices in (H, η) . Let $Q = (x = x_0, x_1, \dots, x_{k-1}, x_k = y)$ be an induced path in G with endpoints x and y . Then, for every edge $e = ab \in E(H)$, it holds that $|V(Q) \cap \eta(e, a)| \leq 1$, and $V(Q) \subseteq \bigcup_{e \in E(H)} \eta(e)$.*

Proof. First, we prove that if there exists $v \in V(H)$ and $x_i \in V(Q)$ such that $x_i \in \eta(v)$, then there exist $p < i < q$ and an edge e incident to v such that $x_p, x_q \in \eta(e, v)$. Indeed, suppose that $x_i \in V(Q) \cap \eta(v)$ for some $v \in V(H)$. Let $p < i$ be maximum such that $x_p \notin \eta(v)$ (this is well-defined since $x_0 = x \notin \eta(v)$ as x is peripheral in (H, η)). Observe that $x_{p+1} \in \eta(v)$. Since $x_p x_{p+1} \in E(G)$, there exists an edge $e \in E(H)$ incident to v such that $x_p \in \eta(e, v)$. Similarly, there exists $q > i$ and an edge $f \in E(H)$ incident to v such that $x_q \in \eta(f, v)$. Since Q is an induced path and $q > p + 1$, there is no edge $x_p x_q$ in G and thus $e = f$. Thus, $x_p, x_q \in \eta(e, v)$. With the same argument, it also holds that if there exists $t \in T(H)$ and $x_i \in V(Q)$ such that $x_i \in \eta(t)$ then there exist $p < i < q$, an edge $e \in E(t)$ and a vertex $v \in V(t)$ such that $x_p, x_q \in \eta(e, v)$. Therefore, it suffices to prove that $|V(Q) \cap \eta(e, a)| \leq 1$ for every edge $e = ab \in E(H)$ to prove that $V(Q) \subseteq \bigcup_{e \in E(H)} \eta(e)$.

For $uv \in E(H)$, let $L(uv, u)$ and $R(uv, u)$ be minimum and maximum values of i such that $x_i \in \eta(uv, u)$, or \perp if no such values exist. We claim that if $|V(Q) \cap \eta(uv, u)| \geq 2$, then $L(uv, v) \leq L(uv, u)$, $R(uv, u) \leq R(uv, v)$ (and, in particular, $|V(Q) \cap \eta(uv, v)| \geq 2$) and prove that as follows. Let $L(uv, u) = i < j = R(uv, u)$. If $x_j \in \eta(uv, v)$, then $R(uv, v) \geq j = R(uv, u)$. If $x_j \notin \eta(uv, v)$, then consider the set $A = \eta(u) \cup \eta(uv) - \eta(uv, v)$ and the smallest $j' > j$ such that $x_{j'} \notin A$. Such j' is well defined as $y \notin A$ since $|\eta(uv, u)| \geq 2$. If $j' > j + 1$, then $x_{j'-1} \in A$ from the minimality of j' , whereas if $j' = j + 1$, then $x_{j'-1} = x_j \in A$. So, in any case $x_{j'-1} \in A$ and we deduce that $x_{j'} \in N(A) \subseteq \eta(uv, v) \cup \bigcup_{uw \in E(H), w \neq v} \eta(uw, u)$. However, if $x_{j'} \in \eta(uw, u)$ for some $w \neq v$, then $x_i x_{j'} \in E(G)$, which is a contradiction as Q is an induced path and $j' > j > i$. Hence, $x_{j'} \in \eta(uv, v) \Rightarrow R(uv, v) \geq j' > R(uv, u)$. In either case, we have that $R(uv, v) \geq R(uv, u)$. Similarly, we conclude $L(uv, v) \leq L(uv, u)$, which proves our auxiliary claim. However, as $|V(Q) \cap \eta(uv, v)| \geq 2$, the roles of u and v can be reversed to show that $L(uv, u) \leq L(uv, v)$ and $R(uv, v) \leq R(uv, u)$, hence we conclude that $L(uv, u) = L(uv, v)$ and $R(uv, u) = R(uv, v)$.

By contradiction, suppose that there exists an edge $e = ab \in E(H)$ which satisfies $|V(Q) \cap \eta(e, a)| \geq 2$. Let $L(ab, b) = i$ and $R(ab, b) = j$. We have that $i < j$, $L(ab, a) = i$ and $R(ab, a) = j$, hence $x_i, x_j \in \eta(ab, a) \cap \eta(ab, b)$. Let us consider the set $A = \eta(a) \cup \eta(b) \cup \eta(ab) \cup \bigcup_{abc \in T(H)} \eta(abc)$ and the smallest $j' > j$ such that $x_{j'} \notin A$. Such j' is well defined as $|\eta(ab, a)|, |\eta(ab, b)| \geq 2$, so $y \notin A$. We have that $x_{j'} \in N(A)$, however $N(A) \subseteq \bigcup_{ac \in E(H), c \neq b} \eta(ac, a) \cup \bigcup_{bc \in E(H), c \neq a} \eta(bc, b)$, therefore $N(A)$ is complete to $\eta(ab, a) \cap \eta(ab, b)$. Hence $x_i x_{j'} \in E(H)$, which is a contradiction since Q is an induced path and $j' > j > i$, which concludes the proof. \blacktriangleleft

► **Corollary 12.** *Let (H, η) be an extended strip decomposition of a graph G and $x, y \in V(G)$ be two distinct peripheral vertices in (H, η) . Let $Q = (x = x_0, x_1, \dots, x_{k-1}, x_k = y)$ be an induced path in G with endpoints x and y . For every $e = ab \in E(H)$ such that $V(Q) \cap \eta(e) \neq \emptyset$, then $G[V(Q) \cap \eta(e)]$ is a path with endpoints in $\eta(e, a)$ and $\eta(e, b)$ and all internal vertices in $\eta(e) - (\eta(e, a) \cup \eta(e, b))$.*

Proof. Since $x, y \in V(G)$ are peripheral in (H, η) , there exist edges $e_x = u_x v_x \in E(H)$, $e_y = u_y v_y \in E(H)$ with v_x, v_y being degree-one vertices, such that $\eta(e_x, v_x) = \{x\}$ and $\eta(e_y, v_y) = \{y\}$.

Let $G_E = G[\bigcup_{e \in E(H)} \eta(e)]$. By Lemma 11 we have that $Q \subseteq G_E$.

Let $e = ab \in E(H)$ such that $V(Q) \cap \eta(e) \neq \emptyset$. Let i be minimum such that $x_i \in \eta(e)$ and j be maximum such that $x_j \in \eta(e)$. If $i = 0$ then $x_i = x$ so $e = e_x$ and up to renaming a and b we have $x_i \in \eta(e, a)$. If $i > 0$ then $x_{i-1} \notin \eta(e)$ and $x_{i-1} x_i \in E(G)$ so up to renaming a and b we have $x_i \in \eta(e, a)$. Similarly, we have that $x_j \in \eta(e, a) \cup \eta(e, b)$.

Let us first consider the cases where $i = j$.

- If $i = j = 0$. Then, $e = e_x = u_x v_x$ and $x_i = x \in \eta(e_x, v_x)$. Furthermore, we have $x_{i+1} = x_{j+1} \notin \eta(e_x)$ by maximality of j , and v_x has degree 1 in H and $\eta(v_x) = \emptyset$ since x is peripheral so $x \in \eta(e, u_x)$.
- If $i = j = k$. This is similar to the previous case.
- If $0 < i = j < k$. In this case we need to prove that $x_i \in \eta(e, b)$. Let us assume by contradiction that this is false. We have that $x_{i-1}, x_{i+1} \in N(x_i) \cap V(G_E) - \eta(e)$. As $x_i \in \eta(e, a) - \eta(e, b)$, we conclude that $x_{i-1}, x_{i+1} \in \bigcup_{ac \in E(H), c \neq b} \eta(ac, a)$. Let us write then that $x_{i-1} \in \eta(ac_1, a)$ and $x_{i+1} \in \eta(ac_2, a)$. If $c_1 = c_2$, then $|V(Q) \cap \eta(ac_1, a)| \geq 2$, which is a contradiction with Lemma 11. If $c_1 \neq c_2$, then $x_{i-1} x_{i+1} \in E(G)$, which is a contradiction with Q being an induced path.



■ **Figure 1** Notations from the proof of Theorem 3. The sets Q_x, Q_y and Q_z all have size $t + 1$.

From now on, let us assume that $i < j$. Since $|V(Q) \cap \eta(e, a)|, |V(Q) \cap \eta(e, b)| \leq 1$ by Lemma 11, we have that $x_i \in \eta(e, a) - \eta(e, b)$, $x_j \in \eta(e, b) - \eta(e, a)$ and no other vertices of Q belong to $\eta(e, a) \cup \eta(e, b)$. Moreover, $\eta(e, a) \cup \eta(e, b)$ separates $\eta(e) - (\eta(e, a) \cup \eta(e, b))$ from $V(G_E) - \eta(e)$, so proving that $x_{i+1} \in \eta(e)$ would imply that all internal vertices of the path $x_i x_{i+1} \dots x_{j-1} x_j$ belong to $\eta(e) - (\eta(e, a) \cup \eta(e, b))$, as desired, so this is what we are going to prove in the remaining cases.

- If $i = 0 < j$. Then $e = e_x = u_x v_x$ and $x_i = x \in \eta(e_x, v_x)$, so we have that $a = v_x, b = u_x$. We recall $x_i \in \eta(e, a) - \eta(e, b)$ and $x_j \in \eta(e, b) - \eta(e, a)$. As a is a degree one vertex and $x_0 \notin \eta(e, b)$, we clearly have that $x_1 \in \eta(e)$, which proves our claim.
- If $i < j = k$. This is similar to the previous case.
- If $0 < i < j < k$. We recall that $x_i \in \eta(e, a) - \eta(e, b)$ and $x_j \in \eta(e, b) - \eta(e, a)$. As $x_i \in \eta(e, a) - \eta(e, b)$, we can use the same reasoning as in the $0 < i = j < k$ case to argue that $x_{i+1} \in \eta(e)$ (which operated under the same assumption that $0 < i$ and $x_i \in \eta(e, a) - \eta(e, b)$, but excluded the possibility that $x_{i+1} \in \eta(e)$). That again completes the argument. ◀

3.3 Proof of Theorem 3

For convenience, we restate the Theorem 3

► **Theorem 3.** *Given a graph $G = (V, E)$ with nonnegative vertex weights $\mathbf{w} : V \rightarrow \mathbb{N}$ and an integer $t \geq 1$, one can in polynomial time either:*

- *output an induced copy of $S_{t,t,t}$ in G , or*
- *output a set \mathcal{S} of size at most $3t + 11$, and a rigid extended strip decomposition of $G - N[\mathcal{S}]$ whose every particle has weight at most $\mathbf{w}(V)/2$.*

Proof. Let Q be a minimal induced path in G given by Theorem 1, that is, a path such that every connected component of $G - N[V(Q)]$ has weight at most $\mathbf{w}(V)/2$ and Q is minimal with this property. If $|V(Q)| \leq 3t + 11$, set $\mathcal{S} = V(Q)$ and observe that every connected component of $G - N[\mathcal{S}]$ has weight at most $\mathbf{w}(V)/2$. Thus, $G - N[\mathcal{S}]$ has a trivial rigid extended strip decomposition whose every particle has weight at most $\mathbf{w}(V)/2$. Assume now that $|V(Q)| > 3t + 11$. Let ℓ be the last vertex of Q . Let Q_2 be the subpath of Q induced by the last $t + 2$ vertices of Q , and denote by z the first vertex of Q_2 . Let z' be the predecessor of z in Q , and y the predecessor of z' in Q . Let x be the first vertex of Q , and Q_1 be the subpath of Q from x to y . Observe that $|V(Q_1)| > 2t + 8$. Let Q_x be the set containing x and the first t successors of x , Q_y be the set containing y and the first t predecessors of y , and Q_z be the set containing z and the first t successors of z .

Let $Y = Q_x \cup Q_y \cup Q_z$. Then, $|Y| = 3t + 3$. Consider the graph G' defined as $G' = G - (N(Y) - V(Q_1) - V(Q_2))$. Note that every vertex of the induced path Q is a vertex of G' , except z' . Let $Z = \{x, y, z\}$. By Theorem 5, we can find in polynomial time either an induced subtree of G' containing all three elements of Z , or a rigid extended strip decomposition of (G', Z) .

We claim that the first outcome is impossible. Observe that for every $\xi \in \{x, y, z\}$, the vertex ξ is of degree 1 in G' and every internal vertex of Q_ξ is of degree 2 in G' . Hence, any induced subtree of G' containing all three elements of Z contains an induced $S_{t,t,t}$ since it must contain x, y , and z as leaves and the three induced paths $G[Q_x], G[Q_y], G[Q_z]$, each containing t edges.

Therefore, we can assume that we find a rigid extended strip decomposition (H, η) of (G', Z) . Note that the size of H is polynomial in the size of G since the algorithm of Theorem 5 runs in polynomial time. Suppose first that every particle of (H, η) has weight at most $\mathfrak{w}(V)/2$. Applying Lemma 9 to $G - N_G[Y]$ and the restriction of (H, η) , with $w = \mathfrak{w}(V)/2$, we can find in polynomial time either a rigid extended strip decomposition of $G - N_G[Y]$ whose every particle has weight at most $\mathfrak{w}(V)/2$, in which case we are done setting $\mathcal{S} = Y$, or a set $X \subseteq V(G - N_G[Y])$ of size at most 2 such that every connected component of $(G - N_G[Y]) - N_{G - N_G[Y]}[X] = G - N_G[Y \cup X]$ has weight at most $\mathfrak{w}(V)/2$, in which case we are done setting $\mathcal{S} = Y \cup X$.

Suppose now that there exists a particle A such that $\mathfrak{w}(A) > \mathfrak{w}(V)/2$. Without loss of generality, we can assume that whenever $p \in V(H)$ is isolated then $G[\eta(p)]$ is connected. Consider such a $p \in V(H)$. Since Q is a path between peripheral vertices then we have $N_G[V(Q)] \cap \eta(p) = \emptyset$ so $G[\eta(p)]$ is a connected component of $G - N_G[V(Q)]$, which implies $\mathfrak{w}(\eta(p)) \leq \mathfrak{w}(V)/2$. Since every nontrivial particle is contained in a full edge particle we can assume without loss of generality that A is a full edge particle: $A = A_{pq}^{pq}$ for some edge $pq \in E(H)$. Applying Lemma 11 to the induced path Q_1 in G' , we obtain that $V(Q_1) \subseteq \bigcup_{e \in E(H)} \eta(e)$.

▷ Claim 13. $V(Q_1) \cap \eta(pq) = \emptyset$.

Proof. By contradiction suppose that $V(Q_1) \cap \eta(pq) \neq \emptyset$.

By Corollary 12, $G'[V(Q_1) \cap \eta(pq)]$ is a path with endpoints $v_p \in \eta(pq, p)$ and $v_q \in \eta(pq, q)$. Therefore, $\eta(pq, p) \neq \{z\}$ and $\eta(pq, q) \neq \{z\}$ so $z \notin \eta(pq)$. Applying Lemma 10 to the induced path Q_2 with endpoint z in G' , we get that if $V(Q_2) \cap N_{G'}[A] \neq \emptyset$, there exists an edge $f \neq pq$, $r \in \{p, q\}$ and $z'' \in \eta(f, r) \cap V(Q_2)$. However, we would then have $z''v_r \in E(G)$ with $z'' \in V(Q_2)$ and $v_r \in V(Q_1)$, contradicting that Q is an induced path in G . Therefore, $V(Q_2) \cap N_{G'}[A] = \emptyset$.

By minimality of Q , $G - N_G[V(Q) - \{\ell\}]$ has a connected component C such that $\mathfrak{w}(V(C)) > \mathfrak{w}(V)/2$. Therefore, ℓ has a neighbor in $V(C)$, and $V(C) \cap A \neq \emptyset$ since $\mathfrak{w}(V(C)), \mathfrak{w}(A) > \mathfrak{w}(V)/2$. Observe also that $C \subseteq V(G) - N_G[V(Q) - \{\ell\}] \subseteq V(G')$. If $u \in V(C) \cap N_{G'}(A)$ then there exists $f \neq pq \in E(H)$ and $r \in \{p, q\}$ such that $u \in \eta(f, r)$. However, this implies that $uv_r \in E(G') \subseteq E(G)$ and therefore $u \in N_G[V(Q_1)]$, which is impossible since $V(C) \cap N_G[V(Q_1)] = \emptyset$. Thus, $V(C) \cap N_{G'}(A) = \emptyset$, and therefore $V(C) \cap A \neq \emptyset$ implies $V(C) \subseteq A$ since C is connected in G' . Hence, ℓ has a neighbor in A , which contradicts $V(Q_2) \cap N_{G'}[A] = \emptyset$.

This finishes the proof of the claim. ◁

From Claim 13 and Lemma 11 we infer that $V(Q_1) \cap A = \emptyset$ as x and y are peripheral. Let $X = N_{G'}[A] \cap V(Q_1)$. If $u \in X$, there is a vertex $a \in A$ such that $ua \in E(G)$, and thus there exists $r \in \{p, q\}$ and an edge f incident to r such that $u \in \eta(f, r)$. However, by Lemma 11, we have $|V(Q_1) \cap \eta(ab, a)| \leq 1$ for every edge $ab \in E(H)$. Furthermore, if $u_1 \in \eta(f_1, r), u_2 \in \eta(f_2, r), u_3 \in \eta(f_3, r)$ for some pairwise distinct edges f_1, f_2, f_3 then $G[\{u_1, u_2, u_3\}]$ induces a triangle. Since Q_1 is an induced path, it follows that $|X| \leq 4$. By Lemma 8, there exists a set X_A of size at most 2 such that $N_{G'}(A) \subseteq N_{G'}(X_A)$.

▷ **Claim 14.** Every connected component of $G'' := G - N_G[Y \cup X \cup X_A \cup \{\ell, z'\}]$ has weight at most $\mathfrak{w}(V)/2$.

Proof. Consider a connected component C of G'' and assume by contradiction that it satisfies $\mathfrak{w}(V(C)) > \mathfrak{w}(V)/2$. Then, $V(C) \cap A \neq \emptyset$. If $v \in V(G'') \cap N_G[A]$ then $v \in V(G')$ and $v \notin N_G[X_A] \supseteq N_{G'}(X_A) \supseteq N_{G'}(A)$. Therefore, $v \in N_{G'}[A] - N_{G'}(A) = A$. Since $V(C) \cap A \neq \emptyset$ and C is connected in G'' then $V(C) \subseteq A$. However, every connected component of $G - N_G[V(Q)]$ has weight at most $\mathfrak{w}(V)/2$. Thus, C is not contained in any connected component of $G - N_G[V(Q)]$ and therefore there exists some vertex $q \in V(Q) - N_G[Y \cup X \cup X_A \cup \{\ell, z'\}]$ which is in $N_G[V(C)]$. Therefore, $q \in V(Q_1) \cap N_G[V(C)] \subseteq V(Q_1) \cap N_G[A] = X$, a contradiction. This finishes the proof of the claim. ◀

Thanks to Claim 14, by setting $\mathcal{S} = Y \cup X \cup X_A \cup \{\ell, z'\}$, we have $|\mathcal{S}| \leq 3t + 11$ and $G - N_G[\mathcal{S}]$ has a trivial rigid extended strip decomposition whose every particle has weight at most $\mathfrak{w}(V)/2$.

Finally, note that this proof is constructive and immediately yields a polynomial time algorithm. ◀

3.4 Proof of Lemma 9

► **Lemma 9.** Let $G = (V, E)$ be a graph with nonnegative vertex weights $\mathfrak{w} : V \rightarrow \mathbb{N}$, $w \geq \mathfrak{w}(V)/2$ and (H, η) be an extended strip decomposition of G whose every particle has weight at most w . Then, one can in polynomial time compute either a rigid extended strip decomposition of G whose every particle has weight at most w , or a set $X \subseteq V(G)$ of size at most 2 such that every connected component of $G - N[X]$ has weight at most w .

Proof. For the purposes of the proof, we introduce the notion of *relaxed extended strip decomposition*, which is similar to that of extended strip decomposition, except that we have a set $\eta(xyz)$ for every triple of vertices xyz in H and not only for every triangle xyz in H . We require the same properties as in an extended strip decomposition when replacing $T(H)$ by $\binom{V(H)}{3}$. The particles of a relaxed extended strip decomposition are the vertex particles, the full edge particles and the triple particles, each defined similarly to the corresponding particle in an extended strip decomposition when replacing $T(H)$ by $\binom{V(H)}{3}$. Note that every extended strip decomposition naturally corresponds to a relaxed extended strip decomposition and that a relaxed extended strip decomposition such that $\eta(xyz) = \emptyset$ whenever $xyz \notin T(H)$ naturally yields an extended strip decomposition.

We view (H, η) as a relaxed extended strip decomposition of G , and we start by ensuring that every edge $xy \in E(H)$ satisfies $\eta(xy, x) \neq \emptyset$. Let $xy \in E(H)$ be such that $\eta(xy, x) = \emptyset$. If $\eta(xy, y) = \emptyset$ too, let H' be the graph obtained from H by removing the edge xy and adding an isolated vertex v . For every $o \in V(H) \cup E(H) \cup \binom{V(H)}{3} - \{xy\}$, set $\eta'(o) = \eta(o)$, and set $\eta'(v) = \eta(xy)$ and $\eta'(vu_1u_2) = \emptyset$ for every newly created triple of vertices. If $\eta(xy, y) \neq \emptyset$, let H' be the graph obtained from H by removing the edge xy and adding a vertex z adjacent only to y . For every $o \in V(H) \cup E(H) \cup \binom{V(H)}{3} - \{xy\}$, set $\eta'(o) = \eta(o)$, and set $\eta'(zy) = \eta'(zy, z) = \eta(xy)$, $\eta'(zy, y) = \eta(xy, y)$, $\eta'(z) = \emptyset$ and $\eta'(zu_1u_2) = \emptyset$ for every newly created triple of vertices.

Note that in both cases (H', η') is still a relaxed extended strip decomposition of G whose every particle has weight at most w . When going from (H, η) to (H', η') , the number of edges with an empty interface decreases strictly. Therefore, by repeating this operation, we eventually get that every edge $xy \in E(H)$ satisfies $\eta(xy, x) \neq \emptyset$.

Suppose now that there exists an isolated vertex x of H and a triple $xyz \in \binom{V(H)}{3}$ such that $\eta(xyz) \neq \emptyset$. If $yz \in E(H)$, let $H' = H$ and for every $o \in V(H) \cup E(H) \cup \binom{V(H)}{3} - \{xyz, yz\}$, set $\eta'(o) = \eta(o)$, set $\eta'(xyz) = \emptyset$ and $\eta'(yz) = \eta(yz) \cup \eta(xyz)$, $\eta'(yz, y) = \eta(yz, y)$, and $\eta'(yz, z) = \eta(yz, z)$. If $yz \notin E(H)$, let H' be the graph obtained from H by adding an isolated vertex v . For every $o \in V(H) \cup E(H) \cup \binom{V(H)}{3} - \{xyz\}$, set $\eta'(o) = \eta(o)$, and set $\eta'(v) = \eta(xyz)$, $\eta'(xyz) = \emptyset$ and $\eta'(vu_1u_2) = \emptyset$ for every newly created triple of vertices.

Observe that in both cases (H', η') is still a relaxed extended strip decomposition of G whose every particle has weight at most w . When going from (H, η) to (H', η') , we create no edge with an empty interface and the number of triples $xyz \in \binom{V(H)}{3}$ such that $\eta(xyz) \neq \emptyset$ and x is isolated in H decreases strictly. Therefore, by repeating this operation, we eventually get that no edge has an empty interface and for every triple $xyz \in \binom{V(H)}{3}$ such that x is isolated, we have $\eta(xyz) = \emptyset$.

Therefore, if x is an isolated vertex of H such that $\eta(x) = \emptyset$ then after removing x from H , what we obtain is again an extended strip decomposition of G with the same properties. By iteratively removing such vertices, we eventually get a rigid relaxed extended strip decomposition whose every particle has weight at most w .

We now show how to convert this rigid relaxed extended strip decomposition (H, η) into a rigid extended strip decomposition whose every particle has weight at most w . If there is no $xyz \in \binom{V(H)}{3}$ such that $\eta(xyz) \neq \emptyset$ and $xyz \notin T(H)$ then restricting η to $V(H) \cup E(H) \cup T(H)$ gives the desired rigid extended strip decomposition. Suppose now that there exists a triple $xyz \in \binom{V(H)}{3}$ such that $\eta(xyz) \neq \emptyset$ and $xyz \notin T(H)$. We consider several cases.

- If $\eta(xyz)$ is not adjacent to $V(G) - \eta(xyz)$, let H' be the graph obtained from H by adding an isolated vertex v . For every $o \in V(H) \cup E(H) \cup \binom{V(H)}{3} - \{xyz\}$, set $\eta'(o) = \eta(o)$, and set $\eta'(v) = \eta(xyz)$, $\eta'(xyz) = \emptyset$ and $\eta'(vu_1u_2) = \emptyset$ for every newly created triple of vertices.
- If $\eta(xyz)$ is adjacent to only one of $\{\eta(xy), \eta(xz), \eta(yz)\}$, say $\eta(xy)$, let $H' = H$ and for every $o \in V(H) \cup E(H) \cup \binom{V(H)}{3} - \{xyz, xy\}$, set $\eta'(o) = \eta(o)$, and set $\eta'(xy) = \eta(xy) \cup \eta(xyz)$ and $\eta'(xyz) = \emptyset$. Importantly, we set $\eta'(xy, x) = \eta(xy, x)$ and $\eta'(xy, y) = \eta(xy, y)$.
- Otherwise, since $xyz \notin T(H)$, only two of $\{xy, xz, yz\}$ are edges of H , say xy and xz . Then, $\eta(xyz)$ is adjacent to both $\eta(xy)$ and $\eta(xz)$. Again, we set $H' = H$, for every $o \in V(H) \cup E(H) \cup \binom{V(H)}{3} - \{xyz, x\}$, set $\eta'(o) = \eta(o)$, and set $\eta'(x) = \eta(x) \cup \eta(xyz)$ and $\eta'(xyz) = \emptyset$.

In each case, (H', η') is a rigid relaxed extended strip decomposition of G and the number of triples $xyz \in \binom{V(H)}{3}$ such that $\eta(xyz) \neq \emptyset$ and $xyz \notin T(H)$ decreases strictly. Thus, by iterating this operation we eventually obtain a rigid relaxed extended strip decomposition of G with no $xyz \in \binom{V(H)}{3}$ such that $\eta(xyz) \neq \emptyset$ and $xyz \notin T(H)$. If this relaxed extended strip decomposition has no particle of weight strictly greater than w then as before we can turn it into the desired rigid extended strip decomposition of G .

Therefore, we can assume that one of the above operations created a particle of weight strictly greater than w . Note that in the first case, we do not increase the weight of any pre-existing particle, and the only non-empty particle we create is $A_v = \eta'(v) = \eta(xyz)$ which has weight at most w by assumption. Similarly, in the second case, we do not increase the weight of any particle. Therefore, we created a particle of weight strictly greater than w going from some rigid relaxed extended strip decomposition (H, η) to another rigid relaxed extended strip decomposition (H', η') using the third case. In this case, all particles keep the same weight except for A_{xyz} which becomes empty, A_x , and the full edge particles A_{xa}^{xa}

($a \notin \{y, z\}$). Since after the modification we have $A_x \subseteq A_{xy}^{xy}$ and the weight of A_{xy}^{xy} did not change then after the modification we still have $\mathfrak{w}(A_x) \leq w$. Thus, there exists a full edge particle $A' = A_{xa}^{xa}$ of (H', η') such that $\mathfrak{w}(A') > w$, and $a \notin \{y, z\}$.

Since (H, η) is rigid, as in Lemma 8 there exists a set $X_{A'}$ of size at most 2 such that $N_G(A') \subseteq N_G(X_{A'})$. Therefore, in $G - N_G[X_{A'}]$, the remaining vertices of A' are disconnected from $V - A'$. Furthermore, $V - A'$ has weight at most $\mathfrak{w}(V) - \mathfrak{w}(A') \leq w$. Let $A = A_{xa}^{xa}$ be the corresponding particle of (H, η) . Then, $A' = A \cup \eta(xyz)$ but since $a \notin \{y, z\}$ then $\eta(xyz)$ is not adjacent to A in G . Hence, for every connected component C of $G - N_G[X_{A'}]$, we have either $V(C) \subseteq V - A'$ or $V(C) \subseteq \eta(xyz)$ or $V(C) \subseteq A$. In all three cases we have $\mathfrak{w}(V(C)) \leq w$ since (H, η) is refined.

Finally, note that this proof is algorithmic and can be implemented in polynomial time (in the sizes of H and G). Indeed, every modification we perform on the relaxed extended strip decomposition can be done in polynomial time, and we can find the next one in polynomial time. Therefore, we just need to argue that the size of the relaxed extended strip decomposition remains within a polynomial factor of the size of the initial extended strip decomposition. First, we perform a polynomial number of modifications to ensure that no edge has an empty interface, and each such modification adds at most one vertex. Then, to ensure that there is no triple xyz such that x is isolated and $\eta(xyz) \neq \emptyset$ we again perform a polynomial number of modifications, each adding at most one vertex. Then, we remove isolated vertices, which cannot increase the size of the relaxed extended strip decomposition. Finally, we perform a polynomial (in the size of the current relaxed extended strip decomposition) number of operations to turn the relaxed extended strip decomposition into a proper extended strip decomposition. ◀

References

- 1 Tara Abrishami, Maria Chudnovsky, Cemil Dibek, and Paweł Rzażewski. Polynomial-time algorithm for maximum independent set in bounded-degree graphs with no long induced claws. In Niv Buchbinder Joseph (Seffi) Naor, editor, *Proceedings of the 2022 ACM-SIAM Symposium on Discrete Algorithms, SODA 2022, Virtual Conference, January 9-12, 2022*, pages 1448–1470. SIAM, 2022. doi:10.1137/1.9781611977073.61.
- 2 Tara Abrishami, Maria Chudnovsky, Cemil Dibek, Stéphan Thomassé, Nicolas Trotignon, and Kristina Vušković. Graphs with polynomially many minimal separators. *J. Comb. Theory, Ser. B*, 152:248–280, 2022. doi:10.1016/j.jctb.2021.10.003.
- 3 Tara Abrishami, Maria Chudnovsky, Marcin Pilipczuk, and Paweł Rzażewski. Max weight independent set in sparse graphs with no long claws. In Olaf Beyersdorff, Mamadou Moustapha Kanté, Orna Kupferman, and Daniel Lokshtanov, editors, *41st International Symposium on Theoretical Aspects of Computer Science, STACS 2024, March 12-14, 2024, Clermont-Ferrand, France*, volume 289 of *LIPICs*, pages 4:1–4:15. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024. doi:10.4230/LIPICs.STACS.2024.4.
- 4 Tara Abrishami, Maria Chudnovsky, Marcin Pilipczuk, Paweł Rzażewski, and Paul D. Seymour. Induced subgraphs of bounded treewidth and the container method. In Dániel Marx, editor, *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms, SODA 2021, Virtual Conference, January 10 - 13, 2021*, pages 1948–1964. SIAM, 2021. doi:10.1137/1.9781611976465.116.
- 5 Vladimir E. Alekseev. The effect of local constraints on the complexity of determination of the graph independence number. *Combinatorial-algebraic methods in applied mathematics*, pages 3–13, 1982.
- 6 Vladimir E. Alekseev. On easy and hard hereditary classes of graphs with respect to the independent set problem. *Discret. Appl. Math.*, 132(1-3):17–26, 2003. doi:10.1016/S0166-218X(03)00387-1.

- 7 Vladimir E. Alekseev. Polynomial algorithm for finding the largest independent sets in graphs without forks. *Discrete Applied Mathematics*, 135(1):3–16, 2004. Russian Translations II. doi:10.1016/S0166-218X(02)00290-1.
- 8 Gábor Bacsó, Daniel Lokshtanov, Dániel Marx, Marcin Pilipczuk, Zolt Tuza, and Erik Jan van Leeuwen. Subexponential-time algorithms for Maximum Independent Set in P_t -free and broom-free graphs. *Algorithmica*, 81(2):421–438, 2019. doi:10.1007/s00453-018-0479-5.
- 9 Maria Chudnovsky, Julien Codsi, Martin Milanič, Daniel Lokshtanov, and Varun Sivashankar. Tree independence number V. Walls and claws, 2025.
- 10 Maria Chudnovsky, Marcin Pilipczuk, Michał Pilipczuk, and Stéphan Thomassé. Quasi-polynomial time approximation schemes for the Maximum Weight Independent Set Problem in H -free graphs. *SIAM Journal on Computing*, 53(1):47–86, 2024. doi:10.1137/20M1333778.
- 11 Maria Chudnovsky and Paul D. Seymour. The structure of claw-free graphs. In Bridget S. Webb, editor, *Surveys in Combinatorics, 2005 [invited lectures from the Twentieth British Combinatorial Conference, Durham, UK, July 2005]*, volume 327 of *London Mathematical Society Lecture Note Series*, pages 153–171. Cambridge University Press, 2005. doi:10.1017/cbo9780511734885.008.
- 12 Maria Chudnovsky and Paul D. Seymour. The three-in-a-tree problem. *Comb.*, 30(4):387–417, 2010. doi:10.1007/s00493-010-2334-4.
- 13 Peter Gartland and Daniel Lokshtanov. Independent set on P_k -free graphs in quasi-polynomial time. In Sandy Irani, editor, *61st IEEE Annual Symposium on Foundations of Computer Science, FOCS 2020, Durham, NC, USA, November 16-19, 2020*, pages 613–624. IEEE, 2020. doi:10.1109/FOCS46700.2020.00063.
- 14 Peter Gartland, Daniel Lokshtanov, Tomáš Masařík, Marcin Pilipczuk, Michał Pilipczuk, and Paweł Rzażewski. Maximum weight independent set in graphs with no long claws in quasi-polynomial time. *CoRR*, abs/2305.15738, 2023. doi:10.48550/arXiv.2305.15738.
- 15 Peter Gartland, Daniel Lokshtanov, Tomáš Masařík, Marcin Pilipczuk, Michał Pilipczuk, and Paweł Rzażewski. Maximum weight independent set in graphs with no long claws in quasi-polynomial time. In Bojan Mohar, Igor Shinkar, and Ryan O'Donnell, editors, *Proceedings of the 56th Annual ACM Symposium on Theory of Computing, STOC 2024, Vancouver, BC, Canada, June 24-28, 2024*, pages 683–691. ACM, 2024. doi:10.1145/3618260.3649791.
- 16 Peter Gartland, Daniel Lokshtanov, Marcin Pilipczuk, Michał Pilipczuk, and Paweł Rzażewski. Finding large induced sparse subgraphs in $C_{>t}$ -free graphs in quasipolynomial time. In Samir Khuller and Virginia Vassilevska Williams, editors, *STOC '21: 53rd Annual ACM SIGACT Symposium on Theory of Computing, Virtual Event, Italy, June 21-25, 2021*, pages 330–341. ACM, 2021. doi:10.1145/3406325.3451034.
- 17 Andrzej Grzesik, Tereza Klimošová, Marcin Pilipczuk, and Michał Pilipczuk. Polynomial-time algorithm for Maximum Weight Independent Set on P_6 -free graphs. In Timothy M. Chan, editor, *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019, San Diego, California, USA, January 6-9, 2019*, pages 1257–1271. SIAM, 2019. doi:10.1137/1.9781611975482.77.
- 18 András Gyárfás. On Ramsey covering-numbers. In *Infinite and finite sets (Colloq., Keszthely, 1973; dedicated to P. Erdős on his 60th birthday)*, Vol. II, number 10 in Colloq. Math. Soc. Janos Bolyai, pages 801–816. North-Holland, Amsterdam, 1975.
- 19 András Gyárfás. Problems from the world surrounding perfect graphs. In *Proceedings of the International Conference on Combinatorial Analysis and its Applications, (Pokrzywna, 1985)*, number 19 in Zastos. Mat., pages 413–441, 1987. doi:10.4064/am-19-3-4-413-441.
- 20 Johan Håstad. Clique is hard to approximate within $n^{1-\epsilon}$. *Acta Math.*, 182(1):105–142, 1999. doi:10.1007/BF02392825.
- 21 Kai-Yuan Lai, Hsueh-I Lu, and Mikkel Thorup. Three-in-a-tree in near linear time. In Konstantin Makarychev, Yury Makarychev, Madhur Tulsiani, Gautam Kamath, and Julia Chuzhoy, editors, *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of*

- Computing, STOC 2020, Chicago, IL, USA, June 22-26, 2020*, pages 1279–1292. ACM, 2020. doi:10.1145/3357713.3384235.
- 22 Daniel Lokshantov, Martin Vatshelle, and Yngve Villanger. Independent set in P_5 -free graphs in polynomial time. In Chandra Chekuri, editor, *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2014, Portland, Oregon, USA, January 5-7, 2014*, pages 570–581. SIAM, 2014. doi:10.1137/1.9781611973402.43.
 - 23 Konrad Majewski, Tomáš Masařík, Jana Masaříková, Karolina Okrasa, Marcin Pilipczuk, Paweł Rzażewski, and Marek Sokołowski. Max Weight Independent Set in Graphs with No Long Claws: An Analog of the Gyárfás’ Path Argument. *The ACM Transactions on Computation Theory*, 16(2), March 2024. doi:10.1145/3636422.
 - 24 George J. Minty. On maximal independent sets of vertices in claw-free graphs. *Journal of Combinatorial Theory, Series B*, 28(3):284–304, 1980. doi:10.1016/0095-8956(80)90074-X.
 - 25 Marcin Pilipczuk, Michał Pilipczuk, and Paweł Rzażewski. Quasi-polynomial-time algorithm for independent set in P_t -free graphs via shrinking the space of induced paths. In Hung Viet Le and Valerie King, editors, *4th Symposium on Simplicity in Algorithms, SOSA 2021, Virtual Conference, January 11-12, 2021*, pages 204–209. SIAM, 2021. doi:10.1137/1.9781611976496.23.
 - 26 Najiba Sbihi. Algorithme de recherche d’un stable de cardinalité maximum dans un graphe sans étoile. *Discrete Mathematics*, 29(1):53–76, 1980. doi:10.1016/0012-365X(90)90287-R.
 - 27 David Zuckerman. Linear degree extractors and the inapproximability of Max Clique and Chromatic Number. *Theory of Computing*, 3(1):103–128, 2007. doi:10.4086/toc.2007.v003a006.