QPTAS for MWIS and finding large sparse induced subgraphs in graphs with few independent long holes

Édouard Bonnet
Univ Lyon, CNRS, ENS de Lyon, Université Claude Bernard Lyon 1, LIP UMR5668, France

Jadwiga Czyżewska
University of Warsaw, Poland

Tomáš Masařík
University of Warsaw, Poland

Marcin Pilipczuk
University of Warsaw, Poland

Paweł Rzążewski
Warsaw University of Technology & University of Warsaw, Poland

Abstract

We present a quasipolynomial-time approximation scheme (QPTAS) for the Maximum Independent Set (MWIS) in graphs with a bounded number of pairwise vertex-disjoint and non-adjacent long induced cycles. More formally, for every fixed s and t, we show a QPTAS for MWIS in graphs that exclude sC_t as an induced minor. Combining this with known results, we obtain a QPTAS for the problem of finding a largest induced subgraph of bounded treewidth with given hereditary property definable in Counting Monadic Second Order Logic, in the same classes of graphs.

This is a step towards a conjecture of Gartland and Lokshtanov which asserts that for any planar graph H, graphs that exclude H as an induced minor admit a polynomial-time algorithm for the latter problem. This conjecture is notoriously open and even its weaker variants are confirmed only for very restricted graphs H.

Funding *Jadwiga Czyżewska*: Supported by Polish National Science Centre SONATA BIS-12 grant number 2022/46/E/ST6/00143.
Tomáš Masařík: Supported by the Polish National Science Centre SONATA-17 grant number 2021/43/D/ST6/03312.
Marcin Pilipczuk: Supported by Polish National Science Centre SONATA BIS-12 grant number 2022/46/E/ST6/00143.
Paweł Rzążewski: Supported by the National Science Centre grant 2024/54/E/ST6/00094.

Acknowledgements The project was initiated at the Structural Graph Theory workshop STWOR in September 2023. This workshop was a part of STRUG: Structural Graph Theory Bootcamp, funded by the "Excellence initiative – research university (2020–2026)" of University of Warsaw. We are grateful to Benjamin Bergougnoux, Linda Cook, and Marek Sokołowski for stimulating discussions on the topic.

1 Introduction

In the Max Independent Set (MIS) problem, one is asked, given a graph G, for a largest *independent set*, i.e., a set of pairwise nonadjacent vertices in G. In the weighted variant of the problem, Max Weight Independent Set (MWIS), the input graph has vertex weights and we ask for an independent set of maximum weight. MIS (and thus MWIS) is a "canonical" hard problem—it is one of Karp's 21 NP-hard problems [21], W[1]-hard (with respect to the solution size) [12] and notoriously hard to approximate [20], even in the parameterized setting [8, 22]. Thus, a natural question to ask when dealing with such a hard problem is as follows: *For what classes of input graphs is the MWIS problem tractable?*

Typically, this question is studied for *hereditary* graph classes, i.e., classes closed under vertex deletion. Each such class can be equivalently characterized by specifying minimal induced subgraphs that do not belong to the class. For graphs G and H, we say that G is *H-free* if it does not contain H as an induced subgraph.

The complexity study of MIS and MWIS in restricted graph classes is among the most active research directions in algorithmic graph theory. Let us list some relevant results, focusing on the case of H-free graphs, for a single graph H. Already in the early 1980s, Alekseev [2] observed that classic NP-hardness reductions imply that MIS remains NP-hard in H-free graphs for many graphs H. First, MIS is NP-hard in subcubic graphs, and thus in H-free graphs whenever H has a vertex of degree at least 4 [14]. The second reduction involves the so-called *Poljak's subdivision trick* [30]: subdividing an edge of a graph twice yields a new graph where the size of a maximum independent set increases by exactly one. Consequently, for a fixed graph H, we can start with an arbitrary instance of MIS and subdivide each edge $2|V(H)|$ times. This way we obtain an equivalent instance that is H-free, whenever H has a cycle or two vertices of degree at least three in the same connected component.

Combining these two observations, we obtain that MIS (and thus, MWIS) remains NP-hard in H-free graphs, unless every component of H is a subcubic tree with at most one vertex of degree 3; let us call the family of such forests $\mathcal{S}$. Let us highlight that in particular $\mathcal{S}$ contains all paths and, more generally, linear forests (i.e., forests of paths).

We do not know any NP-hardness result for MWIS in H-free graphs when $H \in \mathcal{S}$. On the other hand, polynomial-time algorithms are known only for small graphs $H \in \mathcal{S}$ [3, 7, 19, 23–25, 31]. However, general belief in the community is that all cases not excluded by standard NP-hardness reductions mentioned above MWIS is polynomial-time-solvable. This belief is supported by the existence of a *quasipolynomial-time* algorithm. Indeed, for every $H \in \mathcal{S}$, the MWIS problem restricted to n-vertex H-free graphs can be solved in time $n^{\mathcal{O}(\log^{16} n)}$ [17]. Note that this is a strong indication that no $H \in \mathcal{S}$ defines an NP-hard case of MWIS, as otherwise every problem in NP can be solved in quasipolynomial time.

Let us have a closer look at the case of P_t-free graphs, i.e., graphs that exclude a t-vertex path as an induced subgraph. It turns out that in these classes we cannot only solve MWIS in quasipolynomial time [16, 29] (and in polynomial time for $t \leqslant 6$ [19, 23]), but the algorithms can actually be extended to a rich family of problems defined as follows [1, 10, 18]. For an integer r and a CMSO_2 sentence[1] ψ, we define the $(\mathsf{tw} \leqslant r, \psi)$-MWIS problem as follows (here, MWIS stands for *maximum-weight induced subgraph*).

$(\mathsf{tw} \leqslant r, \psi)$-MWIS

Input: A graph G equipped with a weight function $\mathsf{w}\colon V(G) \to \mathbb{R}_{\geqslant 0}$.

Task: Find a set $S \subseteq V(G)$, such that
- $G[S] \models \psi$,
- $\mathsf{tw}(G[S]) \leqslant r$, and
- S is of maximum weight subject to the conditions above,

or conclude that no such set exists.

[1] CMSO_2 is a logic where one can use vertex, edge, and (vertex or edge) set variables, check vertex-edge incidence, quantify over variables, and apply counting predicates modulo fixed integers. See [12] for a formal introduction.

Notable special cases of the $(\mathsf{tw} \leqslant r, \psi)$-MWIS problem are MWIS, Feedback Vertex Set (equivalently, Max Induced Forest), and Even Cycle Transversal (equivalently, Max Induced Odd Cactus).

Interestingly, the quasipolynomial-time algorithm for $(\mathsf{tw} \leqslant r, \psi)$-MWIS in P_t-free graphs can be extended to the class of $C_{\geqslant t}$-free graphs: ones that do not contain an induced cycle with at least t vertices. Note that every P_t-free graph is $C_{\geqslant t+1}$-free. We emphasize that $C_{\geqslant t}$-free graphs form the class defined by an infinite minimal family of forbidden induced subgraphs. Furthermore, we cannot hope for tractability of $(\mathsf{tw} \leqslant r, \psi)$-MWIS in H-free graphs, when H is a single graph other than a linear forest. Indeed, recall that if $H \notin \mathcal{S}$, then already MWIS ($r = 0$) is NP-hard in H-free graphs. On the other hand, if $H \in \mathcal{S}$ but is not a linear forest, then it must contain the *claw*—the three-leaf star, and Max Induced Forest ($r = 1$) is NP-hard in the class of claw-free graphs [26].

This bring us to a question: What is the crucial property of P_t-free graphs that also extends to $C_{\geqslant t}$-free graphs, but not to H-free graphs for any fixed H that is not a linear forest, and allows us to solve $(\mathsf{tw} \leqslant r, \psi)$-MWIS efficiently?

It occurs that we should be looking at these classes from a different angle. For graph G and H, we say that H is an *induced minor* of G if it can be obtained from G by deleting vertices and contracting edges. Equivalently, this means that there is an *induced minor model* of H in G: a collection of $|V(H)|$ pairwise disjoint subsets of $V(G)$, each inducing a connected graph, and a bijection that maps these sets the vertices of H so that there is an edge between two sets if and only if their corresponding vertices are adjacent in H. We say that G is *H-induced-minor-free* if it does not contain H as an induced minor. Note that every H-induced-minor-free graph is in particular H-free. However, if H is a linear forest, then H-free graphs and H-induced-minor-free graphs coincide. Furthermore, $C_{\geqslant t}$-free graphs are precisely C_t-induced-minor-free graphs.

The following conjecture asserts that the tractability of $(\mathsf{tw} \leqslant r, \psi)$-MWIS actually extends to all classes excluding a fixed planar graph as an induced minor.

► **Conjecture A** (Gartland, Lokshtanov [15])**.** *For every planar graph H, every problem expressible as $(\mathsf{tw} \leqslant r, \psi)$-MWIS is polynomial-time-solvable in H-induced-minor-free graphs.*

Let us emphasize that the assumption that we exclude a planar graph is crucial in Conjecture A. Indeed, the class of planar graphs excludes any non-planar graph as an induced minor and MWIS is NP-hard in planar graphs [14].

We are still very far from confirming Conjecture A in full generality. Indeed, even its weakening asking for a quasipolynomial-time algorithm, or even a quasipolynomial-time approximation scheme (QPTAS) seems challenging. Still, in recent years some special cases of Conjecture A were shown [5, 6, 13]. In the context of the current paper, the most relevant one seems to be the result of Bonnet et al. [6] that MWIS can be solved in quasipolynomial time in the class of $(C_4 + sC_3)$-induced-minor-free graphs, for every fixed s. Here, by "+" we mean disjoint union and multiplication by s means s-fold disjoint union. Thus, $C_4 + sC_3$ is the graph with $s + 1$ components: one being a C_4 and the remaining ones being triangles.

Thus, in quasipolynomial time we can solve MWIS (and even $(\mathsf{tw} \leqslant r, \psi)$-MWIS) in graphs that exclude *long* induced cycles (i.e., $C_{\geqslant t}$-free) and in graphs that exclude *many* induced and pairwise non-adjacent cycles (i.e., $(C_4 + sC_3)$-induced-minor-free graphs). In this paper, we are interested in the common generalization of these classes: graphs that exclude *many* induced pairwise non-adjacent *long* cycles, i.e., the class of sC_t-induced-minor-free graphs, for fixed s and t.

While we are not able to prove Conjecture A for sC_t-induced-minor-free graphs, we provide a major step towards such a result. As the main technical contribution, we show a QPTAS for MWIS in the considered classes of graphs.

▶ **Theorem 1.** *Let s, t be positive integers and $\varepsilon \in (0, 1)$ be a real. There is an algorithm that, given a vertex-weighted graph G, in quasipolynomial time returns either:*

- *an induced minor model of sC_t in G, or*
- *an independent set of weight at least $(1-\varepsilon)$ times the maximum possible weight of an independent set in G.*

The aforementioned quasipolynomial-time algorithm for MWIS in H-free graphs for $H \in \mathcal{S}$ [17] was preceded by a QPTAS in the same graph classes [11]. A recent similar result is a QPTAS for MWIS in graphs excluding a fixed wheel as an induced minor [9].

Next, combining Theorem 1 with known results concerning approximation schemes [18], we obtain a QPTAS for the unweighted variant of $(\mathsf{tw} \leqslant r, \psi)$-MWIS under an additional mild assumption that ψ is a *hereditary* CMSO_2 formula: (i) if $G \models \psi$, then $G' \models \psi$ for every induced subgraph G' of G, and (ii) if $G_1 \models \psi$ and $G_2 \models \psi$, then $G_1 + G_2 \models \psi$. We note that many natural graph properties, like e.g., planarity, bounded degeneracy, or excluding a fixed graph as a minor, can be defined by hereditary CMSO_2 formulas.

▶ **Theorem 2.** *Let $r \geqslant 0$, let s, t be positive integers, $\varepsilon \in (0, 1)$ be a real, and ψ be a hereditary CMSO_2 formula. There is an algorithm that, given a graph G, in quasipolynomial time returns one of the following outputs:*

- *an induced minor model of sC_t in G, or*
- *a solution to $(\mathsf{tw} \leqslant r, \psi)$-MWIS of size at least $(1-\varepsilon)$ times the optimum one, or,*
- *a correct conclusion that no solution to $(\mathsf{tw} \leqslant r, \psi)$-MWIS exists.*

2 Preliminaries

For integers k, ℓ, the set $\{k, k+1, \ldots, \ell\}$ is denoted as $[k, \ell]$ and we shorten $[1, k]$ to $[k]$.

We consider here (vertex-)weighted graphs (G, w), where w is a function $V(G) \to \mathbb{R}_{\geqslant 0}$. For a subset X of vertices, we define its weight as $\mathsf{w}(X) = \sum_{v \in X} \mathsf{w}(v)$. For a weighted graph (G, w), by $\alpha(G, \mathsf{w})$ we denote the weight of a maximum-weight independent set in G. We assume that all computations on weights are performed in constant time.

The *open neighborhood* of a vertex v in G, denoted by $N_G(v)$, is the set of vertices adjacent to v. The *closed neighborhood* of v is the set $N_G[v] = N_G(v) \cup \{v\}$. For a subset X of vertices of G, its *open neighborhood* (resp., *closed neighborhood*) we mean the set $N_G(X) = \bigcup_{v \in X} N_G(v) \setminus X$ (resp., $N_G[X] = \bigcup_{v \in X} N_G[v]$. If G is clear from the context, we omit the subscript and simply write $N(\cdot)$ and $N[\cdot]$.

We say that two disjoint sets $X, Y \subseteq V(G)$ are non-adjacent if there is no edge with one endpoint in X and the other in Y. Since all subgraphs in the paper are induced, we often identify such a subgraph with its vertex set.

For a connected graph G and a set $A \subseteq V(G)$, the *BFS-layering of G from A* is the partition of $V(G)$ into sets $L_0, L_1, \ldots, L_h$ called *layers*, where $L_0 = A$ and for all $i \geqslant 1$, we have $L_i = N(L_{i-1}) \setminus \bigcup_{j<i-1} L_j$, and h is the largest possible so that $L_h \neq \emptyset$.

A *hole* in the graph is an induced cycle with at least 4 vertices. The following easy observation allows us to look for holes of certain length.

▶ **Lemma 3.** *Given a graph G and an integer $t \geqslant 4$, one can in time $\mathcal{O}(|V(G)|^t \cdot (|V(G)| + |E(G)|))$ find a shortest hole in G of length at least t, or conclude that no such hole exists.*

Proof. First, in time $\mathcal{O}(|V(G)|^t)$ we exhaustively check whether G has an induced cycle with exactly t vertices. If such a cycle exists, we return it as it is clearly shortest possible.

Next, for every tuple τ of t distinct vertices of G we check whether τ induces a path. Suppose this is the case and let x, y the endvertices of this path. We remove from G the closed neighborhood of all internal vertices of the path, except for x and y. Finally, we search for a shortest path connecting x and y in the obtained graph; together with the vertices from τ it forms an induced cycle with at least t vertices. We return the shortest of all cycles found in this process, or report that no cycle was found. ◀

Finally, let us recall a result that is particularly useful for constructing QPTASes for MWIS. A vertex v of a weighted graph (G, w) is γ-*heavy* with respect to a set $I \subseteq V(G)$ if $\mathsf{w}(N_G[v] \cap I) \geqslant \gamma \cdot \mathsf{w}(I)$.

▶ **Lemma 4** ([11, Lemma 4.1]). *Let (G, w) be a weighted graph on n vertices and $\gamma \in (0, 1)$ be a real number. In time $n^{\mathcal{O}(\log n/\gamma)}$ we can enumerate a family $\mathcal{I}'$ of $n^{\mathcal{O}(\log n/\gamma)}$ independent sets in G, each of size at most $\lceil \gamma^{-1} \log n \rceil$, such that for every independent set I there exists $I' \in \mathcal{I}'$ such that $I' \subseteq I$ and every γ-heavy vertex with respect to I belongs to $N[I']$.*

Intuitively, Lemma 4 is a quasipolynomial approximation-preserving reduction to instances without γ-heavy vertices: we can exhaustively guess $I' \in \mathcal{I}'$ and delete $N[I']$ from the graph.

3 QPTAS for MWIS: Proof of Theorem 1

In this section, we present a QPTAS for MWIS in graphs excluding sC_t as an induced minor.

▶ **Theorem 1.** *Let s, t be positive integers and $\varepsilon \in (0, 1)$ be a real. There is an algorithm that, given a vertex-weighted graph G, in quasipolynomial time returns either:*
- *an induced minor model of sC_t in G, or*
- *an independent set of weight at least $(1 - \varepsilon)$ times the maximum possible weight of an independent set in G.*

Proof. Fix s and t. Since the value of t will be fixed throughout the whole proof, by a *long hole* we mean an induced cycle with at least t vertices. We denote such a long hole shortly by $C_{\geqslant t}$.

Let (G, w) be a weighted graph on n vertices. Let $\varepsilon > 0$ be the desired precision, i.e., we aim for an $(1 - \varepsilon)$-approximation or for finding an induced subgraph isomorphic to $sC_{\geqslant t}$.

Strategy. The algorithm is a typical recursive branching algorithm. Each recursive call is invoked for an induced subgraph G' of G and an integer $s' \leqslant s$; the goal is to either exhibit an induced $s'C_{\geqslant t}$ subgraph of G' or an independent set I' of weight close to $\alpha(G', \mathsf{w})$ (the actual error analysis is made formally later; it is not just merely an $(1 - \varepsilon)$-approximation to $\alpha(G', \mathsf{w})$). The initial call is to $G' = G$ and $s' = s$.

We set

$$\beta := \frac{\varepsilon}{s + t + \log_{6/5} n} \quad \text{and} \quad \gamma := \frac{\beta^3}{1000t}.$$

The computation of one recursive call is embedded in the following claim.

▷ Claim 5. Given an induced subgraph G' of G and an integer $1 \leqslant s' \leqslant s$, in time $n^{\mathcal{O}(\log n/\gamma)}$ one can either report an induced $s'C_{\geqslant t}$ subgraph in G' or enumerate a family $\mathcal{X}$ of pairs (X, J) such that:

1. for every $(X, J) \in \mathcal{X}$, we have $X \subseteq V(G')$ and J is an independent set in G' with $N[J] \subseteq X$;
2. for every $(X, J) \in \mathcal{X}$, every connected component D of $G' - X$ satisfies one of the following conditions:
   a. there is a $C_{\geqslant t}$ in G' which is disjoint and nonadjacent to D; or
   b. D has at most $\frac{5}{6}|V(G')|$ vertices.
3. for every independent set I' in G' of weight $\alpha(G', \mathsf{w})$, there exists $(X, J) \in \mathcal{X}$ with $J \subseteq I'$ and $\mathsf{w}((X \setminus J) \cap I') \leqslant \beta\alpha(G', \mathsf{w})$;
4. the size of $\mathcal{X}$ is bounded by $n^{\mathcal{O}(\log n/\gamma)}$.

We now argue how Claim 5 yields Theorem 1. Consider a recursive call with parameters (G', s'). If $s' = 0$, then we return $\varnothing$ as an induced $s'C_{\geqslant t}$. If $V(G') = \varnothing$, we return the empty (independent) set.

Otherwise, we apply Claim 5 to G' and s'. If an induced $s'C_{\geqslant t}$ is returned, we return it and conclude. Otherwise, we iterate over the obtained family $\mathcal{X}$. For every $(X, J) \in \mathcal{X}$, let $\mathcal{D}_X$ be the set of connected components of $G' - X$. For every $D \in \mathcal{D}_X$, we recurse on $G'_D \coloneqq G'[D]$ and either $s'_D \coloneqq s'$ if $|D| \leqslant \frac{5}{6}|V(G')|$ (Property 2b) or $s'_D \coloneqq s' - 1$ otherwise. If an induced $s'_D C_{\geqslant t}$ is returned in G'_D, we augment it with a $C_{\geqslant t}$ non-adjacent to D in G' if $s'_D = s' - 1$ (such a long hole can be found using Lemma 3) and return. Otherwise, if an independent set I'_D is returned, we compute an independent set $I_{(X,J)} = \bigcup_{D \in \mathcal{D}_X} I'_D \cup J$ (Note that this is an independent set due to condition $N[J] \subseteq X$.) Finally, if no $s'C_{\geqslant t}$ was returned for any $(X, J) \in \mathcal{X}$, we return $I_{(X,J)}$ of maximum weight among all options $(X, J) \in \mathcal{X}$. This completes the description of the algorithm.

Complexity analysis. Let us move to the analysis. For a recursive call (G', s'), let

$$\mu(G', s') \coloneqq s' + \left\lfloor \log_{6/5} |V(G')| \right\rfloor.$$

Observe that every recursive subcall (G'_D, s'_D) satisfies $\mu(G'_D, s'_D) \leqslant \mu(G', s') - 1$. Hence, the depth of the recursion is bounded by $s + \lfloor \log_{6/5} n \rfloor$. Since every recursive call results in $n^{\mathcal{O}(\log n/\gamma)}$ subcalls, the overall running time is as follows:

$$n^{\mathcal{O}\left(\log^2 n/\gamma\right)} = n^{\mathcal{O}\left(\log^2 n/\beta^3\right)} = n^{\mathcal{O}\left(\log^5 n/\varepsilon^3\right)},$$

i.e., quasipolynomial in n (we hide terms depending on s and t in the $\mathcal{O}(\cdot)$-notation).

Correctness and approximation guarantee. Clearly, if any recursive call (G', s') finds an induced $s'C_{\geqslant t}$, it is propagated up in the recursion tree and results in exhibiting an induced $sC_{\geqslant t}$ in the root call. Assume then that every recursive call (G', s') returned an independent set, which we denote $I_{G',s'}$.

Consider a recursive call (G', s') and let I' be an independent set in G' of weight $\alpha(G', \mathsf{w})$. By the promise of Claim 5, there exists $(X, J) \in \mathcal{X}$ in this recursive call with $J \subseteq I'$ and $\mathsf{w}((X \setminus J) \cap I') \leqslant \beta\alpha(G', \mathsf{w}) = \beta\mathsf{w}(I')$. In particular, $\alpha(G' - X, \mathsf{w}) \geqslant (1 - \beta)\alpha(G', \mathsf{w})$. By a standard induction on the depth of the subtree of the recursion tree, we obtain that if the depth of the recursion tree of a call (G', s') is h, then $\mathsf{w}(I_{G',s'}) \geqslant (1 - \beta)^h \alpha(G', \mathsf{w})$. In particular, since the recursion depth at the root is bounded by $s + \lfloor \log_{6/5} n \rfloor$, the root call returns an independent set of weight at least

$$(1 - \beta)^{s + \lfloor \log_{6/5} n \rfloor} \alpha(G, \mathsf{w}) \geqslant \left(1 - \beta(s + \lfloor \log_{6/5} n \rfloor)\right) \alpha(G, \mathsf{w}) \geqslant (1 - \varepsilon)\alpha(G, \mathsf{w}).$$

Thus, it remains to prove Claim 5.

Proof of Claim 5. Let I' be an independent set in G' of weight $\alpha(G', \mathsf{w})$.

We start with a standard application of Lemma 4. We enumerate the family $\mathcal{I}_0$ of $n^{\mathcal{O}(\log n/\gamma)}$ independent sets in G' such that there exists $I_0 \in \mathcal{I}_0$ such that $I_0 \subseteq I'$ and every γ-heavy vertex with respect to I' belongs to $N[I_0]$. Henceforth, we will refer to this choice of I_0 as *the correct choice of* I_0.

Initiate $\mathcal{X} = \varnothing$. We iterate over the elements of $\mathcal{I}_0$. For every $I_0 \in \mathcal{I}_0$ we proceed as follows. Let $X_0 \coloneqq N_{G'}[I_0]$ and $G_0 \coloneqq G' - X_0$.

We find a shortest long hole in G_0 by Lemma 3. If Lemma 3 reports that G_0 is $C_{\geqslant t}$-free, we invoke the exact quasipolynomial-time algorithm of [18] that finds in time $n^{\mathcal{O}(\log^4 n)}$ an independent set J of maximum weight in G_0. We insert $(V(G'), I_0 \cup J)$ into $\mathcal{X}$ and conclude computation for I_0. Note that for the correct choice of I_0, we have $\mathsf{w}(I_0 \cup J) = \alpha(G', \mathsf{w})$.

Assume then that Lemma 3 returns a long hole H. If $|V(H)| \leqslant 2t + 8$, then, assuming the correct choice of I_0, we have

$$\mathsf{w}(N_{G_0}[V(H)] \cap I') \leqslant (2t+8)\gamma \cdot \mathsf{w}(I') \leqslant \beta \mathsf{w}(I').$$

Hence, we can insert $(X \coloneqq X_0 \cup N_{G_0}[V(H)], I_0)$ into $\mathcal{X}$ and conclude the computation for the current I_0, as then $G' - X$ is non-adjacent to H, so every connected component of $G' - X$ satisfies Property 2a.

Therefore, from now on, we assume that $|V(H)| > 2t + 8$.

▷ Claim 6. Every vertex $v \in V(G_0) \setminus V(H)$ has its neighbors in $V(H)$ included in a 3-vertex subpath of H, or has a neighbor in every subpath of H at least $t - 1$ vertices.

Proof. Assume v fails at realizing the latter condition. Then there is a $(t-1)$-vertex subpath P of H, disjoint from the neighborhood of v. Let $u \neq u' \in V(H)$ be such that u and u' are neighbors of v, and delimit a subpath P' of H containing P and containing only two neighbors of v (its endpoints u and u'). Observe that if v has at most one neighbor in $V(H)$, it readily satisfies the first condition of the claim. For $V(P') \cup \{v\}$ not to contradict that H is a shortest long hole of G, vertices u and u' have to be at distance at most 2 in H (along the *other* subpath of H delimited by u and u'), and so v satisfies the first condition. ◁

We first get rid of the vertices satisfying the second condition of Claim 6. Let P be a subpath of H with $t-1$ vertices. We set $X_1 \coloneqq N_{G_0}(P) \setminus V(H)$ and let G_1 be the connected component of $G_0 \setminus X_1$ that contains H.

We will insert $X_0 \cup X_1$ into any output set X in what follows, so one may think of this step as deleting the vertices of X_1. Every connected component of $G_0 \setminus X_1$ except for G_1 satisfies Property 2a, so we need only to focus on G_1.

As $|V(P)| = t - 1$, assuming the correct choice of I_0, we have

$$\mathsf{w}(X_1 \cap I') \leqslant (t-1)\gamma \mathsf{w}(I') \leqslant \frac{\beta}{5}\mathsf{w}(I').$$

Note that every vertex satisfying the second condition of Claim 6 lies in X_1.

Let us now consider BFS layering from H in G_1 with layers $V(H) = L_0,\ L_1,\ L_2, \ldots, L_h$. If $h \leqslant \lceil \frac{5}{\beta} \rceil$, we set $h' \coloneqq h + 1$, and otherwise we iterate over all options of $1 \leqslant h' \leqslant \lceil \frac{5}{\beta} \rceil$. In both cases, there exists a choice of h' with $\mathsf{w}(I' \cap L_{h'}) \leqslant \frac{\beta}{5}\mathsf{w}(I')$; we call it henceforth *the correct choice of* h'. For fixed h', we set $X_2 \coloneqq L_{h'}$ and $G_2 \coloneqq G_1[\bigcup_{i=0}^{h'-1} L_i]$. Note that G_2 is connected and every connected component of $G_1 - X_2$ distinct from G_2 lies in $\bigcup_{i > h'} L_i$ and, consequently, is non-adjacent to H and therefore satisfies point 2a. Hence, in what follows we will insert X_2 into any returned set X and we focus on G_2.

Note that $\mathcal{L} \coloneqq L_0, L_1, \ldots, L_{h'-1}$ is the BFS layering in G_2 from H. The following claims and definitions refer to G_2.

For $i < j$, a *vertical path* between $u \in L_i$ and $v \in L_j$ with respect to $\mathcal{L}$ is a path P such that $V(P)$ intersects exactly once every layer L_k such that $k \in [i, j]$, and no other layer. The *cone* of a vertex $v \in V(G_2)$ (still with respect to $\mathcal{L}$) is the set of vertices

$$\{x \in V(H) \mid \text{there is a vertical path between } x \text{ and } v\}.$$

▷ Claim 7. In G_2, let $x \in L_p$, $y \in L_q$, with $p \leqslant q$ such that there exists a path $P \subseteq \bigcup_{i \geqslant p} L_i$ with ℓ vertices, connecting x and y which inner vertices are non-adjacent to H. Let x_1 and y_1 be two vertices in the cones of x and y, respectively. Then, the distance between x_1 and y_1 along H is at most $p + q + \ell$.

Proof. Let x, y, L_p, L_q, ℓ, x_1, y_1 and P be as in the claim statement.

Let Q_x and Q_y be two vertical paths from x and y to x_1 and y_1, respectively. Let x_2 and y_2 be the penultimate vertices of Q_x and Q_y, respectively. Note that $x_2, y_2 \in L_1$ and if $p = 1$, then $x_2 = x$ and $y_2 = y$. Let R be the shortest path in $G_2[V(P) \cup V(Q_x) \cup V(Q_y) \setminus \{x_1, y_1\}]$ connecting x_2 and y_2. Note that R is non-adjacent to H except for the endpoints. Furthermore, $|E(R)| \leqslant (p-1) + (q-1) + \ell$.

Let A and B be the two subpaths of H connecting x_1 and y_1. Without loss of generality, we assume that $|E(A)| \geqslant |E(B)|$. As $|V(H)| > 2t + 8$, we have $|E(A)| \geqslant t + 5$.

Let x_3 and y_3 be the neighbors of x_2 and y_2 on A, respectively, such that no vertex of A between x_3 and y_3 is adjacent to neither x_2 nor y_2. Let A' be the subpath of A from x_3 to y_3. Since the neighborhood in H of a vertex from L_1 is contained in a subpath of at most three consecutive vertices (by Claim 6), we have $|E(A')| \geqslant |E(A)| - 4 \geqslant t$.

We now consider a cycle H' that is a concatenation of x_2x_3, A', y_3y_2, and R. Note that H' is an induced cycle in G_2. As $|E(A')| \geqslant t$, H' is a long hole. As H is the shortest hole of length at least t, we have

$$|E(A)| + |E(B)| = |E(H)| \leqslant |E(H')| = 2 + |E(A')| + |E(R)| \leqslant |E(A)| + p + q + \ell.$$

Hence, $|E(B)| \leqslant p + q + \ell$. This completes the proof. ◁

▷ Claim 8. In G_2, let $x \in L_p$, $y \in L_q$, with $p \leqslant q$ such that there exists a path $P \subseteq \bigcup_{i \geqslant p} L_i$ with ℓ vertices connecting x and y which is disjoint and non-adjacent to H. Then the union of cones of x and y is contained in a subpath of H with at most $2(p + q + \ell)$ vertices.

Proof. Let Z be the union of the cones of x and y. By Claim 7, any two vertices, one belonging to the cone of x and another to the cone of y respectively, in Z lie within distance at most $p + q + \ell$ along H. Hence, Z lies in the subpath of H with at most $2(p + q + \ell)$ vertices. ◁

In the further proof, we use the following two corollaries of Claim 8:

▷ Claim 9. For any p, the cone of any vertex $v \in L_p$ is contained in a subpath of H with at most $4p$ vertices.

▷ Claim 10. Let $v_1 \in L_p$ and $v_2 \in (L_{p-1} \cup L_p) \cap N(v_1)$. Then the union of the cones of v_1 and v_2 is contained in a subpath of H with at most $4p + 1$ vertices.

For every $u \in V(H)$, let $D(u) \subseteq V(G_2)$ be the vertices whose cone contains u. Let $u_0, u_1, \ldots, u_{|V(H)|-1}$ be a numbering of the vertices of H along the hole H. We set

$$D_{a,b} \coloneqq \bigcup_{k \in [a, a+b-1]} D(u_{k \bmod |V(H)|}) \quad \text{and } b \coloneqq 4\left\lceil \frac{5}{\beta} \right\rceil + 1.$$

As $p \leqslant h' \leqslant \left\lceil \frac{5}{\beta} \right\rceil$, so $4p < b$. The following claim is a direct consequence of Claim 9.

▷ Claim 11. Two sets $D_{a,b}$ and $D_{a',b}$ are disjoint if $|a - a'| \bmod |V(H)| > 2b$.

The removal of $D_{a,b} \cup D_{a',b}$ disconnects G, provided $D_{a,b}$ and $D_{a',b}$ are disjoint.

▷ Claim 12. Let $a, a' \in [0, |V(H)| - 1]$ be such that $|a - a'| \bmod |V(H)| > 2b$. Then the removal of $D_{a,b} \cup D_{a',b}$ disconnects $\tilde{G}$. In particular $u_{a+b \bmod |V(H)|}$ and $u_{a'+b \bmod |V(H)|}$ are in distinct connected components of $\tilde{G} - (D_{a,b} \cup D_{a',b})$.

Proof. Let $G_3 \coloneqq G_2 - (D_{a,b} \cup D_{a',b})$, and assume without loss of generality that $a < a'$. We want to argue that G_3 is disconnected. By Claim 10, no vertex of $V(G_3) \cap \bigcup_{j \in [a+b, a'-1]} D(u_j)$ can be adjacent to a vertex of $V(G_3) \cap \bigcup_{j \in [0,a-1] \cup [a'+b, |V(H)|-1]} D(u_j)$. Thus in particular u_{a+b} and $u_{a'+b \bmod |V(H)|}$ are in distinct connected components of G'. ◁

Let $n_2 \coloneqq |V(G_2)|$. Fix $0 \leqslant a_1 \leqslant a_2 \leqslant |V(H)| - 1$ minimizing $a_2 - a_1$ subject to

$$\left| \bigcup_{j \in [a_1, a_2]} D(u_j) \right| \geqslant \frac{n_2}{6}. \tag{1}$$

Observe that $a_1 = a_2$ or

$$\left| \bigcup_{j \in [a_1, a_2]} D(u_j) \right| < \frac{n_2}{3}. \tag{2}$$

We set $d \coloneqq 2b \cdot \left\lceil \frac{5}{\beta} \right\rceil$, and distinguish two cases: $a_2 - a_1 \leqslant d$ or $a_2 - a_1 > d$.

Case: $a_2 - a_1 \leqslant d$. Let us consider the family of sets

$$\mathcal{F}_3 \coloneqq \{D_{a_1-d,b}, D_{a_1-d+2b,b}, D_{a_1-d+4b,b}, \ldots, D_{a_1-4b,b}, D_{a_1-2b,b}\}.$$

By Claim 11, the $\lceil \frac{5}{\beta} \rceil$ sets of $\mathcal{F}_3$ are pairwise disjoint, thus there exists $X_3 \in \mathcal{F}_3$ such that $\mathsf{w}(X_3 \cap I') \leqslant \frac{\beta}{5}\mathsf{w}(I')$. For the same reason, there exists

$$X_4 \in \mathcal{F}_4 \coloneqq \{D_{a_2+1,b}, D_{a_2+1+2b,b}, D_{a_2+1+4b,b}, \ldots, D_{a_2+d+1-4b,b}, D_{a_2+d+1-2b,b}\}$$

such that $\mathsf{w}(X_4 \cap I') \leqslant \frac{\beta}{5}\mathsf{w}(I')$. We iterate over all choices of X_3 and X_4. Finally, we exploit the fact that $a_2 - a_1$ is small by adding $X_5 \coloneqq \bigcup_{j \in [a_1, a_2]} N(u_j)$ to the set of vertices to remove. As $d \cdot \gamma \leqslant \frac{\beta}{5}$, for the correct choices of I_0 and h' we have $\mathsf{w}(X_5 \cap I') \leqslant \frac{\beta}{5}\mathsf{w}(I')$. We define $X \coloneqq X_0 \cup X_1 \cup X_2 \cup X_3 \cup X_4 \cup X_5$ and insert (X, I_0) into $\mathcal{X}$. We have $\mathsf{w}(X \cap (I' \setminus I_0)) \leqslant \beta \mathsf{w}(I')$.

We already argued that the connected components of $G' - X$ disjoint from $V(G_2)$ satisfy Property 2a. By Claim 12, $G_2 - (X_3 \cup X_4 \cup X_5)$ has two sets of connected components: those intersecting $\bigcup_{j \in [a_1, a_2]} D(u_j)$, and those not. The former kind are not adjacent to H, so satisfy Property 2a. The latter kind has at most $5n_2/6 \leqslant 5|V(G')|/6$ vertices, by design of the interval $[a_1, a_2]$ using (1), so satisfy Property 2b.

Case: $a_2 - a_1 > d$. Recall that by (2) we have $|\bigcup_{j\in[a_1,a_2]} D(u_j)| < n_2/3$. Then $\bigcup_{j\in[a_1-d,a_1-1]} D(u_j)$ and $\bigcup_{j\in[a_2+1,a_2+d]} D(u_j)$ contains each less than $n_2/6$ vertices, by the minimality of a_2-a_1. Hence $|\bigcup_{j\in[a_1-d,a_2+d]} D(u_j)| < 2n_2/3$. By the pigeonhole principle, as $d = 2b\cdot\lceil\frac{5}{\beta}\rceil$, there is $a \in \{a_1-d, a_1-d+2b, \ldots, a_1-2b\}$ (resp., $a' \in \{a_2+2b, a_2+4b, \ldots, a_2+d\}$) such that $\mathsf{w}(D_{a,b} \cap I') \leqslant \frac{\beta}{5}\mathsf{w}(I')$ (resp., $\mathsf{w}(D_{a',b} \cap I') \leqslant \frac{\beta}{5}\mathsf{w}(I')$). We iterate over all choices of a and a' and set $X_3 \coloneqq D_{a,b}$, $X_4 \coloneqq D_{a',b}$. We take $X \coloneqq X_0 \cup X_1 \cup X_2 \cup X_3 \cup X_4$ and insert (X, I_0) into $\mathcal{X}$. For the correct choice of I_0, h', a, a' we have $\mathsf{w}(X \cap (I' \setminus I_0)) \leqslant \beta\mathsf{w}(I')$.

Observe finally that $G_2 - (X_3 \cup X_4)$ has no connected component of size larger than $2n_2/3$, hence larger than $5|V(G')|/6$. This completes the proof of Claim 5 and thus, of Theorem 1. ◀

4 QPTAS for $(\mathsf{tw} \leqslant r, \psi)$-MWIS: Proof of Theorem 2

The proof of Theorem 2 is based on the approach introduced by Gartland et al. [18, Section 4]. For a graph G, we define its *blob graph* G° as follows:

$$V(G^\circ) = \{X \subseteq V(G) \mid G[X] \text{ is connected}\}$$
$$E(G^\circ) = \{XY \mid G[X \cup Y] \text{ is connected}\}.$$

Equivalently, edges join sets that are either non-disjoint, or there is an edge from one set to another. Gartland et al. [18] proved that for every $t \geqslant 4$, a graph G is $C_{\geqslant t}$-free if and only if G° is $C_{\geqslant t}$-free. First, let us extend this result to the setting of H-induced-minor-free graphs, where every component of H is a hole. The proof closely follows the proof of Paesani et al. for the case if H is a linear forest [28].

▶ **Theorem 13.** *Let H be a graph whose every component is a hole. The graph G contains H as an induced minor if and only if G° contains H as an induced minor.*

Proof. As G is an induced subgraph of G°, the forward implication is immediate. We prove the backward implication by induction on the number s of connected components of H. The case $s = 1$ was shown by Gartland et al. [18]. Thus, assume that $s > 1$ and let C be a connected component of H. Let $H' = H - C$.

Suppose that G° contains H as an induced minor. Fix one induced minor model of H in G and let $\mathcal{X}$ be the set of vertices of the model. This means that $G^\circ[\mathcal{X}]$ has s components, each of which is a cycle, and there is a one-to-one mapping from the components of $G^\circ[\mathcal{X}]$ and components of H, so that each cycle is mapped to a cycle of at most its own length.

Let $\mathcal{Y} \subseteq \mathcal{X}$ be the vertices of G° that form the cycle mapped to C; each element of $\mathcal{Y}$ is a subset of $V(G)$. Let $Y \subseteq V(G)$ be the union of all sets in $\mathcal{Y}$. Note that $G^\circ[\mathcal{Y}]$ is an induced subgraph of $(G[Y])^\circ$. Thus, by the inductive assumption, $G[Y]$ contains an induced cycle with at least $|V(C)|$ vertices.

Now let $X \subseteq V(G)$ be the union of sets in $\mathcal{X} \setminus \mathcal{Y}$. Since $\mathcal{Y}$ is the vertex set of one component of $G^\circ[\mathcal{X}]$, there are no edges between $\mathcal{X} \setminus \mathcal{Y}$ and $\mathcal{Y}$. Consequently, there are no edges in G between X and Y, and thus $G^\circ[X \setminus Y]$ is an induced subgraph of $(G - N[Y])^\circ$.

Since $G^\circ[\mathcal{X} \setminus \mathcal{Y}]$, and thus $(G - N[Y])^\circ$, contains H' as an induced minor, by the inductive assumption we know that $G - N[Y]$ contains H' as an induced minor. Combining its model with the model of C in $G[Y]$, we obtain an induced minor model of H in G. ◀

Let $\mathcal{H}$ be a class of graphs. An *induced $\mathcal{H}$-packing* in a graph G is a set $S \subseteq V(G)$ such that every component of $G[S]$ belongs to $\mathcal{H}$. For example, if $\mathcal{H} = \{K_1\}$, then induced $\mathcal{H}$-packing in G if and only if it is an independent set.

The construction of blob graphs allows us to reduce the problem of finding an induced $\mathcal{H}$-packing of maximum size (or weight) in G to solving the MWIS problem in an appropriate induced subgraph of G°. This is expressed in the following lemma whose proof is immediate, see e.g., [18, Section 4].

▶ **Lemma 14.** *Let G be a graph, $\mathsf{w} : V(G) \to \mathbb{R}_{\geqslant 0}$ be a weight function, and $\mathcal{H}$ be any class of graphs. Let $\mathcal{X} = \{X \subseteq V(G) \mid G[X] \in \mathcal{H}$ and is connected$\}$. Let $\mathsf{w}^\circ : \mathcal{X} \to \mathbb{R}_{\geqslant 0}$ be defined as $\mathsf{w}^\circ(X) = \sum_{v \in X} \mathsf{w}(v)$. Then the maximum weight of an induced $\mathcal{H}$-packing in G is equal to $\alpha(G^\circ[\mathcal{X}], \mathsf{w}^\circ)$.*

The problem with using Lemma 14 is that $\mathcal{X}$ might be very large, so the running time of the algorithm solving MWIS on $(G^\circ[\mathcal{X}], \mathsf{w}^\circ)$ is not bounded by a moderate function of $|V(G)|$.

A class $\mathcal{H}$ of graphs is *weakly hyperfinite* if for every $\delta > 0$ there exists $c_\delta \in \mathbb{N}$ such that for any $G \in \mathcal{H}$ there exists a set $S \subseteq |V(G)|$ of size at most $\delta \cdot |V(G)|$ so that every component of $G - S$ has at most c_δ vertices [27, Section 16.2]. The following theorem follows the result of Gartland et al. [18]. Again, we include the proof for completeness.

▶ **Theorem 15.** *Let $\mathcal{H}$ be a non-empty hereditary weakly hyperfinite class of graphs. Let s, t be positive integers and $\varepsilon \in (0, 1)$ be a real. There is an algorithm that, given a graph G, in quasipolynomial time returns one of the following outputs:*

- *an induced minor model of sC_t in G, or*
- *an induced $\mathcal{H}$-packing of size at least $(1 - \varepsilon)$ times the size of a largest induced $\mathcal{H}$-packing in G.*

Proof. Without loss of generality assume that $t \geqslant 4$. Let $\delta = \varepsilon/2$ and let c_δ be the constant witnessing that $\mathcal{H}$ is weakly hyperfinite. Let G be an n-vertex instance instance of the problem. Let

$$\mathcal{X} = \{X \subseteq V(G) \mid |X| \leqslant c_\delta \text{ and } G[X] \text{ is a connected graph from } \mathcal{H}\}.$$

Consider the graph $G^\circ[\mathcal{X}]$; note that it has $\mathcal{O}(n^{c_\delta})$ vertices and can be constructed in time polynomial in n as c_δ is a constant. We call the algorithm from Theorem 1 for $G^\circ[\mathcal{X}]$, weigths w° defined as in Lemma 14, and precision $\varepsilon/2$. It running time is quasipolynomial in $\mathcal{O}(n^{c_\delta})$ and thus in n.

If an induced minor of sC_t is reported, we report an induced minor of sC_t in G. Indeed, recall that $G^\circ[\mathcal{X}]$ is an induced subgraph of G°. Consequently, by Theorem 13, if $G^\circ[\mathcal{X}]$ contains sC_t as an induced minor, so does G.

Thus, suppose that the algorithm from Theorem 1 returns an induced $\mathcal{H}$-packing $\mathcal{S}$ in $G^\circ[\mathcal{X}]$. It is straightforward to verify that $S = \bigcup \mathcal{S}$ is an induced $\mathcal{H}$-packing in G. Let us argue how its size compares to the optimum one.

Let S^* be an optimum induced $\mathcal{H}$-packing in G. Let us construct another induced $\mathcal{H}$-packing S' as follows. We consider every component of $G[S^*]$ separately, let C be the vertex set of one such component. If $|C| \leqslant c_\delta$, we set $C' = C$ and include this set in S'. Otherwise, since $G[C]$ is weakly hyperfinite, there is a set $Q_C \subseteq C$ of size at most $\delta|C| = \varepsilon/2 \cdot |C|$, such that every component of $G[C \setminus Q_C]$ has at most c_δ vertices. We set $C' = C \setminus Q_C$ and include it into S'. Since $\mathcal{H}$ is hereditary, it holds that $G[C'] \in \mathcal{H}$. Note that in both cases we have $|C'| \geqslant (1 - \delta)|C| = (1 - \varepsilon/2)|C|$. Consequently, we obtain

$$|S^*| = \sum_{C\text{: component of } G[S^*]} |C| \leqslant \sum_{C\text{: component of } G[S^*]} \frac{|C'|}{(1 - \varepsilon/2)} = \frac{|S'|}{(1 - \varepsilon/2)}.$$

Now, by Lemma 14 notice that the size of S' is equal to the weight of a largest-weight independent set in $G^\circ[\mathcal{X}]$. Thus, by Theorem 1, we have $|S| \geqslant (1-\varepsilon/2)|S'|$. Summarizing, it holds that

$$|S| \geqslant (1-\varepsilon/2)|S'| \geqslant (1-\varepsilon/2)(1-\varepsilon/2)|S^*| = (1-\varepsilon+\varepsilon^2)|S^*| \geqslant (1-\varepsilon)|S^*|.$$

This completes the proof. ◀

Now let us argue how Theorem 15 implies Theorem 2.

▶ **Theorem 2.** *Let $r \geqslant 0$, let s, t be positive integers, $\varepsilon \in (0,1)$ be a real, and ψ be a hereditary CMSO$_2$ formula. There is an algorithm that, given a graph G, in quasipolynomial time returns one of the following outputs:*

- *an induced minor model of sC_t in G, or*
- *a solution to $(\mathsf{tw} \leqslant r, \psi)$-MWIS of size at least $(1-\varepsilon)$ times the optimum one, or,*
- *a correct conclusion that no solution to $(\mathsf{tw} \leqslant r, \psi)$-MWIS exists.*

Proof. Note that if ψ is *consistent*, i.e., if there is any graph H such that $H \models \psi$, then, since ψ is hereditary, we have $K_1 \subseteq \psi$. Consequently, in such a case, every non-empty instance of $(\mathsf{tw} \leqslant r, \psi)$-MWIS has a solution (for example, a single vertex). Since ψ is a fixed formula, we can verify if it is consistent in constant time. If not, we return the third outcome.

So from now on let us assume that some solution exists. It is known that graphs of treewidth at most r, for any fixed r, form a weakly hyperfinite class of graphs [27, Theorem 16.5]. Furthermore, this class is closed under vertex deletion and disjoint unions. Consequently, graphs of treewidth at most r satisfying ψ form a non-empty class closed under vertex deletion and disjoint unions. Thus, we obtain Theorem 2 as an immediate corollary of Theorem 15. ◀

5 Conclusion

Let us conclude the paper with pointing out two specific problems for further research. First, it would be interesting to strengthen our Theorem 1 by solving the problem *exactly* without significantly increasing the running time.

▶ **Problem P1.** *Show that for every fixed s, t, MWIS can be solved in quasipolynomial time in sC_t-induced-minor-free graphs.*

Second, let us suggest the following possible extension of our Theorem 1. Our starting point was a quasipolynomial-time algorithm for H-induced-minor-free graphs and we aimed to extend this result for H'-induced-minor-free graphs, where every component of H is isomorphic to H. Perhaps this can be done in a more general setting?

▶ **Problem P2.** *Suppose that H_1 and H_2 are graphs such that MWIS is (quasi)polynomial-time solvable in H_i-induced-minor-free graphs for $i \in \{1,2\}$. Show a quasipolynomial-time algorithm (or at least a QPTAS) for MWIS in $(H_1 + H_2)$-induced-minor-free graphs.*

We remark that the solution to Problem P2 is known if instead of forbidding certain graphs as induced minors, we forbid them as induced subgraphs [16]. However, the approach relies on the fact that induced subgraphs are constant-size objects, while an induced minor model of a constant-size graph might be arbitrarily large.

References

1 Tara Abrishami, Maria Chudnovsky, Marcin Pilipczuk, Paweł Rzążewski, and Paul D. Seymour. Induced subgraphs of bounded treewidth and the container method. *SIAM J. Comput.*, 53(3):624–647, 2024.

2 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.

3 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.

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 Marthe Bonamy, Edouard Bonnet, Hugues Déprés, Louis Esperet, Colin Geniet, Claire Hilaire, Stéphan Thomassé, and Alexandra Wesolek. Sparse graphs with bounded induced cycle packing number have logarithmic treewidth. In Nikhil Bansal and Viswanath Nagarajan, editors, *Proceedings of the 2023 ACM-SIAM Symposium on Discrete Algorithms, SODA 2023, Florence, Italy, January 22-25, 2023*, pages 3006–3028. SIAM, 2023.

6 Édouard Bonnet, Julien Duron, Colin Geniet, Stéphan Thomassé, and Alexandra Wesolek. Maximum independent set when excluding an induced minor: $K_1 + tK_2$ and $tC_3 \uplus C_4$. *Algorithmica*, 88(1):16, 2026.

7 Andreas Brandstädt and Raffaele Mosca. Maximum weight independent set for ℓclaw-free graphs in polynomial time. *Discret. Appl. Math.*, 237:57–64, 2018.

8 Parinya Chalermsook, Marek Cygan, Guy Kortsarz, Bundit Laekhanukit, Pasin Manurangsi, Danupon Nanongkai, and Luca Trevisan. From gap-exponential time hypothesis to fixed parameter tractable inapproximability: Clique, dominating set, and more. *SIAM J. Comput.*, 49(4):772–810, 2020.

9 Maria Chudnovsky, Jochen Pascal Gollin, Matjaz Krnc, and Martin Milanič. Dominated balanced separators in wheel-induced-minor-free graphs. *CoRR*, abs/2512.12329, 2025.

10 Maria Chudnovsky, Rose McCarty, Marcin Pilipczuk, Michał Pilipczuk, and Paweł Rzążewski. Sparse induced subgraphs in P_6-free graphs. In David P. Woodruff, editor, *Proceedings of the 2024 ACM-SIAM Symposium on Discrete Algorithms, SODA 2024, Alexandria, VA, USA, January 7-10, 2024*, pages 5291–5299. SIAM, 2024.

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. *SIAM J. Comput.*, 53(1):47–86, 2024.

12 Marek Cygan, Fedor V. Fomin, Lukasz Kowalik, Daniel Lokshtanov, Dániel Marx, Marcin Pilipczuk, Michal Pilipczuk, and Saket Saurabh. *Parameterized Algorithms*. Springer, 2015.

13 Clément Dallard, Martin Milanič, and Kenny Štorgel. Treewidth versus clique number. III. Tree-independence number of graphs with a forbidden structure. *Journal of Combinatorial Theory, Series B*, 167:338–391, 2024.

14 M.R. Garey, D.S. Johnson, and L. Stockmeyer. Some simplified NP-complete graph problems. *Theoretical Computer Science*, 1(3):237–267, 1976.

15 Peter Gartland. *Quasi-Polynomial Time Techniques for Independent Set and Beyond in Hereditary Graph Classes*. PhD thesis, University of California Santa Barbara, 2023.

16 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.

17 Peter Gartland, Daniel Lokshtanov, Tomás Masařík, Marcin Pilipczuk, Michał Pilipczuk, and Paweł Rzążewski. 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.

18 Peter Gartland, Daniel Lokshtanov, Marcin Pilipczuk, Michał Pilipczuk, and Paweł Rzążewski. 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.

19 Andrzej Grzesik, Tereza Klimošová, Marcin Pilipczuk, and Michał Pilipczuk. Polynomial-time algorithm for maximum weight independent set on P_6-free graphs. *ACM Trans. Algorithms*, 18(1):4:1–4:57, 2022.

20 Johan Håstad. Clique is hard to approximate within $n^{(1-\epsilon)}$. In *Acta Mathematica*, pages 627–636, 1996.

21 Richard M. Karp. Reducibility among combinatorial problems. In Raymond E. Miller, James W. Thatcher, and Jean D. Bohlinger, editors, *Complexity of Computer Computations: Proceedings of a symposium on the Complexity of Computer Computations, held March 20–22, 1972, at the IBM Thomas J. Watson Research Center, Yorktown Heights, New York, and sponsored by the Office of Naval Research, Mathematics Program, IBM World Trade Corporation, and the IBM Research Mathematical Sciences Department*, pages 85–103, Boston, MA, 1972. Springer US.

22 Bingkai Lin, Xuandi Ren, Yican Sun, and Xiuhan Wang. Improved hardness of approximating k-Clique under ETH. In *64th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2023, Santa Cruz, CA, USA, November 6-9, 2023*, pages 285–306. IEEE, 2023.

23 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.

24 Vadim V. Lozin and Martin Milanič. A polynomial algorithm to find an independent set of maximum weight in a fork-free graph. *J. Discrete Algorithms*, 6(4):595–604, 2008.

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

26 Andrea Munaro. On line graphs of subcubic triangle-free graphs. *Discret. Math.*, 340(6):1210–1226, 2017.

27 Jaroslav Nešetřil and Patrice Ossona de Mendez. *Sparsity - Graphs, Structures, and Algorithms*, volume 28 of *Algorithms and combinatorics*. Springer, 2012.

28 Giacomo Paesani, Daniël Paulusma, and Paweł Rzążewski. Feedback Vertex Set and Even Cycle Transversal for H-free graphs: Finding large block graphs. *SIAM J. Discret. Math.*, 36(4):2453–2472, 2022.

29 Marcin Pilipczuk, Michał Pilipczuk, and Paweł Rzążewski. 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.

30 Svatopluk Poljak. A note on stable sets and colorings of graphs. *Commentationes Mathematicae Universitatis Carolinae*, 15:307–309, 1974.

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