

Maximum Independent Set when excluding an induced minor: $K_1 + tK_2$ and $tC_3 \uplus C_4$

Édouard Bonnet*  

*Corresponding author: edouard.bonnet@ens-lyon.fr

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

Julien Duron 

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

Colin Geniet  

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

Stéphan Thomassé 

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

Alexandra Wesolek  

Simon Fraser University, Burnaby, BC, Canada

Abstract

Dallard, Milanič, and Štorgel [JCTB '24] ask if, for every class excluding a fixed planar graph H as an induced minor, MAXIMUM INDEPENDENT SET can be solved in polynomial time, and show that this is indeed the case when H is any planar complete bipartite graph, or the 5-vertex clique minus one edge, or minus two disjoint edges. A positive answer would constitute a far-reaching generalization of the state-of-the-art, when we currently do not know if a polynomial-time algorithm exists when H is the 7-vertex path. Relaxing tractability to the existence of a quasipolynomial-time algorithm, we know substantially more. Indeed, quasipolynomial-time algorithms were recently obtained for the t -vertex cycle, C_t [Gartland et al., STOC '21], and the disjoint union of t triangles, tC_3 [Bonamy et al., SODA '23].

We give, for every integer t , a polynomial-time algorithm running in $n^{O(t^5)}$ when H is the friendship graph $K_1 + tK_2$ (t disjoint edges plus a vertex fully adjacent to them), and a quasipolynomial-time algorithm running in $n^{O(t^2 \log n) + f(t)}$, with f a single-exponential function, when H is $tC_3 \uplus C_4$ (the disjoint union of t triangles and a 4-vertex cycle). The former generalizes the algorithm readily obtained from Alekseev's structural result on graphs excluding tK_2 as an induced subgraph [Alekseev, DAM '07], while the latter extends Bonamy et al.'s result.

2012 ACM Subject Classification Theory of computation → Graph algorithms analysis

Keywords and phrases Maximum Independent Set, forbidden induced minors, quasipolynomial-time algorithms

Funding This work was supported by the ANR projects TWIN-WIDTH (ANR-21-CE48-0014) and Digraphs (ANR-19-CE48-0013).

Alexandra Wesolek: Supported by the Vanier Canada Scholarship Program.

Acknowledgements We would like to thank Dibyayan Chakraborty for useful discussions, and anonymous reviewers for their helpful comments, and in particular a nice simplification in the proof of Theorem 2.

1 Introduction

The MAX INDEPENDENT SET (MIS for short) problem asks for a largest *independent set* of its input graph G , i.e., a subset of pairwise non-adjacent vertices in G . In its decision form, the input is a graph G and an integer k , and the question is whether G admits an independent set of size at least k .

Besides the ubiquitous usefulness that such a fundamental problem has within combinatorial optimization, and notably in the areas of packing, scheduling, and coloring, MIS (or equivalently MAXIMUM CLIQUE, the same problem in the complement graph) has a very wide range of applications, as evidenced, for instance, in map labeling [50], coding theory [19], spatial scheduling [27], genetic analysis [1], information retrieval [12], macromolecular docking [29], and sociometry [28] (also see Butenko’s thesis [18]). It is thus unfortunate that this problem is not only hard to solve but also very resistant to approximation. Indeed, the decision version of MIS is NP-complete [30] and W[1]-complete [26], while its optimization version cannot be approximated within a factor $n^{1-\varepsilon}$ on n -vertex graphs, for any $\varepsilon > 0$, unless P=NP [37, 52].

In spite of this, theorists and practitioners have put a lot of effort into designing efficient algorithms for MIS. In parallel with generic approaches via integer programming, high-performance exact and heuristic MIS solvers have emerged in recent years, based on diverse methods such as kernelization and evolutionary approaches [42], deep reinforcement learning [6], graph neural networks [46], and dataless training (where backpropagation is applied to a loss function based instead on the input) [11]. On the theory side, exact exponential algorithms have been developed for decades culminating in a running time below $1.2^n n^{O(1)}$ [51].

Another approach is to try and exploit the structure that the input graphs may have. Indeed, in all the aforementioned applications, inputs are not uniformly sampled over all n -vertex graphs: They instead bear some structural properties, and in some cases, might avoid some specific patterns. Graph theory¹ offers two main notions of *patterns* or containment: the natural and straightforward *subgraphs* (obtained by removing vertices and edges), and the deeper *minors* (further allowing to contract edges). Both notions come with an induced variant, when edge removals are disallowed, bringing the number of containment types to four. It is then sensible to determine the patterns H whose absence makes MIS (more) tractable. It turns out that this question is completely settled for subgraphs and minors.

For the subgraph containment, the argument is the following. By the grid minor theorem [47], the class of graphs excluding H as a subgraph has bounded treewidth if (and only if) all the connected components of H are paths and subdivided *claws* (i.e., stars with three leaves); thus MIS can be solved in polynomial time in this class, for instance by Courcelle’s theorem [22]. If instead H has a connected component that is neither a path nor a subdivided claw, MIS remains NP-complete since such a class either contains all subcubic graphs or the $2|V(H)|$ -subdivision of every graph, two families of graphs on which MIS is known to be NP-complete [7, 10, 45].

For minors, the dichotomy hinges on whether H is planar. Indeed, if H is planar, then the class of graphs excluding H as a minor has bounded treewidth (again, mainly by the grid minor theorem), and MIS can be solved efficiently. If H is non-planar, then the H -minor-free graphs include all planar graphs for which MIS is known to be NP-complete [31].

The question is more intriguing for the induced containments, and the *induced subgraph* case has received a lot of attention. While it has long been known that if H is *not* the disjoint union of paths and subdivided claws, MIS remains NP-complete on graphs without H as an induced subgraph [7, 45], it has been conjectured that MIS is otherwise polynomial-time solvable. This has been proven when H is the 6-vertex path [36], a claw with exactly one edge subdivided [8, 43], or any disjoint union of claws [17]. The latter result extends a polynomial-time algorithm (essentially) due to Alekseev when H is any disjoint union of

¹ We refer the reader to Section 2 for the relevant background in graph theory.

edges [9]. The author indeed proves that the total number of maximal independent sets is polynomially bounded. One can then enumerate all the maximal independent sets in polynomial time (following Alekseev's proof, or using the generic output-sensitive algorithm of Tsukiyama et al. [49]), and thus find a maximum independent set.

While we currently do not know of a polynomial-time algorithm when H is the 7-vertex path, Gartland and Lokshtanov [32] have obtained a quasipolynomial-time algorithm when H is P_t , the t -vertex path, for any positive integer t ; also see [44]. Supporting the existence of a polynomial-time algorithm when H is $S_{i,j,k}$, the claw whose three edges are subdivided $i-1$, $j-1$, and $k-1$ times, respectively, a quasipolynomial-time algorithm [33], and polynomial-time algorithms among bounded-degree graphs [2], and later, among graphs excluding a fixed complete bipartite graph as a subgraph [3] have been found. The parameterized complexity of MIS when excluding a fixed induced subgraph has been studied [14, 15, 23], but the mere statement of which H make the problem fixed-parameter tractable (and which ones keep it W[1]-complete) is unclear [15].

We eventually arrive at the *induced minor* containment, the topic of the current paper. As for minors, the class of all graphs excluding a non-planar graph H as an induced minor contains all planar graphs; hence MIS remains NP-complete in such a class. However we do not know of a *planar* graph H for which MIS remains NP-complete on H -induced-minor-free graphs. This has led Dallard, Milanič, and Štorgel [25] to ask if such classes exist:

► **Question 1.** *Is it true that for every planar graph H , MAX INDEPENDENT SET can be solved in polynomial time in the class of graphs excluding H as an induced minor?*

A first observation is that avoiding H as an induced minor implies avoiding it as an induced subgraph. Thus Question 1 is settled for P_6 , $S_{1,1,2}$, and $tS_{1,1,1}$ (where tG denotes the disjoint union of t copies of G). The same authors [25] further obtain a polynomial-time algorithm when H is K_5^- (the 5-vertex clique minus an edge), $K_{2,t}$ (the complete bipartite graph with 2 vertices fully adjacent to t vertices), and $W_4 = K_1 + C_4$ (a 4-vertex cycle C_4 with a fifth vertex fully adjacent to the cycle). All three cases were shown by bounding the so-called *tree-independence number* (i.e., treewidth where *bag size* is replaced by *independence number of the subgraph induced by the bag*) [25], in which case a polynomial-time algorithm can be derived for MIS using the corresponding tree-decompositions [24]. They also show that this is as far as this sole technique can go: H -induced-minor-free graphs have bounded *tree-independence number* if and only if H is edgeless or an induced minor of K_5^- , $K_{2,t}$, or W_4 [25]. The framework of *potential maximal cliques* [16] and the *container method* have led to a polynomial-time algorithm when $H = C_5$ [4, 21].

Question 1 is a beautiful question and, if true, a very difficult one. Indeed, $H = P_7$, the 7-vertex path, is a very simple planar graph for which we currently do not know such a polynomial-time algorithm. A natural relaxation of Question 1 is to only request a quasipolynomial-time algorithm:

► **Question 2.** *Is it true that for every planar graph H , MAX INDEPENDENT SET can be solved in quasipolynomial time in the class of graphs excluding H as an induced minor?*

We know somewhat more about Question 2. There is a quasipolynomial-time algorithm for MIS in C_t -induced-minor-free graphs [34], building upon the $H = P_t$ case. Recently, Bonamy et al. [13] presented a quasipolynomial-time algorithm when H is tC_3 , i.e., the disjoint union of t triangles. (When $K_{1,t}$, the star with t leaves, is further excluded as an induced subgraph, a polynomial-time algorithm has been obtained [5].) See Table 1 for a summary of the introduction.

H excluded as	subgraph	minor	induced subgraph	induced minor
in P	$tS_{t,t,t}$	planar	$P_6, S_{1,1,2}, tS_{1,1,1}$	$K_5^-, K_{2,t}, W_4, C_5, \boxed{K_1 + tK_2}$
known in QP	—	—	$tS_{t,t,t}$	$C_t (t \geq 6), \boxed{tC_3 \uplus C_4}$
NP-c	$\neg tS_{t,t,t}$	non-planar	$\neg tS_{t,t,t}$	non-planar
open P / NP-c	—	—	P_7, \dots	P_7, C_6, \dots
open QP / NP-c	—	—	—	$C_4 \uplus C_4, \dots$

■ **Table 1** The complexity of MAX INDEPENDENT SET when H is excluded as one of the four main types of patterns. “ $\neg tS_{t,t,t}$ ” means that H is not a subgraph of $tS_{t,t,t}$ for any t . Our results are framed.

Expecting an affirmative solution to Question 1 or Question 2 may seem optimistic. However, regarding precise running time, we do know that for every planar H , MIS is probably not as difficult in H -induced-minor-free graphs as it is in general graphs. Indeed, Korhonen [41] describes a $2^{O(n/\log^c n)}$ -time algorithm, for some constant $c > 0$, to solve MIS on n -vertex graphs excluding a fixed planar graph H as an induced minor. Assuming the Exponential-Time Hypothesis [38], such a running time is impossible in general graphs [39].

Our results. We make some progress regarding Questions 1 and 2. Our first contribution is, for every positive integer t , a polynomial-time algorithm when H is the *friendship graph* $K_1 + tK_2$ (also called *Dutch windmill graph* or *fan*), i.e., t independent edges universally linked to a $(2t + 1)$ -st vertex:

► **Theorem 1.** *For every positive integer t , MAX INDEPENDENT SET can be solved in polynomial-time $n^{O(t^5)}$ in n -vertex $K_1 + tK_2$ -induced-minor-free graphs.*

This extends Alekseev’s result [9] for graphs excluding tK_2 as an induced subgraph, or equivalently, as an induced minor. We indeed use this result to first derive a polynomial-time algorithm in subgraphs of $K_1 + tK_2$ -induced-minor-free graphs G induced by vertices from a bounded number of breadth-first search (BFS) layers of G .

We then consider the connected components of our input graph G when deprived of a subset X of vertices inducing tK_2 and, subject to that property, maximizing the order of the largest connected component in $G - X$. We show that, due to this careful selection of X , every component C of $G - X$ admits an efficiently constructible path-decomposition \mathcal{P} with *bounded adhesion* (i.e., any two distinct bags have a bounded intersection), each bag of which is contained in a bounded number of consecutive BFS layers of C . Hence MIS can be solved efficiently within a bag, by our opening step (see previous paragraph). This part is quite technical, but mostly to justify the existence of \mathcal{P} . The algorithm itself remains simple.

Theorem 1 is then obtained by exhaustively finding X and guessing its intersection X' with a maximum independent set of G , and performing dynamic programming on the connected components of $G - X$, deprived of $N(X')$. The dynamic-programming table is filled via the efficient algorithm when handling an induced subgraph contained in few BFS layers.

Our second contribution is a quasipolynomial-time algorithm when H is $tC_3 \uplus C_4$, the disjoint union of t triangles and a 4-vertex cycle:

► **Theorem 2.** *For every positive integer t , MAX INDEPENDENT SET can be solved in quasipolynomial-time $n^{O(t^2 \log n) + f(t)}$ (where f is single-exponential) in $tC_3 \uplus C_4$ -induced-minor-free graphs.*

We first perform a quasipolynomial branching rule to get rid of holes of length at most 6 (i.e., induced cycles of length 4, 5, or 6). We then assume that the input graph G is not $(t+2)C_3$ -induced-minor-free, for otherwise we conclude with Bonamy et al.'s algorithm [13]. Thus G , being $tC_3 \uplus C_4$ -induced-minor-free, has to admit $(t+2)C_3$ as an induced subgraph, i.e., a collection T_1, \dots, T_{t+2} of $t+2$ pairwise vertex-disjoint and non-adjacent triangles. We define $S_{i,j}$, minimally separating T_i and T_j in the graph G deprived of the neighborhoods of the other triangles T_k (with $k \neq i, j$).

We show that each $S_{i,j}$ induces a clique. So does every intersection $N_{i,j}$ of the neighborhood of two distinct triangles T_i, T_j of the collection (this is where getting rid of the holes of length at most 6 comes into play). We can therefore exhaustively guess the intersection of a maximum independent set with the union of the sets $S_{i,j}$ and $N_{i,j}$ (for every $i < j \in [t+2]$). We finally observe that $G' = G - \bigcup_{i \neq j \in [t+2]} (S_{i,j} \cup N_{i,j})$ is chordal, since the presence of a hole H in G' would imply the existence in G of t independent triangles in the non-neighborhood of H , a contradiction to the $tC_3 \uplus C_4$ -induced-minor-freeness of G . We thus conclude by using a classic algorithm for MIS in chordal graphs [35, 48].

In Section 2 we introduce the relevant graph-theoretic background. In Section 3 we prove Theorem 1, and in Section 4 we prove Theorem 2.

2 Preliminaries

If $i \leq j$ are two integers, we denote by $[i, j]$ the set of integers $\{i, i+1, \dots, j-1, j\}$, and by $[i]$, the set $[1, i]$. We denote by $V(G)$ and $E(G)$ the set of vertices and edges of a graph G , respectively. We denote by $G_1 \simeq G_2$ the fact that the two graphs G_1 and G_2 are *isomorphic*, i.e., equal up to renaming their vertex sets. For $S \subseteq V(G)$, the *subgraph of G induced by S* , denoted $G[S]$, is obtained by removing from G all the vertices that are not in S (together with their incident edges). Then $G - S$ is a shorthand for $G[V(G) \setminus S]$. A graph H is an *induced subgraph* of G if there is an $S \subseteq V(G)$ such that $G[S] \simeq H$.

For G a graph and a set $X \subseteq V(G)$, $E_G(X)$ (or simply $E(X)$) is a shorthand for $E(G[X])$. For G a graph and $X, Y \subseteq V(G)$ two disjoint sets, $E_G(X, Y)$ denotes the set of edges of $E(G)$ with one endpoint in X and the other endpoint in Y . We denote by $N_G(v)$ and $N_G[v]$, the open, respectively closed, neighborhood of v in G . For $S \subseteq V(G)$, we set $N_G(S) := \bigcup_{v \in S} N_G(v) \setminus S$ and $N_G[S] := N_G(S) \cup S$. We may omit the subscript if G is clear from the context. A *connected component* is a maximal connected induced subgraph.

Two cycles C, C' are said to be *independent* if they are vertex-disjoint and there is no edge between C and C' . A collection of cycles is *independent* if they are pairwise independent. Two vertex subsets $X, Y \subseteq V(G)$ *touch* if $X \cap Y \neq \emptyset$ or there is an edge $uv \in E(G)$ with $u \in X$ and $v \in Y$. Then two (or more) cycles are independent if and only if they do *not* touch. We say that $X, Y \subseteq V(G)$ *touch in Z* if $X \cap Y \cap Z \neq \emptyset$ or there is an edge $uv \in E(G)$ with $u \in X \cap Z$ and $v \in Y \cap Z$, or equivalently, if $X \cap Z$ and $Y \cap Z$ touch.

A graph H is an *induced minor* of a graph G if H can be obtained from G by a sequence of vertex deletions and edge contractions. A *minor* is the same but also allows edge deletions. Equivalently an induced minor H —with vertex set, say, $\{v_1, \dots, v_{|V(H)|}\}$ —of G can be defined as a vertex-partition $B_1, \dots, B_{|V(H)|}$ of an induced subgraph of G , such that every $G[B_i]$ is connected and $v_i v_j \in E(H)$ if and only if $E_G(B_i, B_j) \neq \emptyset$ (i.e., when the disjoint sets B_i and B_j touch). Observe indeed that contracting each B_i into a single vertex (which is possible since each B_i induces a connected subgraph) results in H . A graph G (resp. a graph class) is said to be *H -induced-minor-free* if H is not an induced minor of G (resp. no graph of the class admits H as an induced minor).

We denote by C_ℓ the ℓ -vertex cycle, and by K_ℓ , the ℓ -vertex complete graph. A *hole* is an induced cycle of length at least four. A graph is *chordal* if it has no hole. For two disjoint sets $X, Y \subseteq V(G)$ in a graph G , an (X, Y) -separator is a (possibly empty) set $S \subseteq V(G) \setminus (X \cup Y)$ such that there is no path between X and Y in $G - S$. An (X, Y) -separator is *minimal* if no proper subset of it is itself an (X, Y) -separator.

The *disjoint union* $G_1 \uplus G_2$ of two graphs G_1, G_2 has vertex set $V(G_1) \uplus V(G_2)$ and edge set $E(G_1) \uplus E(G_2)$, where $V(G_1) \uplus V(G_2)$ presupposes that the vertex sets of G_1 and G_2 are disjoint. If $t \geq 2$ is an integer and G a graph, tG is the graph $G \uplus (t-1)G$, and $1G$ is simply G . The *join* $G_1 + G_2$ of two graphs G_1, G_2 has vertex set $V(G_1) \uplus V(G_2)$ and edge set $E(G_1) \uplus E(G_2) \uplus \{uv : u \in V(G_1), v \in V(G_2)\}$. In other words, the join of G_1 and G_2 is obtained from their disjoint union by adding all possible edges between G_1 and G_2 .

A *breadth-first search (BFS) layering* in G from a vertex $v \in V(G)$ (or from a connected set $S \subseteq V(G)$) is a partition of the remaining vertices into L_1, L_2, \dots such that every vertex of L_i is at distance exactly i from v (or from S). Such an L_i is called a *BFS layer* of G (from v , or from S). Note that there cannot be an edge in G between L_i and L_j if $|i - j| > 1$.

A *path-decomposition* of a graph G is a list of vertex subsets $\mathcal{P} = (B_1, \dots, B_h)$ such that

- $\bigcup_{1 \leq i \leq h} B_i = V(G)$,
- for every $e \in E(G)$, there is some B_i that contains both endpoints of e , and
- whenever $v \in B_i \cap B_j$ with $i < j$, v is also in all B_k with $i < k < j$.

The sets B_i (for $i \in [h]$) are called the *bags* of \mathcal{P} , and the sets $B_i \cap B_{i+1}$ (for $i \in [h-1]$) the *adhesions* of \mathcal{P} . Path-decomposition \mathcal{P} has maximum adhesion p if all of its adhesions have size at most p . Note that the adhesion $B_i \cap B_{i+1}$, if disjoint from $B_1 \cup B_h$, is a vertex cutset disconnecting B_1 from B_h .

3 Polynomial algorithm in $K_1 + tK_2$ -induced-minor-free graphs

We first show how to solve MAX INDEPENDENT SET in $K_1 + tK_2$ -induced-minor-free graphs of bounded diameter. More generally, we show the following.

► **Lemma 3.** *Let t, h be fixed non-negative integers. Let G be a $K_1 + tK_2$ -induced-minor-free n -vertex graph, and $L_0 \subseteq V(G)$ such that $G[L_0]$ is connected. Let L_i , for any $i \in [h]$, be the subset of vertices of G at distance exactly i from L_0 . Then, given as input G, L_0 , and $S \subseteq \bigcup_{1 \leq i \leq h} L_i$, a maximum independent set of $H := G[S]$ can be computed in polynomial time $n^{(2t-1)h+O(1)}$.*

Proof. For every $j \in [h]$, $G[L_j]$ has no tK_2 induced subgraph (or equivalently, induced minor). Indeed, $G[\bigcup_{0 \leq i \leq j-1} L_i]$ is a connected graph, hence $\bigcup_{0 \leq i \leq j-1} L_i$ can be contracted to a single vertex, and every vertex in L_j has at least one neighbor in L_{j-1} . Therefore a tK_2 induced subgraph in $G[L_j]$ would contradict the $K_1 + tK_2$ -induced-minor-freeness of G .

Fix an arbitrary $S \subseteq \bigcup_{1 \leq i \leq h} L_i$, and consider the induced subgraph $H := G[S]$. In particular $H[L_j \cap S]$ has also no tK_2 induced subgraph, for every $j \in [h]$. Hence, by a classical result of Alekseev [9], $H[L_j \cap S]$ has at most n^{2t-1} maximal independent sets, which can be listed in time $n^{2t+O(1)}$ [49].

We thus exhaustively list every h -tuple (I_1, \dots, I_h) where, for every $j \in [h]$, I_j is a maximal independent set of $H[L_j \cap S]$, in time $n^{(2t-1)h+O(1)}$. Note that if there is an edge in H between L_i and L_j , then $|i - j| \leq 1$. As each I_j (for $j \in [h]$) is an independent set, $H' = H[\bigcup_{j \in [h]} I_j]$ is a bipartite graph as witnessed by the bipartition $(I_1 \cup I_3 \cup \dots, I_2 \cup I_4 \cup \dots)$. A maximum independent set I can thus be computed in polynomial time in H' . Indeed, by the König-Egerváry theorem [40], finding a maximum independent set in a bipartite graph boils down

to finding a maximum matching, which can be done in polynomial time (and now, even in almost linear time [20]) by solving a maximum flow problem. We output the largest independent set I found among every run.

The correctness of the algorithm is based on the observation that a maximum independent set I^* of H intersects every L_i (for $i \in [h]$) in an independent set J_i , which by definition is contained in a maximal independent set I_i of $H[L_i \cap S]$. In the run when every maximal independent set I_i is a superset of J_i , we obtain an independent set with cardinality equal to that of I^* . \blacktriangleleft

We say that G is *reduced* if it does not contain degree-1 vertices or degree-2 vertices with adjacent neighbors (i.e., part of a triangle). If $e = uv$ is an edge of G , let $G \setminus e$ (resp. $S \setminus e$, $S \cup e$, for some $S \subseteq V(G)$) be the induced subgraph $G[V(G) - \{u, v\}] = G - \{u, v\}$ (resp. the sets $S \setminus \{u, v\}$, $S \cup \{u, v\}$). More generally, for a collection $e_1 = u_1v_1, \dots, e_k = u_kv_k$ of edges of G , we denote by $G \setminus \{e_1, \dots, e_k\}$ the induced subgraph $G[V(G) - \{u_1, v_1, \dots, u_k, v_k\}] = G - \{u_1, v_1, \dots, u_k, v_k\}$.

► **Lemma 4.** *Let G be a reduced connected $K_1 + tK_2$ -induced-minor-free graph containing tK_2 as an induced subgraph. Let $X \subseteq V(G)$ maximize the order of a largest component of $G' := G - X$, among those sets X such that $G[X] \simeq tK_2$. Then for any $e \in E(G' - N_G(X))$ contained in a connected component C of G' ,*

1. *$C \setminus e$ is disconnected, and*
2. *each connected component of $C \setminus e$ contains a vertex in $N_G(X)$.*

Proof. As G is reduced, every vertex of X has degree at least 2 (in G). Thus G' cannot be connected, for otherwise, contracting in G the set $V(G')$ to a single vertex would form a $K_1 + tK_2$ induced minor. We thus know that G' has at least two connected components.

Let C' be a largest connected component of G' . Since G is connected, there exists a shortest path P in G from $V(C')$ to $V(G') \setminus V(C')$. Say that P ends in a connected component $C \neq C'$ of G' . Path P has to have some internal vertices in X , but since $G[X] \simeq tK_2$, it follows that there is an edge e^* in $E(X)$ intersecting both $N_G(V(C))$ and $N_G(V(C'))$.

For every edge $e \in E(G' - N_G(X))$ in component C (which is possibly equal to C'), $C \setminus e$ is disconnected. Indeed, for the sake of contradiction, suppose that $C \setminus e$ is connected, and consider $X' := (X \setminus e^*) \cup e$. By assumption, $G[X'] \simeq tK_2$. Furthermore, the connected component of $G - X'$ containing e^* is strictly larger than C' , as it contains $(V(C') \setminus e) \cup e^*$ and intersects $V(C) \setminus e$, which are two disjoint sets. This contradicts the maximality of X , and establishes the first item.

We now prove the second item, also by contradiction. Suppose that there is a connected component D of $C \setminus e$ that does not contain a vertex in $N_G(X)$. We will reach a contradiction by showing that D contains an edge e' not intersecting $N_G(X)$, and such that $D \setminus e'$ is connected (and conclude in light of the previous paragraph).

Let $L_i \subseteq V(D)$ be the i -th neighborhood of e in D , i.e., the vertices at distance i of one endpoint of e , and at least i of the other endpoint. We consider the *last layer* L_k , i.e., such that $L_k \neq \emptyset$ and $L_{k+1} = \emptyset$. If L_k contains an edge e' , then removing the endpoints of this edge does not disconnect D (and hence C) since each vertex in L_k has a neighbor in L_{k-1} (with the convention that L_0 consists of the endpoints of e) and $G[L_0 \cup L_1 \cup \dots \cup L_{k-1}]$ is connected.

If L_k does not contain any edge, then $k \geq 2$. Otherwise (if $k = 1$), then D has a single vertex, say, w , and w has no neighbors in X . Hence, in G , w is only adjacent to one or the two endpoints of e , contradicting that G is reduced.

Furthermore, each vertex in L_k has two neighbors in L_{k-1} . This is because G has minimum degree at least 2 (as G is reduced), and by assumption that no vertex of D has a neighbor in X . Hence, removing the endpoints of any edge e' incident to a vertex in L_k does not disconnect D , nor C , since each vertex in L_k has at least one neighbor in L_{k-1} which is not an endpoint of e' . In either case, e' is an edge of $C - N_G(X)$ that does not disconnect C , which we showed is not possible. \blacktriangleleft

We now prove the main technical result of the section.

► **Proposition 5.** *Let G be a reduced connected n -vertex $K_1 + tK_2$ -induced-minor-free graph containing tK_2 as an induced subgraph. Let $X \subseteq V(G)$ maximize the order of a largest component of $G' := G - X$, among those sets X such that $G[X] \simeq tK_2$. Then for every connected component C of G' , a path-decomposition \mathcal{P} of C such that*

- *every bag of \mathcal{P} is contained in $O(t^4)$ consecutive BFS layers of C , and*
- *every adhesion of \mathcal{P} is of size at most $2t^2$,*

can be computed in time $n^{O(1)}$.

Proof. Let v be a vertex in $N_G(X) \cap V(C)$ and, for any positive integer s , let L_s be the set of vertices at distance s from v in C . Let q be the largest distance between v and a vertex of C . We set $f(t) := (t^2 + 1)(6t^2 + 2) = O(t^4)$. We will show that, for every $s \in [q - f(t)]$, there is a vertex cutset of size at most $2t^2$ separating L_s from $L_{s+f(t)-1}$, and use that fact for $s = 1, f(t) + 1, 2f(t) + 1, \dots$ to build the path-decomposition \mathcal{P} .

We show that any sufficiently long induced path (such as a shortest path from L_s to $L_{s+f(t)-1}$) has some edges with both endpoints in $V(C) \setminus N_G(X)$.

► **Claim 6.** Any induced path P in C contains less than $3t^2$ vertices in $N_G(X)$.

Proof. If there are $3t^2$ vertices on P with a neighbor in X , then there are at least $3t$ vertices w_1, \dots, w_{3t} on P that are neighbors of a fixed edge $e \in E(X)$. For each $i \in [t]$, contract every edge of P between w_{3i-2} and w_{3i-1} but one, say e_i . Contract e , and call z the resulting vertex. The vertex z and the t edges e_i contradict the fact that G is $K_1 + tK_2$ -induced-minor-free. \triangleleft

For an edge $e \in E(C)$ we denote by $\text{dist}(e, v)$ the length of a shortest path in C from an endpoint of e to v . We build a collection of paths Q_i of C , and edges $e_i \in E(Q_i)$, for $i = 1, 2, \dots$, while they are well-defined, in the following way.

Let $s \in [q - f(t)]$ and Q_1 be a shortest path from L_s to $L_{s+f(t)-1}$ in C . Let $e_1 \in E(Q_1)$ minimize $e \mapsto \text{dist}(e, v)$, among those edges of Q_1 with both endpoints in $V(C) \setminus N_G(X)$. By Claim 6 there are less than $6t^2$ edges on Q_1 with an endpoint in $N_G(X)$, hence $\text{dist}(e_1, v) \leq s + 6t^2$. We denote by Q'_1 the maximal subpath of Q_1 starting in L_s and not containing an endpoint of e_1 (that is, stopping just before reaching an endpoint of e_1). Note that Q_1 is possibly empty. For the next iteration, we work in $C \setminus e_1$ (recall that this stands for C deprived of the two endpoints of e_1).

We now describe in general the i -th iteration for $i \geq 2$. Let Q_i be a shortest path from L_s to $L_{s+f(t)-1}$ in $C \setminus \{e_1, e_2, \dots, e_{i-1}\}$. Let e_i be the first edge of Q_i (when starting from L_s) such that $\text{dist}(e_i, v) \geq \text{dist}(e_{i-1}, v) + 2$ and e_i has no endpoint in $N_G(X)$. Note that by Claim 6, $\text{dist}(e_i, v) \leq \text{dist}(e_{i-1}, v) + 2 + 6t^2$. Let Q'_i be the maximal subpath of Q_i starting in L_s and not containing an endpoint of e_i . See Figure 1.

Let e_1, \dots, e_k be the obtained collection of edges. In principle, the while loop stops when one of the following two conditions holds:

(i) L_s is disconnected from $L_{s+f(t)-1}$ in $C \setminus \{e_1, \dots, e_k\}$, or

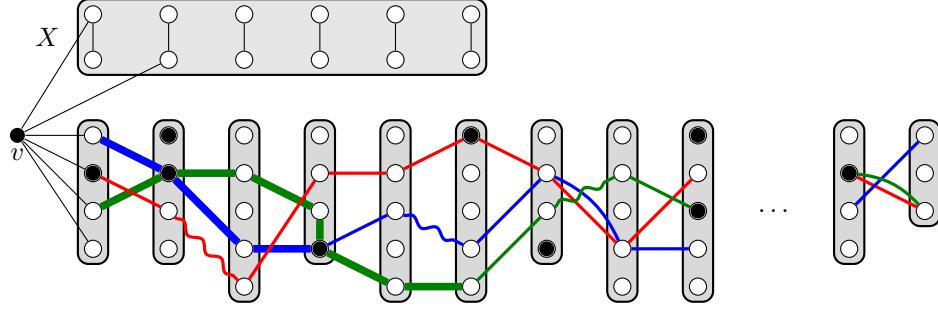


Figure 1 Illustration of X (for $t = 6$) and the first (i.e., with $s = 1$) $f(t)$ layers (from left to right) of the connected component C of $G - X$ rooted at v . Vertices of C with a neighbor in X are filled. We depict the first three iterations: Q_1, e_1 (in red), Q_2, e_2 (in blue), Q_3, e_3 (in dark green). Not to clutter the figure, we do not represent all the edges, and we code the different labels: edges e_i are squiggly, and subpaths Q'_i are thicker.

(ii) there is no edge $e_{k+1} \in E(Q_{k+1})$ such that $\text{dist}(e_{k+1}, v) \geq \text{dist}(e_k, v) + 2$ and e_{k+1} has no endpoint in $N_G(X)$.

▷ **Claim 7.** If case (ii) holds, then $k > t^2$.

Proof. Remark first that for any i we have $\text{dist}(e_i, v) \leq s + i(6t^2 + 2) - 2$. Next, we claim that as long as $\text{dist}(e_i, v) + 2 + 6t^2 \leq s + f(t) - 1$, and case (i) does not occur, the construction of e_{i+1} cannot fail, since Q_{i+1} exists, and there are at least $6t^2$ edges e in Q_{i+1} satisfying $\text{dist}(e, v) \geq \text{dist}(e_i, v)$. Thus if case (ii) occurs, it must be that $(k+1)(6t^2 + 2) - 2 > f(t) - 1$, hence $k > t^2$ since $f(t) = (t^2 + 1)(6t^2 + 2)$. \triangleleft

▷ **Claim 8.** It holds that $k \leq t^2$.

Proof. Assume for the sake of contradiction that $k \geq t^2 + 1$. By Lemma 4 (using both items), for each $i \in [k]$ there exists a vertex $v_i \in V(C) \cap N_G(X)$ disconnected from v in $C \setminus e_i$. Since $k \geq t^2 + 1$, there are $t + 1$ vertices $v_{a_1}, \dots, v_{a_{t+1}}$, with $a_1 < a_2 < \dots < a_t < a_{t+1}$, all adjacent to an endpoint of the same fixed edge $e \in E(X)$. Let C' be the component containing v in $C \setminus \{e_{a_1}, \dots, e_{a_t}\}$, and let us contract (in G) the set $V(C')$ to a single vertex, say w . We will now show that vertex w is adjacent to an endpoint of e .

By construction, the edge $e_{a_{t+1}}$ is reachable from v in $C \setminus \{e_{a_1}, e_{a_2}, \dots, e_{a_t}\}$, hence $e_{a_{t+1}} \in E(C')$. Let P' be a path in C' from v to an endpoint of $e_{a_{t+1}}$. Every path in C from v to $v_{a_{t+1}}$ goes through at least one endpoint of $e_{a_{t+1}}$, by definition of $v_{a_{t+1}}$. Let us consider a shortest path in C from v to $v_{a_{t+1}}$, i.e., intersecting each layer at most once. Consider a suffix P'' of this path starting at an endpoint of $e_{a_{t+1}}$ and ending at $v_{a_{t+1}}$. By the previous remark, P'' cannot contain an edge among $\{e_{a_1}, e_{a_2}, \dots, e_{a_t}\}$, since these edges are in layers with strictly smaller indices than the endpoints of $e_{a_{t+1}}$. Thus $P' \cup P''$ (possibly combined with $e_{a_{t+1}}$) connects v to $v_{a_{t+1}}$ in $C \setminus \{e_{a_1}, e_{a_2}, \dots, e_{a_t}\}$. Therefore, the contracted vertex w contains $v_{a_{t+1}}$; the latter being adjacent to an endpoint of e .

Observe also that, for every $i \in [t]$, each subpath Q'_{a_i} is contained in C' . Vertex w is adjacent to an endpoint u_{a_i} of e_{a_i} , for every $i \in [t]$. Let C_{a_i} be the vertex set of the connected component of $C \setminus e_{a_i}$ containing v_{a_i} . Let R_{a_i} be a path from u_{a_i} to v_{a_i} in the subgraph of G induced by C_{a_i} plus the endpoints of e_{a_i} . For each $i \in [t]$, if R_{a_i} has at least two edges, contract the edge wu_{a_i} , and all the edges of R_{a_i} but the last two; the last one being $f_{a_i} = y_{a_i}v_{a_i}$. Finally contract e , and the edge between e and w . We call the resulting vertex z .

We claim that z and the edges f_{a_i} make a $K_1 + tK_2$ induced minor in G . One can see that z is adjacent to v_{a_i} via $e \in E(X)$, and to y_{a_i} via w and the path R_{a_i} . We shall justify that the edges f_{a_i} form an induced matching in G . Indeed, suppose some f_{a_i} and f_{a_j} touch with $i \neq j \in [t]$. Then in $C \setminus e_{a_i}$, there is a path from v to v_{a_i} via Q'_{a_j} , a contradiction. \triangleleft

By Claims 7 and 8, case (ii) is impossible, and the $2k \leq 2t^2$ endpoints of e_1, \dots, e_k form a vertex cutset disconnecting L_s from $L_{s+f(t)-1}$. We can now build the path-decomposition \mathcal{P} . Recall that the BFS search from v gives rise to q layers, L_1, \dots, L_q (outside $\{v\}$).

For $j \in [|q/f(t)|]$, let S_j be the vertex cutset of size at most $2t^2$ (and obtained as detailed above) disconnecting $L_{(j-1)f(t)+1}$ from $L_{jf(t)}$. We denote by $L_{s \rightarrow s'}$ the set $\bigcup_{s \leq h \leq s'} L_h$.

- Let $B_1 \subseteq V(C)$ consist of v , S_1 , plus all the vertices of connected components of $C[L_{1 \rightarrow f(t)}] - S_1$ that do not intersect $L_{f(t)}$.
- For j going from 2 to $|q/f(t)| - 1$, let $B_j \subseteq V(C)$ consist of $S_j \cup S_{j+1}$ plus the vertices not already present in one of B_1, \dots, B_{j-1} of all the connected components of

$$C[L_{(j-1)f(t)+1 \rightarrow (j+1)f(t)}] - (S_j \cup S_{j+1})$$

that do not intersect $L_{(j+1)f(t)}$.

- Let finally $B_{[q/f(t)]} \subseteq V(C)$ consist of $S_{[q/f(t)]}$ plus the vertices of all the connected components of $C[L_{([q/f(t)]-1)f(t)+1 \rightarrow q}] - S_{[q/f(t)]}$ that intersect L_q .

Let \mathcal{P} be the path-decomposition $(B_1, B_2, \dots, B_{[q/f(t)]})$. \mathcal{P} is indeed a path-decomposition of C , since our process entirely covers $V(C)$, and by virtue of S_j separating $L_{(j-1)f(t)+1}$ from $L_{jf(t)}$. By construction,

- every bag intersects at most $2f(t) = O(t^4)$ layers of BFS from vertex v , and
- for every $j \in [|q/f(t)| - 1]$, $B_j \cap B_{j+1} = S_j$, so every adhesion has size at most $2t^2$.

Finally note that once X is found, one can find the path-decomposition \mathcal{P} of C in time $n^{O(1)}$, since this only involves computing (at most n) shortest paths. \blacktriangleleft

We can now wrap up, using Proposition 5 and Lemma 3.

► **Theorem 1.** *For every positive integer t , MAX INDEPENDENT SET can be solved in polynomial-time $n^{O(t^5)}$ in n -vertex $K_1 + tK_2$ -induced-minor-free graphs.*

Proof. Let G be our $K_1 + tK_2$ -induced-minor-free n -vertex input graph. As including vertices of degree 1 or vertices of degree 2 with adjacent neighbors in the independent set is a safe reduction rule (since it is more generally safe to include a *simplicial* vertex, i.e., one whose neighborhood is a clique), we can assume that G is reduced. By dealing with the possibly several connected components of G separately, we can further assume that G is connected. If G has no tK_2 as an induced subgraph, we conclude by invoking Alekseev's result [9, 49]. Thus we also assume that G has such an induced subgraph. In time $n^{O(t)}$ we find $X \subseteq V(G)$ that maximizes the order of a largest component of $G' := G - X$, among those sets X such that $G[X] \simeq tK_2$.

We exhaustively guess the intersection X' of a fixed maximum independent set of G with the set X , with an extra multiplicative factor of 2^{2t} . We are now left with solving MIS separately in $C' := G - N_G(X')$ for each connected component C of G' . By Proposition 5, we obtain in time $n^{O(1)}$ a path-decomposition $\mathcal{P} = (B_1, \dots, B_p)$ of C' , such that

- every B_i (for $i \in [p]$) is contained in $O(t^4)$ consecutive BFS layers of C , and
- every adhesion $A_i := B_i \cap B_{i+1}$ (for $i \in [p-1]$) is of size at most $2t^2$.

Indeed, removing $N_G(X')$ from C (and from its path-decomposition) preserves those properties.

Let us define A_0, A_p to be empty. We proceed to the following dynamic programming. For $i \in [0, p]$, and for any $S \subseteq A_i$, $T[i, S]$ is meant to eventually contain an independent set I of $C'[\bigcup_{1 \leq j \leq i} B_j]$ of maximum cardinality among those such that $I \cap B_i = S$. We set $T[0, \emptyset] = \emptyset$, and observe that it is the only entry of the form $T[0, \cdot]$.

We fill this table by increasing value of $i = 1, 2, \dots, p$. Assume that all entries of the form $T[i', \cdot]$ are properly filled for $i' < i$. For every $S \subseteq A_i$, $T[i, S]$ is filled in the following way. For every $S' \subseteq A_{i-1}$, if $S \cup S'$ is an independent set, we compute, by Lemma 3 (with L_1, \dots, L_h being the $O(t^4)$ consecutive BFS layers of C containing B_i , and L_0 being the connected set, in C , formed by the union of all the previous layers), a maximum independent set I_i in $C'[B_i] - N[S \cup S']$ in time $n^{O(ht)} = n^{O(t^5)}$. We finally set $T[i, S] = T[i-1, S'] \cup I_i \cup S$ for a run that maximizes the cardinality of $T[i-1, S'] \cup I_i$.

It takes time $p \cdot 2^{O(t^2)} \cdot n^{O(t^5)} = n^{O(t^5)}$ to completely fill T . Eventually $T[p, \emptyset]$ contains a maximum independent set of C' . We return the union of X' and of the maximum independent sets of C' found for each connected component C of G' . The overall running time is $n^{O(t^5)}$. \blacktriangleleft

4 Quasipolynomial algorithm in $tC_3 \uplus C_4$ -induced-minor-free graphs

We first show that the existence of a short hole allows for a quasipolynomial-time branching rule in $tC_3 \uplus C_4$ -induced-minor-free graphs. By branching rule, we mean here a Turing reduction to (quasipolynomially many) subinstances with no short holes.

► **Lemma 9.** *Let G be an n -vertex $tC_3 \uplus C_4$ -induced-minor-free graph, and let $\ell \geq 4$ be a fixed integer. While there is a hole H of length at most ℓ and a tC_3 as an induced subgraph in G , MAX INDEPENDENT SET admits a quasipolynomial branching rule, running in $n^{O_\ell(t \log n)}$.*

Proof. Let $C(G) := \{X \in \binom{V(G)}{3t} : G[X] \simeq tC_3\}$ be the collection of vertex subsets inducing t disjoint triangles, and assume $\mu(G) := |C(G)| > 0$, and H is a hole of G of length at most ℓ . As G is $tC_3 \uplus C_4$ -induced-minor-free, $N[V(H)]$ intersects every $X \in C(G)$. In particular, there is a vertex $v \in V(H)$ such that $N[v]$ intersects at least a $1/\ell$ fraction of the $X \in C(G)$.

We branch on two options: either we take v in an (initially empty) solution, and remove its closed neighborhood from G , or we remove v from G (without adding it to the solution). With the former choice, the measure μ drops by at least a $1/\ell$ fraction (and the number of vertices of G decreases by at least 1), and with the latter choice, the number of vertices drops by 1. This branching is exhaustive. We simply need to argue about its running time.

Note that each option can be done at most n times, while the first option cannot be done more than $\log_\ell(n^{3t}) = O_\ell(t \log n)$ times. Hence the branching tree has at most $\binom{n}{\log_\ell(n^{3t})} = n^{O_\ell(t \log n)}$ leaves. \blacktriangleleft

The previous lemma permits us to get rid of short holes, which turns out useful in some corner case.

► **Theorem 2.** *For every positive integer t , MAX INDEPENDENT SET can be solved in quasipolynomial-time $n^{O(t^2 \log n) + f(t)}$ (where f is single-exponential) in $tC_3 \uplus C_4$ -induced-minor-free graphs.*

Proof. We apply the quasipolynomial branching rule of Lemma 9 with $\ell = 6$, until the input n -vertex graph G no longer has holes of length at most 6, or tC_3 induced subgraph.

In time $n^{O(t)}$, we exhaustively look for a collection of pairwise vertex-disjoint and non-adjacent triangles T_1, T_2, \dots, T_{t+2} in G . If such a collection does not exist, G is $(t+2)C_3$ -induced-minor-free. Indeed, the absence of $(t+2)C_3$ as an induced subgraph implies that at least one of the $(t+2)$ independent cycles realizing a $(t+2)C_3$ induced minor is of length at least four. This is ruled out by the assumption that G is $tC_3 \uplus C_4$ -induced-minor-free. (Note here that only $(t+1)$ independent cycles would suffice.)

We can thus assume that a collection T_1, T_2, \dots, T_{t+2} exists, for otherwise, we can conclude with the quasipolynomial-time algorithm, running in $n^{O(t^2 \log n) + f(t)}$ (where f is single-exponential), of Bonamy et al. [13] for MAX INDEPENDENT SET in graphs with a bounded number of independent cycles (here, $(t+2)C_3$ -induced-minor-free). In turn, as G contains tC_3 (even $(t+2)C_3$) as an induced subgraph, we can, in light of the first paragraph, further assume that all the cycles of G have length either 3 or at least 7. We refer the reader to Figure 2 for a visual summary of the next two paragraphs.

For every pair T_i, T_j (with $i < j \in [t+2]$), consider the subgraph $G_{i,j} := G - \bigcup_{k \in [t+2] \setminus \{i,j\}} N(T_k)$. We claim that $G_{i,j}$ is chordal. Indeed, since $G_{i,j}$ is disjoint from the neighborhood of $\bigcup_{k \in [t+2] \setminus \{i,j\}} T_k$, a hole H in G would form a $tC_3 \uplus C_4$ -induced-minor together with $\{T_k : k \in [t+2] \setminus \{i,j\}\}$. Let now $S_{i,j}$ be a minimal (T_i, T_j) -separator in $G_{i,j}$. A classical argument then shows that $S_{i,j}$ is a clique: suppose for the sake of contradiction that $u, v \in S_{i,j}$ are distinct and non-adjacent. Let X_i, X_j be the two components of $G_{i,j} - S_{i,j}$ containing T_i, T_j . By minimality of $S_{i,j}$, each of u, v is adjacent to both X_i and X_j . Thus we can find two induced uv -paths, whose internal vertices are in X_i and X_j respectively. The union of these paths is a hole H in $G_{i,j}$.

Let $N_{i,j}$ be the set $N(T_i) \cap N(T_j)$ for each pair $i < j \in [t+2]$. Observe that the sets $N_{i,j}$ need not be disjoint, and that when T_i and T_j are at distance at least 3 apart, $N_{i,j}$ is empty. We notice that $N_{i,j}$ is a clique, for otherwise we can exhibit an induced cycle of length 4, 5, or 6 in G (hence a hole of length at most 6).

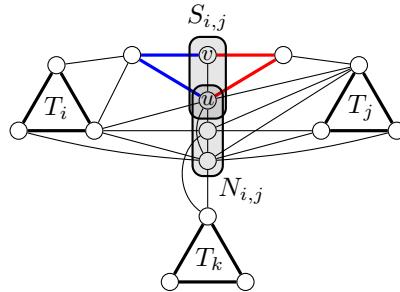


Figure 2 Illustration of the sets $S_{i,j}$ and $N_{i,j}$, and the two uv -paths through the two components, in red and blue. If uv were a non-edge, these paths would form a hole in the non-neighborhood of the other triangles T_k , contradicting $tC_3 \uplus C_4$ -minor-freeness. The absence of hole of length at most 6 implies that $N_{i,j}$ is also a clique.

We claim that $G' := G - (\bigcup_{i < j \in [t+2]} S_{i,j} \cup N_{i,j})$ is chordal. Indeed assume there is a hole H' in G' . The $tC_3 \uplus C_4$ -induced-minor-freeness implies that H' intersects at least two sets $N[T_i]$ and $N[T_j]$. Thus there exists a subpath P of H' whose endpoints are in two distinct $N(T_i)$ and $N(T_j)$. By choosing P minimal, we can furthermore assume that no internal vertex of P lies in some $N[T_k]$ with $k \notin \{i, j\}$. Since G' does not include any vertex of $\bigcup_{i' < j' \in [t+2]} N_{i',j'}$, the endpoints of P are not in some $N[T_k]$ with $k \notin \{i, j\}$ either. Therefore, the path P contradicts that $S_{i,j}$ separates T_i and T_j in $G - \bigcup_{k \in [t+2] \setminus \{i,j\}} N(T_k)$.

We can now describe the rest of the algorithm after the collection T_1, T_2, \dots, T_{t+2} is

found. We greedily compute the minimal separators $S_{i,j}$. We exhaustively try every subset $S \subseteq \bigcup_{i < j \in [t+2]} S_{i,j} \cup N_{i,j}$ that is an independent set. Such a set S contains at most one vertex in each $S_{i,j}$ and each $N_{i,j}$, as we have established that each $S_{i,j}$ and each $N_{i,j}$ form a clique. Hence there are $n^{O(t^2)}$ such sets S . For each S , we compute a maximum independent set I in the *chordal* graph $G - (N[S] \cup \bigcup_{i < j \in [t+2]} S_{i,j} \cup N_{i,j})$ in linear time (see [35, 48]). We finally output the set $S \cup I$ maximizing $|S \cup I|$. Note that the overall running time is $n^{O(t^2 \log n) + f(t)}$. \blacktriangleleft

5 Conclusion

We provided a polynomial-time algorithm for MAXIMUM INDEPENDENT SET on graphs excluding the friendship graph as an induced minor and a quasipolynomial-time algorithm on graphs excluding a disjoint union of t triangles and a 4-cycle as an induced minor. As mentioned in the introduction, it is of interest to study which other graph classes excluding some fixed planar graph H as an induced minor MIS can be solved in (quasi)polynomial time. If H is a subgraph of the friendship graph or of the wheel, some of our methods developed for $K_1 + tK_2$ might extend.

For disjoint unions of cycles, the first open case is when H is $C_4 \uplus C_4$. Our treatment for $H = tC_3 \uplus C_4$ does not (easily) extend to this case. It is thus likely that new methods have to be found. Obtaining a quasipolynomial-time algorithm when $H = tC_t$ for any integer t is a first challenging milestone in the study of MIS on H -induced-minor-free graphs.

COI statement: *The authors have no competing interests as defined by Springer, or other interests that might be perceived to influence the results and/or discussion reported in this paper.*

References

- 1 Kuruvilla J. Abraham and Clara Diaz. Identifying large sets of unrelated individuals and unrelated markers. *Source code for biology and medicine*, 9(1):1–8, 2014.
- 2 Tara Abrishami, Maria Chudnovsky, Cemil Dibek, and Paweł Rzazewski. Polynomial-time algorithm for maximum independent set in bounded-degree graphs with no long induced claws. In Joseph (Seffi) Naor and Niv Buchbinder, editors, *Proceedings of the 2022 ACM-SIAM Symposium on Discrete Algorithms, SODA 2022, Virtual Conference / Alexandria, VA, USA, January 9 - 12, 2022*, pages 1448–1470. SIAM, 2022. doi:10.1137/1.9781611977073.61.
- 3 Tara Abrishami, Maria Chudnovsky, Marcin Pilipczuk, and Paweł Rzazewski. 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. URL: <https://doi.org/10.4230/LIPIcs.STACS.2024.4>, doi:10.4230/LIPIcs.STACS.2024.4.
- 4 Tara Abrishami, Maria Chudnovsky, Marcin Pilipczuk, Paweł Rzazewski, and Paul D. Seymour. Induced subgraphs of bounded treewidth and the container method. In Dániel Marx, editor, *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms, SODA 2021, Virtual Conference, January 10 - 13, 2021*, pages 1948–1964. SIAM, 2021. doi:10.1137/1.9781611976465.116.
- 5 Jungho Ahn, J. Pascal Gollin, Tony Huynh, and O-joung Kwon. A coarse erdős-pósa theorem. In Yossi Azar and Debmalya Panigrahi, editors, *Proceedings of the 2025 Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2025, New Orleans, LA, USA, January 12-15, 2025*, pages 3363–3381. SIAM, 2025. doi:10.1137/1.9781611978322.109.

- 6 Sungsoo Ahn, Younggyo Seo, and Jinwoo Shin. Learning what to defer for maximum independent sets. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pages 134–144. PMLR, 2020. URL: <http://proceedings.mlr.press/v119/ahn20a.html>.
- 7 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.
- 8 Vladimir E. Alekseev. Polynomial algorithm for finding the largest independent sets in graphs without forks. *Discrete Applied Mathematics*, 135(1-3):3–16, 2004. doi:10.1016/S0166-218X(02)00290-1.
- 9 Vladimir E. Alekseev. An upper bound for the number of maximal independent sets in a graph. *Discrete Mathematics and Applications*, 17(4):355–359, 2007. URL: <https://doi.org/10.1515/dma.2007.030> [cited 2023-02-03], doi:doi:10.1515/dma.2007.030.
- 10 Paola Alimonti and Viggo Kann. Some APX-completeness results for cubic graphs. *Theoretical Computer Science*, 237(1):123 – 134, 2000. doi:[https://doi.org/10.1016/S0304-3975\(98\)00158-3](https://doi.org/10.1016/S0304-3975(98)00158-3).
- 11 Ismail R. Alkhouri, George K. Atia, and Alvaro Velasquez. A differentiable approach to the maximum independent set problem using dataless neural networks. *Neural Networks*, 155:168–176, 2022.
- 12 J. Gary Augustson and Jack Minker. An analysis of some graph theoretical cluster techniques. *Journal of the ACM (JACM)*, 17(4):571–588, 1970.
- 13 Marthe Bonamy, Édouard 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 *Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 3006–3028. SIAM, 2023.
- 14 Édouard Bonnet, Nicolas Bousquet, Pierre Charbit, Stéphan Thomassé, and Rémi Watrigant. Parameterized complexity of independent set in H-free graphs. In Christophe Paul and Michał Pilipczuk, editors, *13th International Symposium on Parameterized and Exact Computation, IPEC 2018, August 20-24, 2018, Helsinki, Finland*, volume 115 of *LIPICS*, pages 17:1–17:13. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018. doi:10.4230/LIPICS.IPEC.2018.17.
- 15 Édouard Bonnet, Nicolas Bousquet, Stéphan Thomassé, and Rémi Watrigant. When maximum stable set can be solved in FPT time. In Pinyan Lu and Guochuan Zhang, editors, *30th International Symposium on Algorithms and Computation, ISAAC 2019, December 8-11, 2019, Shanghai University of Finance and Economics, Shanghai, China*, volume 149 of *LIPICS*, pages 49:1–49:22. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019. doi:10.4230/LIPICS.ISAAC.2019.49.
- 16 Vincent Bouchitté and Ioan Todinca. Treewidth and minimum fill-in: Grouping the minimal separators. *SIAM J. Comput.*, 31(1):212–232, 2001. doi:10.1137/S0097539799359683.
- 17 Andreas Brandstädt and Raffaele Mosca. Maximum weight independent set for ℓ -claw-free graphs in polynomial time. *Discrete Applied Mathematics*, 237:57–64, 2018. doi:10.1016/j.dam.2017.11.029.
- 18 Sergiy Butenko. *Maximum independent set and related problems, with applications*. University of Florida, 2003.
- 19 Sergiy Butenko, Panos M. Pardalos, Ivan Sergienko, Vladimir Shylo, and Petro Stetsyuk. Finding maximum independent sets in graphs arising from coding theory. In Gary B. Lamont, Hisham Haddad, George A. Papadopoulos, and Brajendra Panda, editors, *Proceedings of the 2002 ACM Symposium on Applied Computing (SAC), March 10-14, 2002, Madrid, Spain*, pages 542–546. ACM, 2002. doi:10.1145/508791.508897.
- 20 Li Chen, Rasmus Kyng, Yang P. Liu, Richard Peng, Maximilian Probst Gutenberg, and Sushant Sachdeva. Maximum flow and minimum-cost flow in almost-linear time. In *63rd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2022, Denver, CO, USA, October 31 - November 3, 2022*, pages 612–623. IEEE, 2022. doi:10.1109/FOCS54457.2022.00064.

21 Maria Chudnovsky, Marcin Pilipczuk, Michał Pilipczuk, and Stéphan Thomassé. On the maximum weight independent set problem in graphs without induced cycles of length at least five. *SIAM J. Discret. Math.*, 34(2):1472–1483, 2020. doi:10.1137/19M1249473.

22 Bruno Courcelle. The monadic second-order logic of graphs. I. Recognizable sets of finite graphs. *Inf. Comput.*, 85(1):12–75, 1990. doi:10.1016/0890-5401(90)90043-H.

23 Konrad K. Dabrowski, Vadim V. Lozin, Haiko Müller, and Dieter Rautenbach. Parameterized complexity of the weighted independent set problem beyond graphs of bounded clique number. *J. Discrete Algorithms*, 14:207–213, 2012. doi:10.1016/j.jda.2011.12.012.

24 Clément Dallard, Martin Milanič, and Kenny Štorgel. Treewidth versus clique number. II. Tree-independence number. *Journal of Combinatorial Theory, Series B*, 164:404–442, 2024. URL: <https://www.sciencedirect.com/science/article/pii/S0095895623000886>, doi:10.1016/j.jctb.2023.10.006.

25 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. URL: <https://www.sciencedirect.com/science/article/pii/S0095895624000236>, doi:10.1016/j.jctb.2024.03.005.

26 Rodney G. Downey and Michael R. Fellows. Fixed-parameter tractability and completeness II: On completeness for W[1]. *Theor. Comput. Sci.*, 141(1&2):109–131, 1995. doi:10.1016/0304-3975(94)00097-3.

27 Duncan Eddy and Mykel J. Kochenderfer. A maximum independent set method for scheduling earth-observing satellite constellations. *Journal of Spacecraft and Rockets*, 58(5):1416–1429, 2021.

28 Elaine Forsyth and Leo Katz. A matrix approach to the analysis of sociometric data: preliminary report. *Sociometry*, 9(4):340–347, 1946.

29 Eleanor J. Gardiner, Peter Willett, and Peter J. Artymiuk. Graph-theoretic techniques for macromolecular docking. *J. Chem. Inf. Comput. Sci.*, 40(2):273–279, 2000. doi:10.1021/ci990262o.

30 Michael R. Garey and David S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, 1979.

31 Michael R. Garey, David S. Johnson, and Larry J. Stockmeyer. Some simplified NP-complete graph problems. *Theor. Comput. Sci.*, 1(3):237–267, 1976. doi:10.1016/0304-3975(76)90059-1.

32 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.

33 Peter Gartland, Daniel Lokshtanov, Tomás Masarík, 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. doi:10.1145/3618260.3649791.

34 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. doi:10.1145/3406325.3451034.

35 Fánica Gavril. The intersection graphs of subtrees in trees are exactly the chordal graphs. *Journal of Combinatorial Theory, Series B*, 16(1):47–56, 1974.

36 Andrzej Grzesik, Tereza Klimosová, 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. doi:10.1145/3414473.

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

38 Russell Impagliazzo and Ramamohan Paturi. On the complexity of k -SAT. *J. Comput. Syst. Sci.*, 62(2):367–375, 2001. doi:10.1006/jcss.2000.1727.

39 Russell Impagliazzo, Ramamohan Paturi, and Francis Zane. Which problems have strongly exponential complexity? *J. Comput. Syst. Sci.*, 63(4):512–530, 2001. doi:10.1006/jcss.2001.1774.

40 Denés König. Gráfok és mátrixok. *Matematikai és Fizikai Lapok*, 38:116–119, 1931.

41 Tuukka Korhonen. Grid induced minor theorem for graphs of small degree. *J. Comb. Theory, Ser. B*, 160:206–214, 2023. doi:10.1016/j.jctb.2023.01.002.

42 Sebastian Lamm, Peter Sanders, Christian Schulz, Darren Strash, and Renato F. Werneck. Finding near-optimal independent sets at scale. *J. Heuristics*, 23(4):207–229, 2017. doi:10.1007/s10732-017-9337-x.

43 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. doi:10.1016/j.jda.2008.04.001.

44 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. doi:10.1137/1.9781611976496.23.

45 Svatopluk Poljak. A note on stable sets and colorings of graphs. *Commentationes Mathematicae Universitatis Carolinae*, 15(2):307–309, 1974.

46 Thomas Pontoizeau, Florian Sikora, Florian Yger, and Tristan Cazenave. Neural maximum independent set. In *Machine Learning and Principles and Practice of Knowledge Discovery in Databases - International Workshops of ECML PKDD 2021, Virtual Event, September 13-17, 2021, Proceedings, Part I*, volume 1524 of *Communications in Computer and Information Science*, pages 223–237. Springer, 2021. doi:10.1007/978-3-030-93736-2_18.

47 Neil Robertson and Paul D. Seymour. Graph minors. V. Excluding a planar graph. *J. Comb. Theory, Ser. B*, 41(1):92–114, 1986. doi:10.1016/0095-8956(86)90030-4.

48 Donald J. Rose, Robert E. Tarjan, and George S. Lueker. Algorithmic aspects of vertex elimination on graphs. *SIAM J. Comput.*, 5(2):266–283, 1976. doi:10.1137/0205021.

49 Shuji Tsukiyama, Mikio Ide, Hiromu Ariyoshi, and Isao Shirakawa. A new algorithm for generating all the maximal independent sets. *SIAM J. Comput.*, 6(3):505–517, 1977. doi:10.1137/0206036.

50 Bram Verweij and Karen Aardal. An optimisation algorithm for maximum independent set with applications in map labelling. In Jaroslav Nešetřil, editor, *Algorithms - ESA '99, 7th Annual European Symposium, Prague, Czech Republic, July 16-18, 1999, Proceedings*, volume 1643 of *Lecture Notes in Computer Science*, pages 426–437. Springer, 1999. doi:10.1007/3-540-48481-7_37.

51 Mingyu Xiao and Hiroshi Nagamochi. Exact algorithms for maximum independent set. *Inf. Comput.*, 255:126–146, 2017. doi:10.1016/j.ic.2017.06.001.

52 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.