



Few induced disjoint paths for H -free graphs[☆]

Barnaby Martin^a, Daniël Paulusma^{a,*}, Siani Smith^a, Erik Jan van Leeuwen^b

^a Department of Computer Science, Durham University, Durham, UK

^b Department of Information and Computing Sciences, Utrecht University, the Netherlands



ARTICLE INFO

Article history:

Received 13 June 2022

Received in revised form 14 October 2022

Accepted 19 October 2022

Available online 24 October 2022

Communicated by T. Calamoneri

Keywords:

Induced disjoint paths

H -free graph

Complexity dichotomy

ABSTRACT

Paths $P^1, \dots, P^k$ in a graph $G = (V, E)$ are mutually induced if any two distinct P^i and P^j have neither common vertices nor adjacent vertices. For a fixed integer k , the k -INDUCED DISJOINT PATHS problem is to decide if a graph G with k pairs of specified vertices (s_i, t_i) contains k mutually induced paths P^i such that each P^i starts from s_i and ends at t_i . Whereas the non-induced version is well-known to be polynomial-time solvable for every fixed integer k , a classical result from the literature states that even 2-INDUCED DISJOINT PATHS is NP-complete. We prove new complexity results for k -INDUCED DISJOINT PATHS if the input is restricted to H -free graphs, that is, graphs without a fixed graph H as an induced subgraph. We compare our results with a complexity dichotomy for INDUCED DISJOINT PATHS, the variant where k is part of the input.

© 2022 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

We consider problems related to finding paths connecting pre-specified pairs of vertices. A path between vertices s and t in an undirected graph G is an s - t path with *terminals* s and t . Terminal pairs $(s_1, t_1), \dots, (s_k, t_k)$ are *pairwise disjoint* if $\{s_i, t_i\} \cap \{s_j, t_j\} = \emptyset$ for $i \neq j$. The well-known problem k -DISJOINT PATHS is to decide for a graph G and pairwise disjoint terminal pairs $(s_1, t_1), \dots, (s_k, t_k)$, if there are pairwise vertex-disjoint paths $P^1, \dots, P^k$ such that P^i is an (s_i, t_i) -path for $i \in \{1, \dots, k\}$; here k is *fixed*, that is, k is not part of the input.

Shiloach [25] proved that 2-DISJOINT PATHS is polynomial-time solvable. Robertson and Seymour [24] even gave a polynomial-time algorithm for k -DISJOINT PATHS for every integer $k \geq 2$. In contrast, DISJOINT PATHS, the variant where k is part of the input, appeared on Karp's list of NP-complete problems.

Our focus. We consider the *induced* variant of k -DISJOINT PATHS. We say that paths $P^1, \dots, P^k$ in a graph $G = (V, E)$ are *mutually induced* if any two distinct P^i and P^j have neither common vertices nor adjacent vertices, that is, if $i \neq j$ then $V(P^i) \cap V(P^j) = \emptyset$ and $uv \notin E$ for every $u \in V(P^i)$ and $v \in V(P^j)$. This leads to the following problem, where k is a fixed constant.

[☆] An extended abstract of this paper will appear in the proceedings of ISCO 2022 [21].

* Corresponding author.

E-mail addresses: barnaby.d.martin@durham.ac.uk (B. Martin), daniel.paulusma@durham.ac.uk (D. Paulusma), siani.smith@durham.ac.uk (S. Smith), e.j.vanleeuwen@uu.nl (E.J. van Leeuwen).

<https://doi.org/10.1016/j.tcs.2022.10.024>

0304-3975/© 2022 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

k -INDUCED DISJOINT PATHS

Instance: a graph G and pairwise disjoint terminal pairs $(s_1, t_1) \dots (s_k, t_k)$.

Question: Does G have mutually induced paths $P^1, \dots, P^k$ such that P^i is an s_i - t_i path for $i \in \{1, \dots, k\}$?

Note that a graph G with a set of pairwise disjoint terminal pairs $(s_1, t_1) \dots (s_k, t_k)$ is a no-instance of k -INDUCED DISJOINT PATHS if G has an edge between two terminals from different pairs.

In contrast to the previous setting, even 2-INDUCED DISJOINT PATHS is NP-complete, as shown both by Bienstock [4] and Fellows [7]. Restricting the input to some special graph class might help improve our understanding of the hardness of the problem. To do this systematically we focus on hereditary graph classes.

A class of graphs is *hereditary* if it is closed under vertex deletion. This is a natural property and non-surprisingly hereditary graph classes provide a framework that captures many well-known graph classes. In particular, it is not difficult to see that a graph class $\mathcal{G}$ is hereditary if and only if it can be characterized by a (unique) set $\mathcal{F}_{\mathcal{G}}$ of forbidden induced subgraphs. For example, if $\mathcal{G}$ is the class of bipartite graphs, then $\mathcal{F}_{\mathcal{G}}$ is the set of all odd cycles.

The characterization by $\mathcal{F}_{\mathcal{G}}$ allows for a systematic study, which usually starts with the case where $\mathcal{F}_{\mathcal{G}}$ has size 1, say $\mathcal{F}_{\mathcal{G}} = \{H\}$ for some graph H . A graph is *H-free* if it cannot be modified to H by a sequence of vertex deletions, and if $\mathcal{F}_{\mathcal{G}} = \{H\}$ we obtain the class of *H-free graphs*, which we consider in our paper.

1.1. Related work

We first discuss existing results for INDUCED DISJOINT PATHS (where k is part of the input). All the positive results hold for a slightly more general problem definition (see Section 6). Golovach et al. [11,12] proved that INDUCED DISJOINT PATHS is linear-time solvable for circular-arc graphs and polynomial-time solvable for AT-free graphs, respectively. Belmonte et al. [3] showed the latter for chordal graphs, and Jaffke et al. [14] did so for any graph class of bounded mim-width. In contrast, INDUCED DISJOINT PATHS stays NP-complete even for claw-free graphs [8], line graphs of triangle-free chordless graphs [23] and thus for (theta,wheel)-free graphs, and for planar graphs; to prove the latter, use a result of Lynch [19] (see [12]).

The following recent dichotomy is immediately relevant for our paper. Let $G_1 + G_2$ be the disjoint union of two vertex-disjoint graphs G_1 and G_2 , and let sG denote the disjoint union of s copies of a graph G . We write $F \subseteq_i G$ if F is an *induced* subgraph of a graph G , that is, F can be obtained from G by a sequence of vertex deletions. We let P_r denote the path on r vertices. A *linear forest* is the disjoint union of one or more paths.

Theorem 1 ([22]). *For a graph H , INDUCED DISJOINT PATHS on H -free graphs is polynomial-time solvable if $H \subseteq_i sP_3 + P_6$ for some $s \geq 0$; NP-complete if H is not a linear forest; and quasipolynomial-time solvable otherwise.*

We return to Theorem 1 later, and we now fix k . Radovanović et al. [23] proved that k -INDUCED DISJOINT PATHS is polynomial-time solvable for (theta,wheel)-free graphs. Fiala et al. [8] proved the same result for claw-free graphs. Note that both results complement the aforementioned hardness results when k is part of the input. Golovach et al. [10] showed that INDUCED DISJOINT PATHS is even FPT with parameter k for claw-free graphs. The same holds for planar graphs [15], and even for graph classes of bounded genus, as shown by Kobayashi and Kawarabayashi [17]. Let C_r denote the r -vertex cycle. It follows (using Lemma 3) from a result of Leveque et al. [18] that 2-INDUCED DISJOINT PATHS is NP-complete for H -free graphs if $H = C_r$ for every $r \geq 3$ with $r \neq 6$.

The generalization from paths to connected subgraphs joining sets of terminals instead of pairs has also been considered, but these results do not impact upon our work in this paper; we refer to [22] for further details. Moreover, the restriction to H -free graphs has also been studied for DISJOINT PATHS (recall that if k is fixed this problem is polynomial in general [24]); see [16] for a complexity classification of DISJOINT PATHS for H -free graphs, subject to a set of three unknown cases.

1.2. Our results

To explain our results we first introduce some extra terminology. For $r \geq 1$, the graph $K_{1,r}$ is the $(r+1)$ -vertex *star*, i.e., the graph with vertices $x, y_1, \dots, y_r$ and edges xy_i for $i = 1, \dots, r$. The graph $K_{1,3}$ is known as the *claw*. The *subdivision* of an edge uw removes uw and replaces it with a new vertex v and edges uv, vw . A *subdivided claw* is a tree with one vertex x of degree 3 and exactly three leaves. For $1 \leq h \leq i \leq j$, let $S_{h,i,j}$ be the subdivided claw whose three leaves are of distance h, i and j from the vertex of degree 3. Note that $S_{1,1,1} = K_{1,3}$. The graph $S_{1,1,2}$ is called the *chair* (or *fork*). Let $\mathcal{S}$ be the set of graphs, each connected component of which is a path or a subdivided claw.

Using the above terminology we can now present our main theorem.

Theorem 2. *Let $k \geq 2$. For a graph H , k -INDUCED DISJOINT PATHS restricted to H -free graphs is polynomial-time solvable if H is a subgraph of the disjoint union of a linear forest and a chair, and it is NP-complete if H is not in $\mathcal{S}$.*

Comparing Theorems 1 and 2 shows that the problem becomes tractable for an infinite family of graphs H after fixing k . As the class of claw-free graphs is contained in the class of chair-free graphs, Theorem 2 extends the aforementioned

polynomial-time result of Fiala et al. [8] for claw-free graphs. Moreover, the case $H = C_6$ (the 6-vertex cycle) fills a gap in the aforementioned result of Leveque et al. [18]. As we shall explain in Section 3, the NP-hardness construction relies on their gadget but also requires significant additional work. Before doing this we first prove the polynomial-time part of Theorem 2 in Section 2. Then, in Section 4, we prove Theorem 2.

In Section 5 we consider the problem from a parameterized complexity viewpoint. Recall that Golovach et al. [10] proved that INDUCED DISJOINT PATHS is FPT for claw-free graphs when parameterized by the number k of paths. We consider the class of P_r -free graphs. This gives us another natural parameter, namely r . However, we show by adapting a construction of Haas and Hoffmann [13] that even 2-INDUCED DISJOINT PATHS is W[1]-hard for P_r -free graphs when parameterized by r .

In Section 6 we summarize our findings and give a number of relevant open problems. In particular we discuss some open problems on the parameterized complexity of INDUCED DISJOINT PATHS.

2. Polynomial-time algorithms

In this section we prove the polynomial-time part of Theorem 2. We first show the following general result that we will need as a lemma.

Lemma 1. *For every linear forest F and every integer $k \geq 2$, if the k -INDUCED DISJOINT PATHS problem is polynomial-time solvable for H -free graphs for some graph H , then it is so for $(F + H)$ -free graphs.*

Proof. Let H be a graph such that k -INDUCED DISJOINT PATHS is polynomial-time solvable for H -free graphs. Let (G, T) be an instance of k -INDUCED DISJOINT PATHS, where G is an $(F + H)$ -free graph on n vertices and T is a set of k terminal pairs (s_i, t_i) .

Let $r = 2|V(F)| - 1$. Note that F is an induced subgraph of P_r . We check in $O(n^{k(r+1)})$ time (by brute force) if there exists a solution $(P^1, \dots, P^k)$ for (G, T) in which each path has at most $r + 1$ vertices. As k and r are constants, this takes polynomial time.

Suppose we have not found a solution yet. Then if a solution $(P^1, \dots, P^k)$ exists, at least one of the paths P^i in it has $r + 2$ or more vertices. We guess which path P^i will have at least $r + 2$ vertices. This leads to k branches. We guess the first $r + 1$ vertices $u_1, \dots, u_{r+1}$ on P^i after $s_i = u_0$. This leads to $O(n^r)$ further branches. We remove $s_i, u_1, \dots, u_r$ and all their neighbours from G , except for u_{r+1} . Let G' be the resulting graph. In the pair (s_i, t_i) , we replace s_i by u_{r+1} to obtain a new instance (G', T') . As F is an induced subgraph of P_r , we have that G' is H -free. Hence, by our assumption, we can solve k -INDUCED DISJOINT PATHS on (G', T') in polynomial time. As the total number of branches is polynomial, the total running time is polynomial. $\square$

We need two known results for proving the polynomial part of Theorem 2 in Lemma 2.

Theorem 3 ([8]). *For every $k \geq 2$, k -INDUCED DISJOINT PATHS is polynomial-time solvable for claw-free graphs.*

Theorem 4 ([2]). *If a connected chair-free graph G contains an induced claw and an induced path P on at least eight vertices, then G has a vertex adjacent to all vertices of P .*

Lemma 2. *Let $k \geq 2$. For every linear forest F , k -INDUCED DISJOINT PATHS is polynomial-time solvable for $(F + \text{chair})$ -free graphs.*

Proof. By Lemma 1, it remains to consider chair-free graphs. Let (G, T) be an instance of k -INDUCED DISJOINT PATHS, where G is a chair-free graph on n vertices and $T = \{(s_1, t_1), \dots, (s_k, t_k)\}$ is a set of terminal pairs. Let $(P^1, \dots, P^k)$ be a solution for (G, T) (if it exists). We call a path P^i *long* if it has at least eight vertices; else we call it *short*. We first guess which of the paths of a solution for (G, T) will be short. There are 2^k options for doing this, which is a constant number as k is a constant. We will consider each of these options one by one.

Suppose we consider the option where $T' \subseteq T$ is the subset of terminal pairs that will be in short solution paths. Let $|T'| = k' \leq k$. We guess all $O(n^{5k'}) = O(n^{5k})$ options of choosing the inner vertices of the solution paths for the terminal pairs in T' . We discard an option if two of the guessed solution paths contain an edge between them or if a guessed solution path contains a vertex with a neighbour in some $(s_i, t_i) \notin T'$. Otherwise, we continue as follows.

We first delete all vertices of the guessed solution paths and also their neighbours from G . We denote the new instance by (G, T) again and also write $T = \{(s_1, t_1), \dots, (s_k, t_k)\}$. Assuming our guess was correct, (G, T) only has solutions $(P^1, \dots, P^k)$ in which each P^i is long. Hence, from G , we can safely remove for every $i \in \{1, \dots, k\}$, every vertex that is adjacent to both s_i and t_i .

We now check in polynomial time if there are two terminal s_i and t_i that belong to different connected components of the resulting graph G' . If so, then we can discard this branch. Else, we let $(G'_1, T'_1), \dots, (G'_r, T'_r)$ be the connected components of G' , together with the terminal pairs subsets of T they contain.

We consider each (G'_j, T'_j) as a separate instance. We check in polynomial time if G'_j is claw-free. First suppose that we find that G'_j contains an induced claw. We claim that we may discard the branch, as (G'_j, T'_j) has no solution. For a




Fig. 1. The part of the graph G^1 that corresponds to a vertex x in a graph G that has exactly two neighbours x_1 and x_2 .

contradiction, suppose (G'_j, T'_j) has a solution. Recall that all paths in any solution for (G'_j, T'_j) must be long. Consider a path P^i in this solution. As P^i must be long, Theorem 4 tells us that G'_j must contain a vertex adjacent to all vertices of P^i . However, by construction, G'_j contains no vertices adjacent to both s_i and t_i , which both belong to P^i , a contradiction.

Now suppose that we find that G'_j is claw-free. Otherwise, we apply Theorem 3 to check in polynomial time if (G'_j, T'_j) has a solution. If for some (G'_j, T'_j) no solution exists, then we move to the next branch; otherwise, we return a yes-answer.

As the number of branches is polynomial and processing each branch takes polynomial time, the total running time of our algorithm is polynomial. $\square$

3. NP-completeness results

In this section we prove the NP-completeness part of Theorem 2 (see Section 4 for details on how we combine the several hardness results proven in this section). As explained below, we base our proofs on a hardness result of Leveque et al. [18] for 2-INDUCED CYCLE, which is to decide if a graph has an induced cycle containing two *distinguished* vertices x and y (we assume without loss of generality that the induced cycle is a *hole*, meaning it has at least four vertices).

Theorem 5 ([18]). *The 2-INDUCED CYCLE problem is NP-complete even for instances (G, x, y) where G is a graph of maximum degree 3, and x and y each have degree 2.*

Now let (G, x, y) be an instance of 2-INDUCED CYCLE, where G has maximum degree 3, and x and y have degree 2. Let x and y have neighbours x_1, x_2 and y_1, y_2 respectively. We replace x and its incident edges by the following. Create vertices $p_1, q_1, r_1, p_2, q_2, r_2, s_1, s_2$. Add edges $s_1p_1, p_1q_1, q_1x_1, q_1r_1, r_1s_2$ and $s_2r_2, r_2q_2, q_2x_2, p_2s_1, p_2q_2$. In the same way, we replace y and its incident edges by vertices $a_1, b_1, c_1, a_2, b_2, c_2, t_1, t_2$ and edges $a_1t_1, a_1b_1, b_1y_1, b_1c_1, c_1t_2$ and $a_2t_1, a_2b_2, b_2y_2, b_2c_2, c_2t_2$. Note that any vertex that is introduced has degree at most 3. For an integer $d \geq 0$, we now subdivide every incident edge of every newly introduced vertex d times. We denote the new graph by G^d and say that G^d is the d -replacement for G . Note that G^d has maximum degree 3 as well. See also Fig. 1, where we display the replacement of x for $d = 1$.

We now derive the following relation.

Lemma 3. *For a graph G of maximum degree 3 that has two vertices x and y of degree 2 and for every integer d , it holds that (G, x, y) is a yes-instance of 2-INDUCED CYCLE if and only if $(G^d, \{(s_1, t_1), (s_2, t_2)\})$ is a yes-instance of 2-INDUCED DISJOINT PATHS. Moreover, G^d has maximum degree 3.*

Proof. It is readily seen that for every integer $d \geq 0$, G^d has maximum degree 3.

To prove the remainder of the lemma, first assume that $d = 0$. We observe that the paths s_1, p_1, q_1, x_1 and s_2, r_2, q_2, x_2 are mutually induced. The paths s_1, p_2, q_2, x_2 and s_2, r_1, q_1, x_1 are also mutually induced. Moreover, these are the only two options that can co-exist, in the following sense. A path from s_1 to x_1 that uses only edges of this gadget and that does not have s_2 as an internal vertex has to pass through q_1 . Similarly, a path from s_2 to x_2 that uses only edges of this gadget and that does not have s_1 as an internal vertex has to pass through q_2 . See also Fig. 2. The same observations hold with respect to the gadget replacing vertex y . Hence, G has a hole containing x and y if and only if G^d has mutually induced paths between s_1 and t_1 and between s_2 and t_2 .




Fig. 2. The path from s_2 to x_1 (left) and the path from s_1 to x_2 (right) from the proof of Lemma 3 if $d = 1$.




Fig. 3. A drawing of H_1 (left) and H_3 (right).

Note that any incident edges of every vertex of G^d that does not belong to G can be subdivided an arbitrary number of times without affecting the correctness of the reduction (q_1, q_2, b_1, b_2 remain bottlenecks). So, the claim also holds for $d \geq 1$. $\square$

Our first two results require a single change to the construction of [18]. The first rectifies, by using Lemma 3, a potential issue with the same claim made in [10].

Lemma 4. For every $k \geq 2$, the k -INDUCED DISJOINT PATHS problem is NP-complete for $K_{1,4}$ -free graphs.

Proof. First let $k = 2$. By Theorem 5 it holds that 2-INDUCED CYCLE is NP-complete on graphs of maximum degree 3 where the two distinguished vertices have degree 2. Apply the reduction of Lemma 3 for $d = 0$. The graph G^0 has maximum degree 3; thus, it is $K_{1,4}$ -free.

Now suppose $k \geq 3$. We reduce from 2-INDUCED DISJOINT PATHS restricted to $K_{1,4}$ -free graphs. Let G be a $K_{1,4}$ -free graph that forms, together with two disjoint terminal pairs (s_1, t_1) and (s_2, t_2) , an instance of 2-INDUCED DISJOINT PATHS. We add $k - 2$ isolated edges $s_3 t_3, \dots, s_k t_k$ to G (so these isolated edges consist of new vertices $s_3, \dots, s_k$ and $t_3, \dots, t_k$). Let G' be the new graph. It is readily seen that G together with terminal pairs $(s_1, t_1), (s_2, t_2)$ has a solution if and only if G' together with terminal pairs $(s_1, t_1), \dots, (s_k, t_k)$ has a solution. $\square$

Lemma 5. For every $s \geq 3$ with $s \neq 6$, 2-INDUCED DISJOINT PATHS is NP-complete for C_s -free graphs.

Proof. Leveque et al. [18] proved that 2-INDUCED CYCLE is NP-complete on C_s -free graphs for every $s \geq 3$ with $s \neq 6$, where the two distinguished vertices have degree 2. We apply Lemma 3 for a sufficiently large integer d . In this way, we ensure that G^d is C_s -free. $\square$

Our third result requires a significant overhaul of the construction in [18].

3.1. Omitting “H”-graphs and Six-Vortex Cycles. Let H_1 be the “H”-graph on six vertices formed by an edge joining the middle vertices of two paths on three vertices. For $\ell \geq 2$, let H_ℓ be the graph obtained from H_1 by subdividing the crossing edge (which is the edge whose endpoints both have degree 3) $\ell - 1$ times. See Fig. 3 for two examples.

We prove that for every $\ell \geq 1$, 2-INDUCED DISJOINT PATHS is NP-complete for (C_6, H_ℓ) -free graphs. To this end, we consider the hardness reduction by Leveque et al. [18] for 2-INDUCED CYCLE in more detail. We very closely follow their notation and the proof of our main Lemma 6 mimics the proof of their Lemma 2.6. We show how their construction can be modified so that it becomes H_ℓ -free for any fixed $\ell \geq 1$ and C_6 -free.

Let ϕ be an instance of 3-SATISFIABILITY consisting of m clauses $C_1, \dots, C_m$ on n variables $z_1, \dots, z_n$. For each clause C_j of the form $y_{3j-2} \vee y_{3j-1} \vee y_{3j}$ where $y_i, i \in [3m]$, is a literal from $\{z_1, \dots, z_n, \bar{z}_1, \dots, \bar{z}_n\}$. Let $\ell \geq 1$ be given. We will construct a graph G_ϕ^ℓ with two specified vertices x and y of degree 2 so that G_ϕ^ℓ has a hole containing x and y if and only if there is a truth assignment satisfying ϕ .




Fig. 4. The literal gadget (dashed lines indicate paths of length ℓ).




Fig. 5. The clause gadget together with its interface with the literal gadget (drawn above). Dashed lines indicate paths of length ℓ .

For each literal y_j , prepare a graph $G^\ell(y_j)$ (literal gadget) as drawn in Fig. 4 where the corresponding labelled vertices inherit a subscript j . Numerous vertices on paths will remain unlabelled. Our literal gadget is more elaborate than that in [18] as we need to forbid, as induced subgraphs, the C_6 and for any fixed ℓ , every H_ℓ . The idea is that two induced disjoint paths may be drawn through this gadget either crossing the edges $(\alpha^{2+}, \alpha^{3+})$ and (β^{2+}, β^{3+}) ; or the edges $(\alpha^{2-}, \alpha^{3-})$ and (β^{2-}, β^{3-}) . All other possibilities are forbidden.

For each clause C_j , prepare a graph $G^\ell(C_j)$ (clause gadget) as drawn in the bottom of Fig. 5 where the corresponding labelled vertices inherit a subscript j . Numerous vertices on paths will remain unlabelled. Our clause gadget is exactly the same as in [18] except we replaced the edges by paths. The idea is that a path may be drawn through this gadget in precisely one of three ways selecting the literal that is true.

For each variable z_i , prepare a graph $G^\ell(z_i)$ (variable gadget) as in Fig. 6 consisting of two internally disjoint paths P_i^+ (top) and P_i^- (bottom). The idea in Fig. 6 is that full edges and dashed edges alternate on this diagram and the length is enough for m full edges. The end points of the full edges are labelled $(p_{i,1}^+, p_{i,1}^{++}), \dots, (p_{i,2m}^+, p_{i,2m}^{++})$ on the top; and $(p_{i,1}^-, p_{i,1}^{--}), \dots, (p_{i,2m}^-, p_{i,2m}^{--})$ on the bottom. Our variable gadget is exactly the same as in [18] except we lengthened some paths. The idea is that a path may be drawn through this gadget in precisely one of two ways selecting whether the variable is evaluated true or false.




Fig. 6. The variable gadget. Dashed lines indicate paths of length ℓ . Dotted lines indicate a continuation of the gadget.

We now construct the final graph G_ϕ^ℓ . We do this in a manner similar to Leveque et al. [18] from the disjoint union of all the graphs $G^\ell(y_j)$ (literal gadgets), $G^\ell(C_j)$ (clause gadgets) and $G^\ell(x_i)$ (variable gadgets). Compared to the construction [18], we replace some edges by paths of length ℓ (so the construction would be the same when $\ell = 1$) and sometimes use slightly different vertex names. In three places, however, the modifications go further beyond the construction of Lemma 2.6 in [18] and we will indicate those specifically below. In particular, the top of Fig. 5 shows how a clause gadget interacts with a literal gadget in our construction; see point 6 below. A variable gadget interacts with a clause gadget in a similar way; see points 4 and 5 below.

1. For $j = 1, \dots, 3m - 1$, add paths of length ℓ between α'_j and α_{j+1} and between β'_j and β_{j+1} .
2. For $j = 1, \dots, m - 1$, add paths of length ℓ between c_j^{0-} and c_{j+1}^{0+} .
3. For $i = 1, \dots, m - 1$, add paths of length ℓ between d_i^- and d_{i+1}^+ .
4. For $i = 1, \dots, n$, let $y_{n_1}, \dots, y_{n_{z_i^-}}$ be the occurrences of $\bar{z}_i$ over all literals. For $j = 1, \dots, z_i^-$, delete the edge $p_{i,j}^+ p_{i,j}^{++}$ and add the four edges $p_{i,j}^+ \alpha_{n_j}^{2+}$, $p_{i,j}^+ \beta_{n_j}^{2+}$, $p_{i,j}^{++} \alpha_{n_j}^{3+}$, $p_{i,j}^{++} \beta_{n_j}^{3+}$. Additionally to these edges, which were in [18], we also add: $p_{i,j}^+ \alpha_{n_j}^{3+}$, $p_{i,j}^+ \beta_{n_j}^{3+}$, $p_{i,j}^{++} \alpha_{n_j}^{2+}$, $p_{i,j}^{++} \beta_{n_j}^{2+}$.
5. For $i = 1, \dots, n$, let $y_{n_1}, \dots, y_{n_{z_i^+}}$ be the occurrences of z_i over all literals. For $j = 1, \dots, z_i^+$, delete the edge $p_{i,j}^- p_{i,j}^{--}$ and add the four edges $p_{i,j}^- \alpha_{n_j}^{2+}$, $p_{i,j}^- \beta_{n_j}^{2+}$, $p_{i,j}^{--} \alpha_{n_j}^{3+}$, $p_{i,j}^{--} \beta_{n_j}^{3+}$. Additionally to these edges, which were in [18], we also add: $p_{i,j}^- \alpha_{n_j}^{3+}$, $p_{i,j}^- \beta_{n_j}^{3+}$, $p_{i,j}^{--} \alpha_{n_j}^{2+}$, $p_{i,j}^{--} \beta_{n_j}^{2+}$.
6. For $i = 1, \dots, m$ and $j = 1, 2, 3$, add the edges $\alpha_{3(i-1)+j}^{2-} c_i^{j+}$, $\alpha_{3(i-1)+j}^{3-} c_i^{j-}$, $\beta_{3(i-1)+j}^{2-} c_i^{j+}$, $\beta_{3(i-1)+j}^{3-} c_i^{j-}$. Additionally to these edges, which were in [18], we also add: $\alpha_{3(i-1)+j}^{3-} c_i^{j+}$, $\alpha_{3(i-1)+j}^{2-} c_i^{j-}$, $\beta_{3(i-1)+j}^{3-} c_i^{j+}$, $\beta_{3(i-1)+j}^{2-} c_i^{j-}$. See the top of Fig. 5.
7. Add paths of length ℓ between α'_{3m} and d_1^+ and between β'_{3m} and c_1^{0+} .
8. Add the vertex x . Add paths of length ℓ between x and α_1 and between x and β_1 .
9. Add the vertex y . Add paths of length ℓ between y and c_m^{0-} and between y and d_n^- .

Claim 1. ϕ is satisfied by a truth assignment if and only if G_ϕ^ℓ contains a hole passing through x and y .

Proof. The idea is that any hole emanating from x and moving rightwards towards y (see Fig. 4) must traverse the literal gadgets in precisely one of two ways (upper path on top and bottom; or bottom path on the top and bottom). Now, subsequently the paths building the hole may return to these literal gadgets but they can never leave them as each $\alpha, \beta, \alpha', \beta'$ are already traversed. Subsequently, the paths do indeed not return to the literal gadgets to ensure their consistent evaluation with the variables and that one in each clause is true.

More formally, first assume that ϕ is satisfied by a truth assignment $\xi \in \{0, 1\}^n$. We pick a set of vertices that induce a hole containing x and y .

1. Pick vertices x and y .
2. For $i = 1, \dots, 3m$, pick $\alpha_i, \alpha'_i, \beta_i, \beta'_i$.
3. For $i = 1, \dots, 3m$, if y_i is satisfied by ξ , then pick $\alpha_i^{1+}, \alpha_i^{1++}, \alpha_i^{2+}, \alpha_i^{3+}, \alpha_i^{4++}, \alpha_i^{4+}$ and any vertices on a direct path between these. Else, pick $\alpha_i^{1-}, \alpha_i^{1--}, \alpha_i^{2-}, \alpha_i^{3-}, \alpha_i^{4--}, \alpha_i^{4-}$ and any vertices on a direct path between these.
4. For $i = 1, \dots, n$, if $\xi(i) = 1$, then pick all vertices of P_i^+ and all the neighbours of the vertices in P_i^+ of the form α_k^{2+} (or one could choose α_k^{3+} , but only one among the two) for any k . Additionally pick any vertices on a direct path between these.
5. For $i = 1, \dots, n$, if $\xi(i) = 0$, then pick all the vertices of P_i^- and all the neighbours of the vertices in P_i^- of the form α_k^{2+} (or one could choose α_k^{3+} , but only one among the two) for any k . Additionally pick any vertices on a direct path between these.

6. For $i = 1, \dots, m$, pick the vertices c_i^{0+} and c_i^{0-} . Choose any $j \in \{3i - 2, 3i - 1, 3i\}$ such that ξ satisfies y_j . Pick vertices α_j^{2-} and α_j^{3-} .
- If $j = 3i - 2$, then pick $c_i^{12+}, c_i^{1+}, c_i^{1-}, c_i^{12-}$ as well as all vertices on a path between: c^{0+} and c_i^{12+} ; c_i^{12+} and c_i^{1+} ; c^{0-} and c_i^{12-} ; c_i^{12-} and c_i^{1-} .
 - If $j = 3i - 1$, then pick $c_i^{12+}, c_i^{2+}, c_i^{2-}, c_i^{12-}$ as well as all vertices on a path between: c^{0+} and c_i^{12+} ; c_i^{12+} and c_i^{2+} ; c^{0-} and c_i^{12-} ; c_i^{12-} and c_i^{2-} .
 - If $j = 3i$, then pick c_i^{3+}, c_i^{3-} as well as all vertices on a path between: c^{0+} and c_i^{3+} ; c^{0-} and c_i^{3-} .

It suffices to show that the chosen vertices induce a hole in G_ϕ^ℓ containing x and y . The only potential problem is that for some k , one of the vertices $\alpha_k^{2+}, \alpha_k^{3+}, \alpha_k^{2-}, \alpha_k^{3-}$ was chosen more than once. If α_k^{2+} and α_k^{3+} were picked in Step 3, then y_k is satisfied by ξ . Therefore, α_k^{2+} and α_k^{3+} were not chosen in Step 4 or Step 5. Similarly, if α_k^{2-} and α_k^{3-} were picked in Step 6, then y_k is satisfied by ξ . Therefore, α_k^{2-} and α_k^{3-} were not chosen in Step 3. Thus, the chosen vertices induce a hole in G_ϕ^ℓ containing x and y .

Now assume that G_ϕ^ℓ has a hole including x and y . The hole must contain α_1 and β_1 since they are the only neighbours of x . Next, either both α_1^{1+} and β_1^{1+} are in the hole or both α_1^{1-} and β_1^{1-} . Without loss of generality, let α_1^{1+} and β_1^{1+} be in the hole (the same reasoning will apply in the other case). Since $\alpha_1^{1-}, \beta_1^{1-}, \alpha_1^{1--}, \beta_1^{1--}$ are all neighbours of two vertices in the hole, they cannot themselves be in the hole. Thus, $\alpha_1^{1++}, \beta_1^{1++}$ must be in the hole. Now, since, $\alpha_1^{1-}, \beta_1^{1-}, \alpha_1^{1--}, \beta_1^{1--}$ are all neighbours of two vertices in the hole, they cannot themselves be in the hole. Thus, $\alpha_1^{2+}, \beta_1^{2+}$, and the paths that lead to them, must be in the hole. Since $\alpha_1^{2+}, \beta_1^{2+}$ have the same neighbourhood outside of $G(y_1)$ it follows that $\alpha_1^{3+}, \beta_1^{3+}$ must be in the hole. Indeed, applying this reasoning once more, so must also $\alpha_1^{4++}, \beta_1^{4++}$, and the paths that lead to them. Now, since $\alpha_1^{4-}, \beta_1^{4-}, \alpha_1^{4--}, \beta_1^{4--}$ are all neighbours of both α_1^{4++} and β_1^{4++} , α_1^{4+} and β_1^{4+} must be in the hole. Finally, note that $\alpha_1^{4-}, \beta_1^{4-}$ are not in the hole, as they are adjacent to both α_1^{4++} and β_1^{4++} . So it must contain instead $\alpha_1', \beta_1', \alpha_2, \beta_2$. By induction, we see for $i \in [3m]$ that the hole must contain $\alpha_i, \beta_i, \alpha_i', \beta_i'$. Also, for each i , the hole must contain $\alpha_i^{1+}, \alpha_i^{1++}, \dots, \alpha_i^{2+}, \alpha_i^{3+}, \dots, \alpha_i^{4++}, \alpha_i^{4+}$ or $\alpha_i^{1-}, \alpha_i^{1--}, \dots, \alpha_i^{2-}, \alpha_i^{3-}, \dots, \alpha_i^{4--}, \alpha_i^{4-}$. Hence, the hole contains d_1^+ and c_1^{0+} .

By symmetry we may assume the hole contains d_1^{++} , and the path to $p_{1,1}^+$, and α_k^{2+} for some k . As α_k^{1++} is adjacent to two vertices in the hole, the hole must contain one of α_k^{2+} and α_k^{3+} . Similarly, the hole cannot proceed on a path to α_k^{4++} , so it must contain $p_{1,2}^+$ and $p_{1,2}^{++}$. By induction, we see that the hole contains $p_{1,i}^+, p_{1,i}^{++}$, for $i \in [n]$, and d_1^- . If the hole contains d_1^{--} , then the hole must contain $p_{1,i}^-, p_{1,i}^{--}$, for $i \in [n]$, and eventually d_1^{+-} , a contradiction. Thus, the hole must contain d_2^+ . By induction, for $i \in [n]$, we see that the hole contains all the vertices of the path P_i^+ or P_i^- and, by symmetry, we assume that the hole contains neighbours of the vertices in P_i^+ or P_i^- , one among α_k^{2+} and α_k^{3+} , for each k .

Similarly, for $i \in [m]$, it follows that the hole must contain c_i^{0+} and c_i^{0-} . The hole also contains one of the following:

- $c_i^{12+}, c_i^{1+}, c_i^{1-}, c_i^{12-}$, and the paths between, and either one of $\alpha_j^{2-}, \alpha_j^{3-}$; or one of $\beta_j^{2-}, \beta_j^{3-}$.
- $c_i^{12+}, c_i^{2+}, c_i^{2-}, c_i^{12-}$, and the paths between, and either one of $\alpha_j^{2-}, \alpha_j^{3-}$; or one of $\beta_j^{2-}, \beta_j^{3-}$.
- c_i^{3+} , and the path between, and either one of $\alpha_j^{2-}, \alpha_j^{3-}$; or one of $\beta_j^{2-}, \beta_j^{3-}$.

We now recover the satisfying assignment ξ . For $i \in [n]$, set $\xi(i) = 1$ if the vertices of P_i^+ are in the hole; otherwise set $\xi(i) = 0$. By construction, at least one literal in every clause is satisfied by ξ , so indeed ξ is a satisfying assignment. $\square$

Claim 2. The graph G_ϕ^ℓ is C_6 -free and H_i -free for every $i \in [\ell]$.

Proof. Owing to the length of the ℓ paths that populate our construction and are drawn as dashed edges in our figures, we need only verify the omission of the relevant graphs on the connected components of the graph G_ϕ after the removal of these ℓ paths that are dashed edges. That would suffice for C_6 , but H_i has a pendant edge, so for these we must leave a pendant edge from the corresponding connected component at the extremities of an instance of these ℓ paths that are drawn as dashed edges. In this fashion, we only need to check for omission of the given graphs in the non-trivial cases drawn in Fig. 7. It can be readily observed that these graphs are P_7 free (but they are not P_6 -free). Hence, we need not test beyond H_3 . This task was accomplished by a program testing subgraph isomorphism whose code we provide a link to.¹

We will give an explicit argument for the case of C_6 -freeness, which is simpler as C_6 has numerous symmetries (a transitive automorphism group). Let us begin with the graph depicted on the left-hand side of Fig. 7. This graph has an automorphism that swaps α and β at the same time as $+$ and $-$. It also has an automorphism that only swaps $+$ and

¹ See <https://github.com/barnabymartin/InducedSubgraph>.




Fig. 7. Cases that need to be checked for omission of the graphs C_6 and H_i ($1 \leq i \leq \ell$).

–. Any subgraph that induces a C_6 cannot contain any of the unlabelled vertices, nor α nor β . This leaves eight vertices that may be involved. We will consider the case where the C_6 contains α^{1+} . Owing to the two automorphisms we have described, this argument would equally apply to α^{1-} and β^{1-} . But any C_6 must involve one of these vertices as there were only eight to choose from. Thus, when we have considered this case, our work is done:

Subcase A. The C_6 contains α^{1+} and α^{1++} . All other neighbours of α^{1+} (except α) are adjacent to α^{1++} . No C_6 can be formed here.

Subcase B. The C_6 contains α^{1+} and α^{1--} . Any C_6 involving a path α^{1--} to α^{1+} must next go to β^{1-} . We cannot continue this cycle.

Subcase C. The C_6 contains α^{1+} and β^{1--} . Any C_6 involving a path β^{1--} to α^{1+} must next go to α^{1-} or α^{1--} . We cannot continue this cycle.

Subcase D. The C_6 contains α^{1+} and β^{1-} . Now, β^{1-} can have as the next in the cycle either of β^{1+} or β^{1++} . We cannot continue this cycle.

Subcase E. The C_6 contains α^{1+} and α^{1-} . Now, α^{1-} can have as the next in the cycle either of β^{1+} or β^{1++} . We cannot continue this cycle.

Now we consider the graph depicted on the right-hand side of Fig. 7. Any C_6 cannot contain any of the unlabelled vertices. It follows that it must use all six remaining vertices. But this induced graph has a triangle, so we are finished. $\square$

We note that the construction in [18] omits all cycles other than C_6 , and they note specifically this lacuna, which we have remedied.

We now prove our result.

Lemma 6. For every integer $\ell \geq 1$, 2-INDUCED DISJOINT PATHS is NP-complete for (C_6, H_ℓ) -free graphs.

Proof. We give a reduction from an instance ϕ of 3-SATISFIABILITY. First, we construct G_ϕ^ℓ . By Claim 1, G_ϕ^ℓ has a hole through x and y if and only if ϕ is satisfiable. Moreover, G_ϕ^ℓ is (C_6, H_ℓ) -free by Claim 2. We now apply the reduction of Lemma 3. As we can choose the integer d in Lemma 3 sufficiently large, the resulting graph is still (C_6, H_ℓ) -free. $\square$

4. The proof of Theorem 2

We now use the results from the previous two sections to prove our main theorem.

Theorem 2 (restated). Let $k \geq 2$. For a graph H , k -INDUCED DISJOINT PATHS restricted to H -free graphs is polynomial-time solvable if H is a subgraph of the disjoint union of a linear forest and a chair, and it is NP-complete if H is not in $\mathcal{S}$.

Proof. If H has a cycle C_s , apply Lemma 5 for $s \neq 6$ or Lemma 6 for $s = 6$. Then we may assume H is a forest. If H has a vertex of degree at least 4, then every $K_{1,4}$ -free graph is H -free, so apply Lemma 4. Suppose H has maximum degree at most 3. If H has a connected component with at least two vertices of degree 3, then H has an induced H_ℓ , so apply Lemma 6 again. Else, H is in $\mathcal{S}$. If H is a subgraph of the disjoint union of a linear forest and a chair, then H is a linear forest, or the disjoint union of a linear forest and a chair, or the disjoint union of a linear forest and a claw. Hence, we can apply Lemma 2. $\square$

5. Parameterized complexity

We prove that 2-INDUCED DISJOINT PATHS is W[1]-hard for P_r -free graphs when parameterized by r . In order to do this we adapt a reduction of Haas and Hoffmann [13]. They prove that finding an induced path through three specified vertices is W[1]-hard when parameterized by the length of the path. The graph of their construction can potentially contain arbitrarily long induced paths, but we propose a modification which guarantees that the length of any induced path in the construction is bounded. Below, we discuss the original construction and our modification; the proof of correctness is almost exactly the same.

Theorem 6. 2-INDUCED DISJOINT PATHS is W[1]-hard on P_r -free graphs, parameterized by r .

Proof. The reduction is from INDEPENDENT SET, which is known to be W[1]-hard when parameterized by the size k of the solution [6]. Let $G = (V, E)$ be an instance of this problem, where $V = \{v_1, \dots, v_n\}$, and let k be the parameter. The main ingredient of the construction by Haas and Hoffmann [13] is a set of k vertex sets, which are called *choice diamonds*. The i -th choice diamond consists of n vertices $v_1^i, \dots, v_n^i$, and two vertices s^i and t^i which are adjacent to all $v_1^i, \dots, v_n^i$. Our modification is to make $v_1^i, \dots, v_n^i$ a clique, instead of an independent set as in the original construction of [13] (the reason for this modification is to ensure that the final gadget we construct is P_r -free for $r = 4k + 4$). Afterwards, we identify s^{i-1} and t^i for each $2 \leq i \leq k$ and call the resulting graph G_{VC} .

Now create two copies of G_{VC} . Denote the vertices in the first copy by s^i, v_j^i , and t^i and in the second copy by σ^i, φ_j^i , and τ^i , respectively. Add an edge between t^k and τ^k and subdivide it once (this is a minor modification with respect to the original construction, in which t^k and τ^k are identified; however, we want t^k and τ^k to be terminals from different pairs). Call the resulting graph G'' .

Note that G'' is a union of $2k + 2$ cliques, one of which has size 1. In the remainder, we will not add any more vertices, only edges. Since any induced path can contain at most two vertices of any clique, the graph is and will remain P_r -free for $r = 4k + 4$.

From G'' , construct the graph G' by adding the following edges (again, following Haas and Hoffmann [13]):

- add a *consistency edge* between v_j^i and φ_ℓ^i , for all $1 \leq i \leq k$ and all $1 \leq j, \ell \leq n$ with $j \neq \ell$ (thus $kn(n - 1)$ consistency edges are added in total);
- add *independence edges* between $\{v_p^i, \varphi_p^i\}$ and $\{v_q^j, \varphi_q^j\}$ for each edge $v_p v_q \in E$ and for all $1 \leq i, j \leq k$ with $i \neq j$ (thus $4k(k - 1) \cdot |E|$ independence edges are added in total);
- add *set edges* between $\{v_\ell^i, \varphi_\ell^i\}$ and $\{v_\ell^j, \varphi_\ell^j\}$ for all $1 \leq i, j \leq k$ with $i \neq j$ and all $1 \leq \ell \leq n$ (thus $4k(k - 1)n$ set edges are added in total).

This completes the construction. Let the vertex pairs for the instance be (s^1, t^k) and (σ^1, τ^k) .

To show correctness, we essentially repeat the arguments of Haas and Hoffmann [13, Lemma 7, Theorem 8]. Let P^1, P^2 be mutually induced disjoint (s^1, t^k) - and (σ^1, τ^k) -paths respectively. By shortcutting if necessary, we may assume that P^1 and P^2 are induced paths of G' . By construction, P^1 must contain one of $\{v_1^1, \dots, v_n^1\}$, say v_j^1 , and P^2 must contain one of $\{\varphi_1^1, \dots, \varphi_n^1\}$, say φ_ℓ^1 . By the consistency edges, $j = \ell$; otherwise, P^1 and P^2 would not be mutually induced. Also note that since $v_1^1, \dots, v_n^1$ and $\varphi_1^1, \dots, \varphi_n^1$ both induce a clique, P^1 cannot follow a consistency edge from v_j^1 and P^2 cannot follow a consistency edge from φ_j^1 , or P^1 and P^2 would not be mutually induced. Also note that the independence and set edges incident on v_j^1 and φ_j^1 lead to exactly the same vertices, and thus these vertices cannot be part of P^1 nor P^2 , or they would not be mutually induced. Hence, P^1 must continue to $t^1 = s^2$ and P^2 must continue to $\tau^1 = \sigma^2$. Since $\{v_1^1, \dots, v_n^1\}$ is a clique and P^1 is an induced path, P^1 must continue to one of $\{v_1^2, \dots, v_n^2\}$. Similarly, P^2 must continue to one of $\{\varphi_1^2, \dots, \varphi_n^2\}$. By repeating the same argument, we can argue that P^1 and P^2 only use edges of the choice diamonds, and none of the consistency, independence, and set edges.

In particular, the preceding implies that the paths P^1 and P^2 use vertices $s^1, v_{\gamma_1}^1, t^1, \dots, v_{\gamma_k}^k, t^k$ and $\sigma^1, \varphi_{\delta_1}^1, \tau^1, \dots, \varphi_{\delta_k}^k, \tau^k$ respectively. The consistency edges ensure that $\gamma_i = \delta_i$ for $1 \leq i \leq k$, as argued previously. The set edges ensure that $v_{\gamma_1}, \dots, v_{\gamma_k}$ is a set of size k ; indeed, if $\gamma_i = \gamma_j$ for some $1 \leq i, j \leq k$ with $i \neq j$, then $v_{\gamma_i}^i \varphi_{\gamma_j}^j = v_{\gamma_i}^i \varphi_{\gamma_i}^j$ is an edge of G' by construction, but $v_{\gamma_i}^i$ is in P^1 and $\varphi_{\gamma_j}^j$ is in P^2 , contradicting that P^1 and P^2 are mutually induced. The independence edges ensure that $v_{\gamma_1}, \dots, v_{\gamma_k}$ form an independent set; indeed, if $v_{\gamma_i} v_{\gamma_j} \in E$ for some $1 \leq i, j \leq k$, then $v_{\gamma_i}^i \varphi_{\gamma_j}^j$ is an edge of G' by construction, but $v_{\gamma_i}^i$ is in P^1 and $\varphi_{\gamma_j}^j$ is in P^2 , contradicting that P^1 and P^2 are mutually induced. Hence, $v_{\gamma_1}, \dots, v_{\gamma_k}$ is an independent set of size k .

For the converse, an independent set $I = \{v_{\gamma_1}, \dots, v_{\gamma_k}\}$ in G can be transformed in a solution to the instance: by following a similar reasoning as above, the vertices $s^1, v_{\gamma_1}^1, t^1, \dots, v_{\gamma_k}^k, t^k$ form an (s^1, t^k) -path that is mutually disjoint from the (σ^1, τ^k) -path formed by $\sigma^1, \varphi_{\gamma_1}^1, \tau^1, \dots, \varphi_{\gamma_k}^k, \tau^k$. $\square$

6. Conclusions

We showed new tractable and hard results for k -INDUCED DISJOINT PATHS for H -free graphs and extended a number of known results in this way. The open cases all involve graphs H that are not the disjoint union of some linear forest and the chair but that do belong to the family $\mathcal{S}$; we recall that $\mathcal{S}$ consists of all graphs, every connected component of which is a path P_r or a subdivided claw $S_{h,i,j}$.

Due to the above, k -INDUCED DISJOINT PATHS belongs to a set of several other problems whose complexity is open for $S_{h,i,j}$ -free graphs for many h, i, j . The best-known problem of this set of problems is INDEPENDENT SET, which is to decide if a graph has an independent set of size at least p for some integer p . Alekseev [1] proved that if a graph H is not in $\mathcal{S}$, then INDEPENDENT SET is NP-complete for H -free graphs. If $H \in \mathcal{S}$, only a restricted number of cases are known to be polynomial-time solvable for INDEPENDENT SET. Another example is 3-COLOURING for H -free graphs of bounded diameter (see [20]). We do not know any graph $S_{h,i,j}$ and integer d such that 3-COLOURING for $S_{h,i,j}$ -free graphs of diameter d is NP-complete (and only a small number of polynomial cases exist). Hence, in order to make further progress on k -INDUCED DISJOINT PATHS and these other problems we must better understand the structure of $S_{h,i,j}$ -free graphs.

It would also be interesting to consider the parameterized complexity of the problem in more detail. First recall that Golovach et al. [10] showed that INDUCED DISJOINT PATHS is FPT for claw-free graphs when parameterized by k . The class of claw-free graphs is properly contained in the class of chair-free graphs. From Theorem 2 we know that INDUCED DISJOINT PATHS is XP for chair-free graphs when parameterized by k . Is the problem even FPT for chair-free graphs when parameterized by k ?

Moreover, in Theorem 6 we proved that even 2-INDUCED DISJOINT PATHS is W[1]-hard on P_r -free graphs when parameterized by r . But what is the parameterized complexity of INDUCED DISJOINT PATHS for P_r -free graphs if r is a constant but k is the parameter? So far, we only know that the problem is in XP by Theorem 2 and that we may assume that $r \geq 7$ due to Theorem 1. Again we may draw a parallel to the situation for INDEPENDENT SET on P_r -free graphs. This problem can be solved in polynomial time for $r < 7$ [9] and is trivially in XP when parameterized by the size of the solution. However, it is open whether it is FPT for constant $r \geq 7$ when parameterized by the size of the solution (see also [5]). Progress on this problem would help to advance our understanding of k -INDUCED DISJOINT PATHS on P_r -free graphs (see [22] for a close relationship between both problems when the input is restricted to P_r -free graphs for some integer $r \geq 1$).

Finally, in some previous works, a slightly more general definition is used (see also Section 1). Given a graph G , *vertex-disjoint paths* $P^1, \dots, P^k$, for some integer $k \geq 1$ are *flexibly mutually induced paths* of G if there is no edge between two vertices from different P^i and P^j except possibly between the endpoints of the paths. If k is in the input, the complexity of the corresponding decision problem and ours is most likely different for P_r -free graphs. Namely, FLEXIBLY INDUCED DISJOINT PATHS is NP-complete for P_{14} -free graphs [22], whilst INDUCED DISJOINT PATH is quasipolynomial-time solvable for P_{14} -free graphs by Theorem 1. However, it is readily seen that all polynomial-time results in Theorem 2 (so, for fixed $k \geq 2$) also hold for FLEXIBLE k -INDUCED DISJOINT PATHS. This is even in the case if we also allow that two different paths P^i and P^j share a terminal (this even more general variant has been considered in the literature as well).

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

No data was used for the research described in the article.

References

- [1] V.E. Alekseev, The effect of local constraints on the complexity of determination of the graph independence number, in: *Combinatorial-Algebraic Methods in Applied Mathematics*, 1982, pp. 3–13 (in Russian).
- [2] V.E. Alekseev, Polynomial algorithm for finding the largest independent sets in graphs without forks, *Discrete Appl. Math.* 135 (2004) 3–16.
- [3] R. Belmonte, P.A. Golovach, P. Heggernes, P. van't Hof, M. Kaminski, D. Paulusma, Detecting fixed patterns in chordal graphs in polynomial time, *Algorithmica* 69 (3) (2014) 501–521.
- [4] D. Bienstock, On the complexity of testing for odd holes and induced odd paths, *Discrete Math.* 90 (1991) 85–92.
- [5] E. Bonnet, N. Bousquet, P. Charbit, S. Thomassé, R. Watrigant, Parameterized complexity of independent set in H -free graphs, *Algorithmica* 82 (2020) 2360–2394.
- [6] R.G. Downey, M.R. Fellows, Fixed-parameter tractability and completeness II: on completeness for W[1], *Theor. Comput. Sci.* 141 (1–2) (1995) 109–131.
- [7] M.R. Fellows, The Robertson-Seymour theorems: a survey of applications, in: *Proc. AMS-IMS-SIAM Joint Summer Research Conference*, in: *Contemporary Mathematics*, vol. 89, 1989, pp. 1–18.
- [8] J. Fiala, M. Kamiński, B. Lidický, D. Paulusma, The k -in-a-path problem for claw-free graphs, *Algorithmica* 62 (2012) 499–519.
- [9] A. Grzesik, T. Klimošová, M. Pilipczuk, M. Pilipczuk, Polynomial-time algorithm for maximum weight independent set on P_6 -free graphs, *ACM Trans. Algorithms* 18 (1) (2022) 4.
- [10] P.A. Golovach, D. Paulusma, E.J. van Leeuwen, Induced disjoint paths in claw-free graphs, *SIAM J. Discrete Math.* 29 (2015) 348–375.
- [11] P.A. Golovach, D. Paulusma, E.J. van Leeuwen, Induced disjoint paths in circular-arc graphs in linear time, *Theor. Comput. Sci.* 640 (2016) 70–83.
- [12] P.A. Golovach, D. Paulusma, E.J. van Leeuwen, Induced disjoint paths in AT-free graphs, *J. Comput. Syst. Sci.* 124 (2022) 170–191.

- [13] R. Haas, M. Hoffmann, Chordless paths through three vertices, *Theor. Comput. Sci.* 351 (2006) 360–371.
- [14] L. Jaffke, O. Kwon, J.A. Telle, Mim-width I. Induced path problems, *Discrete Appl. Math.* 278 (2020) 153–168.
- [15] K. Kawarabayashi, Y. Kobayashi, A linear time algorithm for the induced disjoint paths problem in planar graphs, *J. Comput. Syst. Sci.* 78 (2012) 670–680.
- [16] W. Kern, B. Martin, D. Paulusma, S. Smith, E.J. van Leeuwen, Disjoint paths and connected subgraphs for H -free graphs, *Theor. Comput. Sci.* 898 (2022) 59–68.
- [17] Y. Kobayashi, K. Kawarabayashi, Algorithms for finding an induced cycle in planar graphs and bounded genus graphs, in: *Proc. SODA 2009*, 2009, pp. 1146–1155.
- [18] B. Lévêque, D.Y. Lin, F. Maffray, N. Trotignon, Detecting induced subgraphs, *Discrete Appl. Math.* 157 (2009) 3540–3551.
- [19] J. Lynch, The equivalence of theorem proving and the interconnection problem, *SIGDA Newsl.* 5 (1975) 31–36.
- [20] B. Martin, D. Paulusma, S. Smith, Colouring H -free graphs of bounded diameter, in: *Proc. MFCS 2019*, in: *LIPIcs*, vol. 138, 2019, 14.
- [21] B. Martin, D. Paulusma, S. Smith, E.J. van Leeuwen, Few induced disjoint paths for H -free graphs, in: *Proc. ISCO 2022*, in: *LNCS*, in press.
- [22] B. Martin, D. Paulusma, S. Smith, E.J. van Leeuwen, Induced disjoint paths and connected subgraphs for H -free graphs, in: *Proc. WG 2022*, in: *LNCS*, vol. 13453, 2022, pp. 398–411.
- [23] M. Radovanović, N. Trotignon, K. Vušković, The (theta,wheel)-free graphs part IV: induced paths and cycles, *J. Comb. Theory, Ser. B* 146 (2021) 495–531.
- [24] N. Robertson, P.D. Seymour, Graph minors XIII. The disjoint paths problem, *J. Comb. Theory, Ser. B* 63 (1995) 65–110.
- [25] N. Shibi, Algorithme de recherche d'un stable de cardinalité maximum dans un graphe sans étoile, *Discrete Math.* 29 (1980) 53–76.