



Induced Disjoint Paths in AT-free graphs[☆]

Petr A. Golovach^a, Daniël Paulusma^{b,*}, Erik Jan van Leeuwen^c

^a Department of Informatics, University of Bergen, PB 7803, 5020 Bergen, Norway

^b School of Engineering and Computing Sciences, Durham University, Science Laboratories, South Road, Durham DH1 3LE, UK

^c Department of Information and Computing Sciences, Utrecht University, Princetonplein 5, 3584 CC Utrecht, The Netherlands

ARTICLE INFO

Article history:

Received 14 February 2021

Received in revised form 15 September 2021

Accepted 18 October 2021

Available online 29 October 2021

Keywords:

Disjoint paths

AT-free graph

Co-bipartite graph

Polynomial-time algorithm

W[1]-hardness

Induced topological minor

k -in-a-Path

k -in-a-Tree

ABSTRACT

Paths P_1, \dots, P_k are mutually induced if any two distinct P_i and P_j have neither common vertices nor adjacent vertices (except perhaps their end-vertices). The 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 connects s_i and t_i . This problem is NP-complete even for $k = 2$. We prove that it can be solved in polynomial time for AT-free graphs even when k is part of the input. Consequently, the problem of deciding if an AT-free graph contains a fixed graph H as an induced topological minor admits a polynomial-time algorithm. We show that such an algorithm is essentially optimal by proving that the problem is W[1]-hard with parameter $|V_H|$, even for a subclass of AT-free graphs, namely cobipartite graphs. We also show that the problems k -IN-A-PATH and k -IN-A-TREE are polynomial-time solvable on AT-free graphs even if k is part of the input.

© 2021 Elsevier Inc. All rights reserved.

1. Introduction

We study the INDUCED DISJOINT PATHS problem. This problem can be seen as the induced version of the well-known DISJOINT PATHS problem, which is to decide if a graph G with k pairs of specified vertices (s_i, t_i) contains a set of k mutually vertex-disjoint paths P_1, \dots, P_k , where each P_i starts from s_i and ends at t_i . The DISJOINT PATHS problem is one of the problems in Karp's list of NP-complete problems [21]. If k is a *fixed* integer, that is, not part of the input, then the problem can be solved in $O(n^3)$ time for n -vertex graphs, as shown by Robertson and Seymour [39] (see [23] for an algorithm with $O(n^2)$ running time). In contrast, the INDUCED DISJOINT PATHS problem is NP-complete even if $k = 2$, as shown both by Fellows [12] and Bienstock [2].

The INDUCED DISJOINT PATHS problem is formally defined as follows. The problem takes as input a graph G and a set $S = \{(s_1, t_1), \dots, (s_k, t_k)\}$ of k distinct unordered pairs of specified vertices of a graph G , which we call *terminals*. It asks if G contains k paths P_1, \dots, P_k that satisfy the following conditions:

1. for every i , P_i has end-vertices s_i and t_i ;
2. if $i \neq j$, then P_i and P_j may share at most one vertex, which must be an end-vertex of both P_i and P_j ;

[☆] This work is supported by EPSRC (EP/G043434/1), Royal Society (JP100692), the European Research Council (ERC) via the grant Rigorous Theory of Preprocessing (reference 267959), and ERC StG project PAAL no. 259515. A preliminary abstract of this paper appeared in the proceedings of SWAT 2012 [14].

* Corresponding author.

E-mail addresses: petr.golovach@ii.uib.no (P.A. Golovach), daniel.paulusma@durham.ac.uk (D. Paulusma), e.j.van.leeuwen@uu.nl (E.J. van Leeuwen).




Fig. 1. An example of an instance (G, S) of INDUCED DISJOINT PATHS, which is a yes-instance as the three paths P_1, P_2, P_3 satisfy conditions 1–3.

- 3. if $i \neq j$, then no inner vertex u of P_i is adjacent to a vertex v of P_j except when v is an end-vertex of both P_i and P_j ;
- 4. for every i , P_i is an induced path.

A collection of paths P_1, \dots, P_k that satisfies the aforementioned conditions 1–4 for a given graph G with a set S of terminals pairs (s_i, t_i) is called a *solution* for (G, S) . An example of a graph with three paths that satisfy all conditions is shown in Fig. 1.

We make two observations about the problem definition.

Remark 1. We require the terminal pairs to be distinct, but do allow in conditions 2 and 3 that two terminals from distinct pairs coincide or are adjacent. Hence, conditions 2 and 3 are slightly more relaxed compared to the usual definition of INDUCED DISJOINT PATHS by means of mutually induced paths where paths in a solution have neither common nor adjacent vertices. For general graphs, we can make these relaxations without loss of generality, because we can perform some elementary graph operations to ensure that two terminals from different pairs are distinct and non-adjacent. However, these operations may break membership of a graph class. Hence, our definition with the relaxed conditions yields a more general problem to solve if we restrict the input to some special graph class. We need this more general problem for our applications, as we will explain in Section 1.3.¹

Remark 2. As we can take shortcuts if necessary, we may omit condition 4 to obtain an equivalent problem. However, we added condition 4 explicitly, as this will be convenient for algorithmic purposes.

Research question. As the INDUCED DISJOINT PATHS problem is NP-complete even if $k = 2$ [2,12], we study the following research question to increase our understanding in the computational complexity of the problem:

What are natural graph classes for which INDUCED DISJOINT PATHS is polynomial-time solvable?

1.1. Known results

For planar graphs, INDUCED DISJOINT PATHS stays NP-complete. This result can be obtained by subdividing each edge of a planar input graph of DISJOINT PATHS. As planar graphs are closed under edge subdivision, this yields a planar input graph of INDUCED DISJOINT PATHS. We then apply the result of Lynch [35] who proved that DISJOINT PATHS is NP-complete for planar graphs. Kobayashi and Kawarabayashi [25] presented an algorithm that solves INDUCED DISJOINT PATHS on planar graphs that runs in linear time for any fixed k ; see [24] for an extension of this result to graph classes of bounded genus.

For claw-free graphs, INDUCED DISJOINT PATHS stays NP-complete as well. This is shown by Fiala et al. [13] even for line graphs, which form a subclass of claw-free graphs, and very recently by Radovanović, Trotignon and Vušković [38] even for line graphs of triangle-free chordless graphs. In addition, Fiala et al. [13] showed that INDUCED DISJOINT PATHS can be solved in polynomial time for claw-free graphs if k is fixed. In a previous paper [16], we improved the latter result by showing that INDUCED DISJOINT PATHS is fixed-parameter tractable for claw-free graphs when parameterized by k . Radovanović, Trotignon and Vušković [38] proved that INDUCED DISJOINT PATHS can be solved in polynomial time for (theta,wheel)-free graphs if k is fixed. However, the problem stays NP-complete for (theta,wheel)-free graphs if k is part of the input due to the aforementioned NP-completeness result for the subclass of line graphs of triangle-free chordless graphs [38].

From the above we observe that being planar, claw-free or (theta,wheel)-free does not help with respect to polynomial-time solvability of INDUCED DISJOINT PATHS when k is part of the input. However, for chordal graphs, INDUCED DISJOINT PATHS

¹ We refer to Open problem 1 in Section 6 for a discussion on the even more general problem variant where two terminal pairs are allowed to be the same.

is polynomial-time solvable when k is part of the input, as shown by Belmonte et al. [1], whereas we showed that for circular-arc graphs, INDUCED DISJOINT PATHS is even linear-time solvable [15].

More recently, Jaffke, Kwon and Telle [20] proved that INDUCED DISJOINT PATHS is polynomial-time solvable for any graph class of bounded mim-width. Their results imply that in addition to the above graph classes, the problem is polynomial-time solvable for the following classes: bi-interval graphs, circular permutation graphs, convex graphs, circular p -trapezoid graphs, p -polygon graphs, Dilworth- p graphs, and co- p -degenerate graphs.

1.2. Our main result

In this paper, we prove that INDUCED DISJOINT PATHS is polynomial-time solvable on a large class of graphs, namely the class of *asteroidal triple-free graphs* (or *AT-free graphs*), even if k is part of the input. An *asteroidal triple* is a set of three mutually non-adjacent vertices, such that each two of them are joined by a path that avoids the neighbourhood of the third; AT-free graphs are exactly those graphs that contain no such triple [31]. The class of AT-free graphs is easily recognized in polynomial time, widely studied (see e.g. [22,26–28,40]), and contains many well-known classes, such as cobipartite graphs, cocomparability graphs, cographs, interval graphs, permutation graphs, and p -trapezoid graphs (cf. [9]). In contrast to these subclasses, AT-free graphs miss the advantage of having a known geometric intersection model (which can be exploited for the design of polynomial-time algorithms). The class of AT-free graphs is incomparable to chordal graphs, circular-arc graphs, and graphs of bounded mim-width, which were covered by previous work, as mentioned above. Hence, with our main result, we make significant progress on answering the research question.

It is interesting to observe here that INDUCED DISJOINT PATHS seems to behave substantially different from DISJOINT PATHS. Generally, the induced variant of a containment relation problem is computationally just as hard or even harder than its non-induced variant. For example, MATCHING can be solved in polynomial time, but INDUCED MATCHING is NP-hard even for bipartite graphs [4]. In contrast, DISJOINT PATHS is NP-complete on split graphs [19], but INDUCED DISJOINT PATHS is polynomial-time solvable on the strictly larger class of chordal graphs [1]. Similarly, DISJOINT PATHS is NP-complete on interval graphs [37], but INDUCED DISJOINT PATHS is polynomial-time solvable on the strictly larger classes of circular-arc graphs [15] and chordal graphs [1]. With our main result, we now add the strictly and substantially larger class of AT-free graphs to the list of classes for which INDUCED DISJOINT PATHS is polynomial-time solvable, but DISJOINT PATHS is not [37], unless $P=NP$.

In order to solve INDUCED DISJOINT PATHS in polynomial time on AT-free graphs, we first apply some general preprocessing rules (Section 3.1). We then provide a thorough exploration of the structure of AT-free graphs, in particular in relation to the INDUCED DISJOINT PATHS problem (Section 3.2). Finally, we use these structural results for a dynamic programming algorithm (Section 3.3). Here we heavily rely on the fact that AT-free graphs have a path that dominates the graph and that AT-free graphs cannot contain long induced cycles, so that we can always focus on the neighbourhood of a constant number of vertices in the graph.

Our approach is substantially different from the approach of Belmonte et al. [1] for solving this problem in polynomial time on chordal graphs, as it appears that the tree-decomposition-based approach of [1] does not work for AT-free graphs. Our proof techniques are also quite different from the irrelevant vertex technique used by Kobayashi and Kawarabayashi [25] for planar graphs, which only work for fixed k in contrast to our result, and from the proof techniques for claw-free graphs [13,16]. The latter techniques are based on the characterization of claw-free graphs of Chudnovsky and Seymour [5] and also only work when k is a fixed constant instead of being part of the input.

1.3. Applications to other containment relation problems

We show that our algorithm for INDUCED DISJOINT PATHS for AT-free graphs can be used as a subroutine to solve two categories of containment relation problems in polynomial-time for AT-free graphs. Recall that a containment relation problem asks to detect if one graph is contained in some other graph by some specified graph containment relation. For general graphs, INDUCED DISJOINT PATHS has no direct applications for other containment problems, as it is NP-complete when $k=2$ [2,12]. However, such applications are possible for graph classes for which INDUCED DISJOINT PATHS is polynomial-time solvable, such as AT-free graphs.

Below we provide a definition, survey previous work, and state our results for both considered categories of containment relation problems. Proof details can be found in Sections 4 and 5.

H -Induced Topological Minor. A graph G contains a graph H as a *topological (induced) minor* if G contains an (induced) subgraph that is isomorphic to a subdivision of H , that is, to a graph obtained from H by a number of edge subdivisions. The problems that are to decide whether a given graph contains some fixed graph H as a topological minor or induced topological minor are called H -TOPOLOGICAL MINOR and H -INDUCED TOPOLOGICAL MINOR, respectively.

There is a great disparity in what we know about H -TOPOLOGICAL MINOR and H -INDUCED TOPOLOGICAL MINOR. Robertson and Seymour [39] showed that H -TOPOLOGICAL MINOR can be solved in polynomial time for any fixed graph H . Grohe et al. [18] improved this result to cubic time. In contrast, for H -INDUCED TOPOLOGICAL MINOR, Lévêque et al. [32] obtained some partial results by showing both polynomial-time and NP-complete cases. In particular they proved that H -INDUCED TOPOLOGICAL MINOR is NP-complete for $H = K_5$. Afterwards, Le [30] proved that H -INDUCED TOPOLOGICAL MINOR is polynomial-time

solvable for $H = K_4$, which was known to be a stubborn case (see also [33]). However, the complexity classification for H -INDUCED TOPOLOGICAL MINOR is still far from complete.

In contrast to the situation for general graphs, we can determine the complexity of H -INDUCED TOPOLOGICAL MINOR for any graph H when we restrict the input to AT-free graphs. For doing this, we make explicit use of conditions 2 and 3 for the terminals in the definition of INDUCED DISJOINT PATHS. We first consider the anchored version of H -INDUCED TOPOLOGICAL MINOR. This variant is to decide if a graph G contains a graph H as an induced topological minor such that the vertices of H are mapped to specified vertices of G . We show that the anchored version of H -INDUCED TOPOLOGICAL MINOR can be solved in polynomial time on AT-free graphs even when H is an arbitrary graph that is part of the input. This result cannot be generalized to the original, non-anchored version of this problem, which we show to be NP-complete even for a subclass of AT-free graphs, namely for the class of cobipartite graphs.

However, our result for the anchored version still implies a polynomial-time algorithm for H -INDUCED TOPOLOGICAL MINOR on AT-free graphs as long as H is fixed. This may be the best result one can hope for: besides proving NP-completeness, we also prove in Section 4 that INDUCED TOPOLOGICAL MINOR is W[1]-hard when parameterized by $|V_H|$ even for cobipartite graphs.

We note that by demanding conditions 2 and 3 and following the above approach based on the anchored version of the problem, H -INDUCED TOPOLOGICAL MINOR (for fixed H) is polynomial-time solvable for any graph class for which INDUCED DISJOINT PATHS is polynomial-time solvable or becomes polynomial-time solvable after fixing k ; see, for example, [1,13,20,38] where this has been shown for chordal graphs, claw-free graphs, graphs of bounded mim-width and (theta,wheel)-free graphs, respectively.

k -in-a-Subgraph. The problem of detecting an induced subgraph of a certain type containing a set of k specified vertices (called terminals as well) has also been extensively studied, in particular for the case where the induced subgraph is required to be a tree, cycle, or path. Then this problem is called k -IN-A-TREE, k -IN-A-CYCLE, and k -IN-A-PATH, respectively. Derhy and Picouleau [7] proved that k -IN-A-TREE is NP-complete when k is part of the input even for planar bipartite cubic graphs. Chudnovsky and Seymour [6] proved that 3-IN-A-TREE is polynomial-time solvable. Very recently, Lai, Lu and Thorup [29] gave a faster, near-linear algorithm for 3-IN-A-TREE and Gomes et al. [17] proved various parameterized complexity results for k -IN-A-TREE. However, the (classical) complexity of k -IN-A-TREE is still open for every fixed $k \geq 4$. In contrast, the problems 2-IN-A-CYCLE and 3-IN-A-PATH are NP-complete; this follows from the aforementioned results of Fellows [12] and Bienstock [2].

We note that for every fixed k , the problems k -IN-A-CYCLE, k -IN-A-PATH and k -IN-A-TREE can be reduced to a polynomial number of instances of INDUCED DISJOINT PATHS due to the relaxed conditions 2 and 3 in the problem definition (see, for example, [38] for details). Hence, all the aforementioned polynomial-time results for the latter problem restricted to special graph classes carry over to these three k -in-a-subgraph problems for every fixed integer $k \geq 1$. It is also known that k -IN-A-TREE can be solved in polynomial time for graphs of girth at least k [34] even if k is part of the input (see [8] for the case where $k = 3$). Lévêque et al. [32] proved that 2-IN-A-CYCLE is polynomial-time solvable for graphs not containing an induced path or induced subdivided claw on some fixed number of vertices. Radovanović, Trotignon and Vušković [38] proved that k -IN-A-CYCLE is fixed-parameter tractable for (theta,wheel)-free graphs when parameterized by k .

In our paper we add to these positive results by showing that the three problems k -IN-A-TREE, k -IN-A-CYCLE, and k -IN-A-PATH are polynomial-time solvable on AT-free graphs even when k is part of the input. Note that solving k -IN-A-CYCLE in polynomial time is straightforward, because AT-free graphs do not contain induced cycles on six or more vertices. For the remaining two problems, we again apply our main result and make explicit use of conditions 2 and 3 for the terminals in the definition of INDUCED DISJOINT PATHS. The standard brute-force approach that reduces the instance to a polynomial number of instances of INDUCED DISJOINT PATHS will only yield this result for any fixed integer k , because in the worst case it must process all $k!$ orderings of the terminals. However, we are able to give a more direct approach for AT-free graphs, based on dynamic programming, that works even when k is part of the input.

2. Preliminaries

Throughout the paper, we consider only finite, undirected graphs without multiple edges and self-loops. We denote the vertex set and edge set of a graph G by V_G and E_G , respectively; we may omit subscripts if no confusion is possible.

Let $G = (V, E)$ be a graph. For $U \subseteq V$, the graph $G[U]$ denotes the subgraph of G induced by the vertices in U . We denote the (open) neighbourhood of a vertex u by $N(u) = \{v \mid uv \in E\}$ and its closed neighbourhood by $N[u] = N(u) \cup \{u\}$. We denote the (open) neighbourhood of a set $U \subseteq V$ by $N(U) = \{v \in V \setminus U \mid uv \in E \text{ for some } u \in U\}$ and its closed neighbourhood by $N[U] = N(U) \cup U$. We let $d(u) = |N(u)|$ denote the *degree* of a vertex u . Whenever it is not clear from the context, we may add an extra subscript G to these notations.

The *length* of a path is its number of edges. 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 (except perhaps their end-vertices). Let $uw \in E$. The *edge subdivision* of uw removes uw and adds a new vertex v with edges uv and vw . A graph G' is a *subdivision* of G if G' can be obtained from G by a sequence of edge subdivisions. A graph H is an *induced topological minor* of G , if G has an induced subgraph that is isomorphic to a subdivision of H .

Let G be a graph in which we specify r distinct vertices p_1, \dots, p_r . Let H be a graph in which we specify r distinct vertices q_1, \dots, q_r . Then G contains H as an induced topological minor *anchored* in $(p_1, q_1), \dots, (p_r, q_r)$ if G contains an

induced subgraph isomorphic to a subdivision of H such that the isomorphism maps p_i to q_i for $i = 1, \dots, r$. The graphs G and H may have common vertices. If $p_i = q_i$ for $i = 1, \dots, r$, we speak of “being anchored in p_1, \dots, p_k ” instead.

A subset $U \subseteq V$ dominates G if $u \in U$ or $u \in N(U)$ for each $u \in V$. A pair of vertices $\{x, y\}$ is a dominating pair of a graph G if the vertex set of every (x, y) -path dominates G . A path P dominates a vertex u if $u \in N[V_P]$. A path P dominates a vertex set U if it dominates each $u \in U$. Corneil, Olariu and Stewart proved the following result, which we will use as a lemma.

Lemma 1 ([9,10]). *Every connected AT-free graph has a dominating pair and such a pair can be found in linear time.*

Throughout the paper, we let $T = \bigcup_{i=1}^k \{s_i, t_i\}$ denote the set of terminals, whereas $S = \{(s_1, t_1), \dots, (s_k, t_k)\}$ denotes the set of terminal pairs. Note that a vertex u may be in more than one terminal pair. We say that a terminal s_i or t_i is represented by a vertex u if $u = s_i$ or $u = t_i$, respectively. We also say that u is a terminal vertex. If a vertex u does not represent any terminal, then we call u a non-terminal vertex.

3. Induced disjoint paths

In this section we present our polynomial-time algorithm for the INDUCED DISJOINT PATHS problem. Our algorithm has as input a pair (G, S) , where G is an AT-free graph and S is a set of terminal pairs. It consists of the following three phases.

Phase 1. Preprocess (G, S) to derive a number of convenient properties. For instance, afterwards two terminals of the same pair are non-adjacent.

Phase 2. Derive a number of structural properties of (G, S) . The algorithm constructs an auxiliary graph H , which is obtained from the subgraph of G induced by the terminal vertices by adding a path of length 2 between each pair of terminals. It then checks whether H satisfies some necessary conditions, such as being AT-free and being an anchored topological minor of G . The latter condition is also shown to be sufficient and demands the construction of another auxiliary graph in order to describe how induced paths connecting terminal pairs may interfere with each other.

Phase 3. Perform dynamic programming using the information from Phases 1–2.

3.1. Phase 1: preprocessing

Let (G, S) be an instance of INDUCED DISJOINT PATHS, where $G = (V, E)$ is an AT-free graph on n vertices with a set S of k terminal pairs $(s_1, t_1), \dots, (s_k, t_k)$. Recall that we assume that the terminal pairs in S are pairwise distinct if we consider them to be unordered. In other words, although a vertex may represent more than one terminal, we do not allow two terminal pairs to coincide, that is, $\{s_i, t_i\} \neq \{s_j, t_j\}$ for all $i \neq j$. We can safely make this assumption, as the algorithm for detecting induced topological minors that uses our algorithm for INDUCED DISJOINT PATHS as a subroutine does not require two terminal pairs to coincide, as we shall see. It may happen though that in this application two terminals of the same pair are the same. However, we will show that we can easily work around this, and to avoid any further technicalities in our algorithm for INDUCED DISJOINT PATHS we assume from now on that this is not the case either. Both assumptions are summarized in the following condition.

Condition i) For all $i \neq j$, $\{s_i, t_i\} \neq \{s_j, t_j\}$, and for all i , $s_i \neq t_i$.

After imposing Condition i) on S , we continue with the following preprocessing steps. We first apply Step 1, then Step 2, and then Step 3, where we perform each step as long as possible. In the below, $M(u)$ denotes the set of all terminals that form a terminal pair with vertex u . Recall also that $T = \bigcup_{i=1}^k \{s_i, t_i\}$ denotes the set of terminals.

Step 1. For all $u, v \in T$ with $uv \in E$, remove all common neighbours of u and v that are not terminals from G , that is, remove all vertices of $N(u) \cap N(v) \cap (V \setminus T)$.

Step 2. Let U denote the set of all terminal vertices u such that u only represents terminals whose partners are in $N(u)$; that is, $U = \{u \in T \mid M(u) \subseteq N(u)\}$. Remove U and all non-terminal vertices of $N(U)$ from G . Remove also all terminal pairs (s_j, t_j) from S for which $s_j \in U$ or $t_j \in U$.

Step 3. For all (s_i, t_i) with $s_i t_i \in E$, remove (s_i, t_i) from S .

Observe that Steps 1–3 preserve AT-freeness. Steps 1–3 are identical to Rules 1–3 in [16] for claw-free graphs. As such, the correctness of these three steps, as stated in our next lemma, can be proved by using exactly the same arguments as in [16].

Lemma 2. *Applying Steps 1–3 takes polynomial time and results in a new instance (G', S') , where G' is an induced (and hence AT-free) subgraph of G and $S' \subseteq S$, such that (G', S') is a yes-instance of INDUCED DISJOINT PATHS if and only if (G, S) is a yes-instance of INDUCED DISJOINT PATHS.*




Fig. 2. A graph G with four terminal pairs and the corresponding graphs G_1, \dots, G_4 .




Fig. 3. The graph H that is constructed from the graph G shown in Fig. 2.

For convenience we denote the obtained instance by (G, S) as well, and we also assume that $|S| = k$.

For $i = 1, \dots, k$, let G_i denote the subgraph obtained from G after removing all terminal vertices not equal to s_i or t_i , together with all of their neighbours not equal to s_i or t_i (should s_i or t_i be adjacent to a terminal from some other pair), that is, G_i is the subgraph of G induced by $(V \setminus (\bigcup_{v \in T \setminus \{s_i, t_i\}} N[v])) \cup \{s_i, t_i\}$ (see Fig. 2). The following lemma is straightforward to see.

Lemma 3. *If s_i and t_i are in different connected components of G_i for some $1 \leq i \leq k$, then (G, S) is a no-instance of INDUCED DISJOINT PATHS.*

Hence, we can add the following preprocessing step, which we can perform in polynomial time.

Step 4. If some s_i and t_i are in two different connected components of G_i , then return no.

Summarizing, applying Steps 1–4 takes polynomial time and results in the following additional conditions for our instance:

Condition ii) The terminals of each pair (s_i, t_i) are not adjacent.

Condition iii) The terminals of each pair (s_i, t_i) are in the same connected component of G_i .

3.2. Phase 2: obtaining structural results

Let G be an AT-free graph with a set $S = \{(s_1, t_1), \dots, (s_k, t_k)\}$ of terminal pairs satisfying Conditions i)–iii). In $G[T]$, there is no edge between any two terminals of the same pair due to Condition ii). We extend $G[T]$ by joining the terminals of each pair (s_i, t_i) by a path P_i of length 2, that is, for each pair (s_i, t_i) we introduce a new vertex that we make adjacent only to s_i and t_i . We denote the resulting graph by H ; see Fig. 3 for an example. Note that $G[T]$ is an induced subgraph of H . The inner vertex of each P_i is called a *path-vertex*, and the two edges of each P_i are called *path-edges*.

We now prove the following lemma; note that in this lemma $H' = H$ is possible.

Lemma 4. *The pair (G, S) is a yes-instance of INDUCED DISJOINT PATHS if and only if G contains an induced subgraph that is isomorphic to H' , where H' is obtained from H after subdividing a number of path-edges one or more times, such that the isomorphism maps every terminal of G to the same terminal of H .*

Proof. Suppose that (G, S) is a yes-instance of INDUCED DISJOINT PATHS and let P_1, \dots, P_k be a solution to (G, S) . Observe that the only terminal vertices to which any internal vertices of a path P_i are adjacent are the terminal vertices that

represent s_i and t_i . Hence, the union of $G[T]$ and the paths P_1, \dots, P_k forms a graph H' that satisfies the conditions of the lemma. The converse is readily seen. \square

Lemma 4 implies that the pair (G, S) is a yes-instance of INDUCED DISJOINT PATHS if and only if G contains H as an induced topological minor anchored in the terminals of T . Moreover, Lemma 4 immediately implies the next lemma.

Lemma 5. *If (G, S) is a yes-instance of INDUCED DISJOINT PATHS then H is AT-free.*

Lemma 5 yields the following step, which can be straightforwardly implemented in polynomial time by checking for the existence of an asteroidal triple.

Step 5. If H is not an AT-free graph, then return no.

From now on, we assume that H is AT-free. We deduce the following two lemmas.

Lemma 6. *The graph $G[T]$ is triangle-free.*

Proof. In order to obtain a contradiction, assume that $G[T]$ contains a triangle. Then $H[T]$ also contains this triangle. By Condition ii), each vertex of the triangle is contained in at least one terminal pair that contains a vertex outside the triangle. Then let $(s_i, t_i), (s_j, t_j), (s_r, t_r)$ be three terminal pairs for which, without loss of generality, s_i, s_j, s_r induce the triangle. By Condition ii), s_i, s_j, s_r are not adjacent to t_i, t_j, t_r , respectively. In H , let x_i, x_j, x_r be the path-vertices of the paths P_i, P_j, P_r , respectively. Recall that path-vertices have degree 2. Then x_i, x_j, x_r form an asteroidal triple in H , a contradiction to the AT-freeness of H . \square

Lemma 7. *Every vertex $u \in T$ is included in at most five terminal pairs.*

Proof. In order to obtain a contradiction, assume that T contains a vertex u that is included in six terminal pairs $(s_1, t_1), \dots, (s_6, t_6)$. We assume without loss of generality that $u = s_1 = \dots = s_6$. Among the vertices t_1, \dots, t_6 , it follows from the Ramsey number $R(3, 3) = 6$ that we have either three pairwise adjacent or three pairwise non-adjacent vertices. In the first case, $G[T]$ contains a triangle. However, this is not possible due to Lemma 6. In the second case, let $t_{i_1}, t_{i_2}, t_{i_3}$ denote the three pairwise non-adjacent vertices. Then we find that $t_{i_1}, t_{i_2}, t_{i_3}$ form an asteroidal triple in H , because $t_{i_1}, t_{i_2}, t_{i_3}$ are joined to u by three mutually induced paths $P_{i_1}, P_{i_2}, P_{i_3}$, a contradiction to the AT-freeness of H . \square

Let H_1, \dots, H_r be the connected components of H . We observe that the terminals of each pair are in the same connected component of H due to the paths P_i . Hence, we can define the set $S_i \subseteq S$ of terminal pairs in a connected component H_i . We write $T_i = V_{H_i} \cap T$ to denote the set of terminals in H_i .

Lemma 8. *Each H_i has a dominating pair $\{x_i, y_i\}$ with $x_i, y_i \in T_i$. Moreover, such a dominating pair can be found in linear time.*

Proof. Each connected component H_i is a connected AT-free graph. Hence, H_i has a dominating pair of vertices due to Lemma 1. Moreover, by the same lemma, there exists a linear-time algorithm that finds such a dominating pair $\{u, v\}$. If both u and v are in T_i , then we are done. Suppose that this is not the case, say $u \notin T_i$. Then u must be the path-vertex of some (s_j, t_j) -path P_j in H_i .

We claim that $\{s_j, v\}$ or else $\{t_j, v\}$ is a dominating pair. This can be seen as follows. If $\{s_j, v\}$ is not a dominating pair, then there is an (s_j, v) -path Q_1 that does not dominate H_i . This path Q_1 does not use vertex u ; otherwise it would contain a subpath from u to v , which dominates H_i because $\{u, v\}$ is a dominating pair. Hence, we can extend Q_1 by u to obtain the path uQ_1 . This path is an (u, v) -path. Consequently, uQ_1 dominates H_i . As u is only adjacent to s_j and t_j , and Q_1 does not dominate H_i , we find that Q_1 does not dominate t_j . In other words, Q_1 avoids t_j and the neighbourhood of t_j . By the same arguments, there exists a (t_j, v) -path Q_2 that avoids s_j and the neighbourhood of s_j . Condition ii) tells us that s_j and t_j are not adjacent. As $v \in V_{Q_1} \cup V_{Q_2}$, we find that v is not adjacent to s_j, t_j . Moreover, u is only adjacent to s_j and t_j . Hence, s_j, t_j, v form an asteroidal triple, a contradiction to the AT-freeness of H_i .

From the above, it is clear that we can decide in linear time if $\{s_j, v\}$ or $\{t_j, v\}$ is a dominating pair. We only have to check whether there exists an (s_j, v) -path that avoids t_j and the neighbourhood of t_j . If not, then $\{s_j, v\}$ is a dominating pair, and otherwise $\{t_j, v\}$ is a dominating pair. If $v \notin T_i$, then v is a path-vertex, and we can repeat the same arguments as we used for u to replace v by one of its neighbours in T_i . This completes the proof of Lemma 8. \square

For the dominating pairs $\{x_i, y_i\}$ found by Lemma 8, we compute in linear time a shortest (x_i, y_i) -path D_i in H_i for $i = 1, \dots, r$. The next two lemmas show a number of properties of these paths D_i .

Lemma 9. *Each D_i contains at least one of the terminals of every pair in S_i .*

Proof. The path D_i dominates all vertices of H_i . In particular, it dominates the path-vertex of the (s_j, t_j) -path for each $(s_j, t_j) \in S_i$. This vertex is only adjacent to s_j and t_j . Hence, $s_j \in V_{D_i}$ or $t_j \in V_{D_i}$. \square

Lemma 10. *Each vertex of D_i is adjacent to at most five path-vertices of H_i that are not on D_i and to at most two terminals that are not on D_i .*

Proof. Let u be a vertex of D_i . If u is not a terminal, then u is a path-vertex, and consequently, has degree 2. Hence, we may assume that u is a terminal. Lemma 7 tells us that u can represent at most five terminals. This implies that u is adjacent to at most five path-vertices not on D_i .

We now show that u is adjacent to at most two terminals not on D_i . To obtain a contradiction, suppose that u is adjacent to three terminals. Since at least one terminal of each terminal pair belongs to D_i due to Lemma 9, these three terminals are from three different pairs. Hence we may without loss of generality denote them by $s_{i_1}, s_{i_2}, s_{i_3}$. Let v, w, z be the path-vertices of the (s_{i_1}, t_{i_1}) -, (s_{i_2}, t_{i_2}) - and (s_{i_3}, t_{i_3}) -paths $P_{i_1}, P_{i_2}, P_{i_3}$, respectively. Note that $u \notin \{t_{i_1}, t_{i_2}, t_{i_3}\}$, because H_i contains no edge between two terminals of the same pair by Condition (ii). This means that no vertex of $\{v, w, z\}$ is adjacent to u . Hence v, w, z form an asteroidal triple, a contradiction to the AT-freeness of H_i . \square

Recalling Lemma 4 and using Lemma 9, we obtain the following.

Lemma 11. *Let H'_i be an AT-free graph obtained from H_i by subdividing a number of path-edges one or more times. Let P'_j and D'_i be the resulting paths obtained from the paths P_j and D_i , respectively. Then,*

- a) *the length of each P'_j that is not a subpath of D'_i is at most 3 (implying that every internal vertex of P'_j is adjacent to at least one of s_j, t_j);*
- b) *D'_i dominates all but at most two vertices of H'_i .*

Proof. First, we prove a). In order to obtain a contradiction, assume that P'_j has length at least 4. Let u be an internal vertex of P'_j that is at distance at least 2 from s_j and t_j . By Lemma 9, at least one of the vertices s_j, t_j is a vertex of D'_i . Assume that $s_j \in V_{D'_i}$. Then either $t_j \in V_{D'_i}$ or t_j is adjacent to a vertex $z \in V_{D'_i}$.

First suppose that $t_j \in V_{D'_i}$. Let Q be an (s_j, t_j) -subpath of D'_i . Then Q has length at least 2, because Condition ii) tells us that s_j and t_j are not adjacent in G , and consequently, not in H' . Since u has no neighbours outside P'_j , we find that Q avoids the neighbourhood of u . However, then u, s_j, t_j form an asteroidal triple, a contradiction to the AT-freeness of H_i .

Now suppose that t_j is adjacent to a vertex $z \in V_{D'_i}$. We choose z in such a way that t_j is not adjacent to any vertex on the subpath Q of D'_i from s_j to z . The path Qt_j has length at least 2, because s_j and t_j are not adjacent in H . Moreover, Qt_j avoids the neighbourhood of u . Hence, u, s_j, t_j form an asteroidal triple, a contradiction to the AT-freeness of H .

To prove b), assume that D'_i does not dominate three vertices $u_1, u_2, u_3 \in V_{H'_i}$. Since we only subdivide path-edges and D_i dominates each terminal, we find that u_1, u_2, u_3 are internal vertices of some (s_j, t_j) -paths P'_j . By Lemma 9, at least one of the vertices s_j, t_j is a vertex of D'_i for each such P'_j , that is, D'_i dominates at least one internal vertex of P'_j . Then by a), each P'_j can have at most one internal vertex not dominated by D'_i . Hence, u_1, u_2, u_3 are internal vertices of distinct paths $P'_{j_1}, P'_{j_2}, P'_{j_3}$, respectively. We assume without loss of generality that $s_{j_1}, s_{j_2}, s_{j_3} \in V_{D'_i}$. Then the path composed of the (u_1, s_{j_1}) -subpath of P'_{j_1} , the (s_{j_1}, s_{j_2}) -subpath of D'_i , and the (s_{j_2}, u_2) -subpath of P'_{j_2} is a (u_1, u_2) -path in H'_i that avoids the neighbourhood of u_3 . By the same arguments, H'_i has a (u_1, u_3) -path that avoids the neighbourhood of u_2 and a (u_2, u_3) -path that avoids the neighbourhood of u_1 . We conclude that u_1, u_2, u_3 form an asteroidal triple in H'_i , a contradiction to the AT-freeness of H'_i . This completes the proof of Lemma 11. \square

Recall that Condition iii) tells us that the terminals s_i and t_i are in the same connected component of the graph G_i for $i = 1, \dots, k$. Two terminal pairs (s_i, t_i) and (s_j, t_j) are *interfering* if there is an induced (s_i, t_i) -path Q_i in G_i and an induced (s_j, t_j) -path Q_j in G_j , respectively, such that Q_i and Q_j are not mutually induced; see Fig. 4 for an example. We say that there is *interference* between two sets of terminal pairs S_i and S_j if a terminal pair from S_i and a terminal pair from S_j are interfering.

Lemma 12. *Let (s_i, t_i) and (s_j, t_j) be interfering terminal pairs from two different connected components of H . Let Q_i and Q_j be induced (s_i, t_i) - and (s_j, t_j) -paths in G_i and G_j , respectively, such that Q_i and Q_j are not mutually induced. Then Q_i and Q_j are vertex-disjoint. Moreover, for any edge $uv \in E_G$ with $u \in V_{Q_i}$ and $v \in V_{Q_j}$, both $u \in N(s_i) \cup N(t_i)$ and $v \in N(s_j) \cup N(t_j)$ hold.*

Proof. To obtain a contradiction, we first assume that Q_i and Q_j have a common vertex w . As (s_i, t_i) and (s_j, t_j) are terminal pairs from two different connected components of H , the four vertices s_i, t_i, s_j, t_j are pairwise distinct. As Q_i is a path in G_i , we find that Q_i avoids $N[s_j]$ and $N[t_j]$. Symmetrically, Q_j avoids $N[s_i]$ and $N[t_i]$. This means that w is not




Fig. 4. The graph I that is constructed from the graph G shown in Fig. 2. The terminal pairs (s_1, t_1) and (s_2, t_2) of G are interfering. Note that there is no interference between (s_2, t_2) and (s_3, t_3) .

adjacent to any of the vertices s_i, t_i, s_j, t_j . However, then s_i, t_i, s_j form an asteroidal triple, a contradiction to the AT-freeness of H .

Now assume that there is an edge $uv \in E_G$ such that $u \in V_{Q_i}, v \in V_{Q_j}$ and $u \notin N(s_i) \cup N(t_i)$. Recall that Q_i avoids $N[s_j]$ and $N[t_j]$, and that Q_j avoids $N[s_i]$ and $N[t_i]$. Then we obtain a contradiction, because s_i, t_i, s_j again form an asteroidal triple. \square

Lemma 12 implies that any interference between two terminal pairs (s_i, t_i) and (s_j, t_j) from different connected components of H stems from an edge whose endpoints are in $N(s_i) \cup N(t_i)$ and $N(s_j) \cup N(t_j)$, respectively. We now obtain the following lemma.

Lemma 13. *It is possible to check in polynomial time whether two terminal pairs (s_i, t_i) and (s_j, t_j) from different connected components of H are interfering.*

Proof. By Lemma 12, we only need to check if there exist four vertices $u_1 \in N(s_i), u_2 \in N(t_i), v_1 \in N(s_j), v_2 \in N(t_j)$ such that

- there exists a path Q_i in G_i from u_1 to u_2 that does not contain s_i or t_i ;
- there exists a path Q_j in G_j from v_1 to v_2 that does not contain s_j or t_j ; and
- there exists an edge uv in G with $u \in \{u_1, u_2\}$ and $v \in \{v_1, v_2\}$.

For each 4-tuple (u_1, u_2, v_1, v_2) , checking the above three conditions can be done in polynomial time. \square

Lemma 13 enables us to construct in polynomial time an auxiliary graph I with vertices $1, \dots, r$ and edges ij if and only if there is interference between S_i and S_j ; see Fig. 4 for an example. This leads to the following lemma.

Lemma 14. *The graph I is a disjoint union of paths.*

Proof. First, we prove that $d_I(i) \leq 2$ for all $1 \leq i \leq r$. To obtain a contradiction, assume that $d_I(i) \geq 3$. Let $j_1, j_2, j_3 \in \{1, \dots, r\}$ be three vertices of I adjacent to i . For any pair of terminals (s_p, t_p) from $S_i \cup S_{j_1} \cup S_{j_2} \cup S_{j_3}$, there is an (s_p, t_p) -path R_p in G_p by Condition iii). Moreover, for each $l \in \{1, 2, 3\}$, there are interfering terminal pairs (s_{i_l}, t_{i_l}) and (s_{j_l}, t_{j_l}) in S_i and S_{j_l} respectively. Hence, by definition, there exist an (s_{i_l}, t_{i_l}) -path Q_{i_l} in G_{i_l} and an (s_{j_l}, t_{j_l}) -path Q_{j_l} in G_{j_l} , respectively, such that Q_{i_l} and Q_{j_l} are not mutually induced.

By definition of the graphs G_i , the paths R_p, Q_{i_l} , and Q_{j_l} each avoid the closed neighbourhood of terminal vertices that are not their endpoints. In particular, for any pair S_a, S_b among S_i, S_{j_1}, S_{j_2} , and S_{j_3} , the paths among R_p, Q_{i_l} , and Q_{j_l} whose endpoints are in T_a avoid the closed neighbourhood of all terminal vertices in T_b , because H_i, H_{j_1}, H_{j_2} , and H_{j_3} are connected components of H .

Let R be the subgraph of G induced by the vertices of all the paths R_p, Q_{i_l} and Q_{j_l} . Let u_1, u_2, u_3 be arbitrary terminal vertices from $T_{j_1}, T_{j_2}, T_{j_3}$ respectively, and let v be a terminal from T_i . For each $l \in \{1, 2, 3\}$, there is a (v, u_l) -path in R that avoids the neighbourhoods of terminals u_p for $p \in \{1, 2, 3\}, p \neq l$. In particular, recalling the observations of the previous paragraph:

- the connectedness of H_i implies that using the edges of the paths R_p there is a path in R from v to s_{i_l} that avoids the closed neighbourhood of any terminal vertex in $T_{j_1}, T_{j_2}, T_{j_3}$;
- the interference of Q_{i_l} and $Q_{j_{l'}}$ implies there is a path in R from s_{i_l} to $s_{j_{l'}}$ that avoids the closed neighbourhood of any terminal vertex in $T_{j_{l''}}$ for $l' \neq l$;
- the connectedness of $H_{j_{l'}}$ implies that using the edges of the paths R_p there is a path in R from $s_{j_{l'}}$ to u_l that avoids the closed neighbourhood of any terminal vertex in $T_{j_{l''}}$ for $l' \neq l$ and in T_i .

Hence, u_1, u_2, u_3 form an asteroidal triple, a contradiction.




Fig. 5. The graphs G_1^* and G_2^* that are constructed from the graph G shown in Fig. 2.

It remains to prove that I has no cycles. In order to obtain a contradiction, assume that $i_0 i_1 \dots i_q$ with $i_0 = i_q$ is a cycle in I . For any pair of terminals (s_p, t_p) from $S_{i_1} \cup \dots \cup S_{i_q}$, there is an (s_p, t_p) -path R_p in G_p by Condition iii). Moreover, for each $l \in \{1, \dots, q\}$, there are interfering terminal pairs (s_{l_1}, t_{l_1}) and (s_{l_2}, t_{l_2}) in $S_{i_{l-1}}$ and S_{i_l} respectively. Hence, by definition, there exists an (s_{l_1}, t_{l_1}) -path Q_{l_1} in G_{l_1} and an (s_{l_2}, t_{l_2}) -path Q_{l_2} in G_{l_2} , respectively, such that Q_{l_1} and Q_{l_2} are not mutually induced.

Let R be the subgraph of G induced by the vertices of all the paths R_p, Q_{l_1} and Q_{l_2} . Let u_1, \dots, u_q be arbitrary terminal vertices from H_{i_1}, \dots, H_{i_q} , respectively. Observe that for each $j \in \{1, \dots, q\}$, there is a (u_{j-1}, u_j) -path in R that avoids the neighbourhood of u_l for each $l \in \{1, \dots, q\}, l \neq j - 1, j$. The existence of such a path follows along the same lines as before. Hence, u_1, u_2, u_3 form an asteroidal triple, a contradiction. \square

Let I_1, \dots, I_l be the connected components of I , and let J_1, \dots, J_l be their vertex sets, respectively. For $h = 1, \dots, l$, we define $X_h = \{(s_j, t_j) \mid (s_j, t_j) \in \bigcup_{p \in J_h} S_p\}$, and we let G_h^* be the graph obtained from G by removing all vertices of the closed neighbourhoods of all the terminals that are not included in X_h ; see Fig. 5 for an example. Lemma 15 shows how G is related to the graphs G_h^* .

Lemma 15. *The instance (G, S) is a yes-instance of INDUCED DISJOINT PATHS if and only if (G_h^*, X_h) is a yes-instance of INDUCED DISJOINT PATHS for all $1 \leq h \leq l$.*

Proof. Clearly, a solution for (G, S) induces a solution for the instances (G_h^*, X_h) . Suppose that we have a solution for each (G_h^*, X_h) . Then two paths from two different solutions are mutually induced, because no two terminal pairs from two different sets X_h and $X_{h'}$ are interfering. Hence, the union of these solutions for the instances (G_h^*, X_h) forms a solution for (G, S) . \square

Lemma 15 gives us the following step.

Step 6. If I is disconnected, then solve INDUCED DISJOINT PATHS for each (G_h^*, X_h) . Return `yes` if the return value is `yes` for all of these instances and return `no` otherwise.

By Step 6, we may assume that I is connected. Then, by Lemma 14, we may assume that I is a path. This leads to the following new condition:

Condition iv) There is interference between two sets S_i and S_j for some $1 \leq i < j \leq r$ if and only if $j = i + 1$.

We are now ready to be a bit more precise, which is necessary for our algorithm. For $i = 1, \dots, r - 1$, let W_i be the set of all vertices $u \in V_G$, such that there is an edge $uv \in E_G$ with the following property: there are interfering terminal pairs (s_p, t_p) and (s_q, t_q) in S_i and S_{i+1} , respectively, such that G_p has an induced (s_p, t_p) -path containing u and G_q has an induced (s_q, t_q) -path containing v . Using this definition we can state the next lemma.

Lemma 16. *For $i = 1, \dots, r - 1$, there is a set $Z_i \subseteq T_i$ of at most two terminals such that $W_i \subseteq N(Z_i)$.*

Proof. By Lemma 12, each vertex from W_i is in the neighbourhood of some terminal vertex in T_i . By Lemma 6, $G[T]$ is triangle free. Hence, to obtain a contradiction, it is sufficient to suppose that there are two non-adjacent terminal vertices z_1, z_2 in T_i and vertices $u_1 \in (N(z_1) \setminus N(z_2)) \cap W_i, u_2 \in (N(z_2) \setminus N(z_1)) \cap W_i$. For any pair of terminals (s_j, t_j) from S_i or S_{i+1} , there is an (s_j, t_j) -path R_j in G_j by Condition iii). Moreover, by the definition of W_i , there are vertices v_1 and v_2 adjacent to u_1 and u_2 , respectively, such that there are two terminal pairs (s_p, t_p) and (s_q, t_q) in S_{i+1} and there are (s_p, t_p) and (s_q, t_q) -paths Q_p and Q_q in G_p and G_q that pass through v_1 and v_2 , respectively.

Consider the subgraph R of G induced by all the vertices of the paths R_j, Q_p, Q_q , and the vertices u_1, u_2 . Let w be an arbitrary terminal in S_{i+1} . It remains to observe that:

- the connectedness of H_i implies that using the paths R_j there is a path in R from z_1 to z_2 that avoids the closed neighbourhood of any terminal vertex in T_{i+1} , and in particular, of w ;

- the fact that $z_2 \notin T_{i+1}$ and $u_1 \notin N(z_2)$ implies that using the vertices u_1, v_1 , and the path Q_p there is a path in R from z_1 to s_p that avoids the closed neighbourhood of z_2 and of any terminal in $T_{i+1} \setminus \{s_p\}$;
- similarly, there is a path in R from z_2 to s_q that avoids the closed neighbourhood of z_1 and of any terminal in $T_{i+1} \setminus \{s_q\}$;
- the connectedness of H_{i+1} implies that using the paths R_j there is a path in R from s_p to w and from s_q to w that avoids the closed neighbourhood of any terminal vertex in T_i , and in particular, of z_1 and z_2 .

Hence, z_1, z_2, w form an asteroidal triple, a contradiction. \square

We obtain the following result on the sets Z_1, \dots, Z_{r-1} defined in Lemma 16.

Lemma 17. *The sets Z_1, \dots, Z_{r-1} can be found in polynomial time.*

Proof. Let $1 \leq i \leq r - 1$. By Lemma 12 we can compute the set W_i in polynomial time (see also the proof of Lemma 13). By Lemma 16, it suffices to determine as Z_i a set of at most two of terminals in T_i that are adjacent to every vertex of W_i . This can be done in polynomial time as well. \square

3.3. Phase 3: dynamic programming

We are now ready to give a dynamic-programming algorithm for INDUCED DISJOINT PATHS. For simplicity, we solve the decision problem here, that is, we only check for the existence of paths, but the algorithm can be modified to get the paths themselves (if they exist).

Our algorithm is based on the information obtained from Phases 1 and 2, which leads to the following intuition. By Condition iv), there is only interference between sets S_i and S_{i+1} for $1 \leq i < r$. By Lemma 12, this interference is restricted to the neighbourhoods of the terminals of the pairs in S_i and S_{i+1} . Even stronger, by the definition of W_i and Lemma 16, the interference caused by terminals in T_i is restricted to the neighbourhoods of the terminal vertices in Z_i , where $|Z_i| \leq 2$. By Lemma 7, each terminal vertex represents at most five terminals. Hence, there are at most ten vertices in $N(Z_i)$ that cause interference between S_i and S_{i+1} . Unfortunately, we do not know these ten vertices. However, since $|Z_i| \leq 2$ and Z_i can be found in polynomial time (by Lemma 17), we can enumerate all possibilities for these ten vertices in polynomial time. In fact, it suffices to enumerate only the set Y of vertices among the ten vertices that are actually used by the paths in a solution.

This intuition leads to the following dynamic programming algorithm: for each $i \in \{1, \dots, r\}$ and $Y \subseteq N(Z_i)$ of size at most 10, we define a routine that solves INDUCED DISJOINT PATHS for the graph $F_i = G[V \setminus (\bigcup_{j \in \{i+1, \dots, r\}} \bigcup_{u \in T_j} N[u])]$ with the set of terminal pairs $S'_i = \bigcup_{j=1}^i S_j$ under the following additional condition: the set of non-terminal vertices from $N(Z_i)$ used by the paths in a solution is a subset of Y . Call this routine $\text{Sol}(i, Y)$. We execute this routine sequentially for $i = 1, \dots, r$, and we are clearly interested in deciding whether $\text{Sol}(r, Y)$ returns yes for some $Y \subseteq N(Z_r)$. The crux then is to show what the return value for $\text{Sol}(i, Y)$ for a particular value of i should be, using the return values for $\text{Sol}(i - 1, \cdot)$.

To this end, we construct the subroutine $\text{Component}(i, X, Y)$: for each $i \in \{1, \dots, r\}$ and for any two sets $X \subseteq N(Z_{i-1})$, $Y \subseteq N(Z_i)$ of size at most 10 each, the subroutine solves INDUCED DISJOINT PATHS for the graph $F_i = G[V \setminus (\bigcup_{j \in \{1, \dots, r\}, j \neq i} \bigcup_{u \in T_j} N[u])]$ with set of terminal pairs S_i under the following two additional conditions:

- a) the paths from a solution are not adjacent to the vertices of X , and
- b) the set of non-terminal vertices from $N(Z_i)$ used by the paths in a solution is a subset of Y .

Then, to determine the return value of $\text{Sol}(i, Y)$, we consider all $X \subseteq N(Z_{i-1})$ for which $\text{Sol}(i - 1, X)$ returns yes and let $\text{Sol}(i, Y)$ return yes if $\text{Component}(i, X, Y)$ returns yes. Otherwise, we let $\text{Sol}(i, Y)$ return no.

The correctness of the above algorithm is clear from the preceding discussion, the description of the dynamic programming routine Sol , and the description of the subroutine Component . Moreover, the algorithm will run in polynomial time if the desired subroutine Component can be constructed to run in polynomial time. We claim that Component can indeed be so constructed, and the remainder of this section is dedicated to proving this claim.

3.3.1. Constructing the component subroutine

We are looking for an induced subgraph H'_i in F_i obtained by subdivisions P'_i of the paths P_i of H_i (under some additional constraints). We call this a *solution*. Recall that D_i is a dominating path of H_i that joins terminals x_i and y_i , which means that D_i is a shortest path between x_i and y_i such that each vertex of H_i lies on D_i or has a neighbour on D_i . We construct D_i by a breadth-first search that puts path vertices in its queue first (this is only done to improve the running time and does not influence correctness).

In order to find H'_i , we will attempt to “trace” the (x_i, y_i) -path D'_i in F_i that is in the solution H'_i and is a subdivision of D_i . By Lemma 10, we are interested in only a bounded number of vertices adjacent to the vertices of D'_i , and by Lemma 11, the paths outside D'_i have length at most 3 and D'_i dominates almost all vertices of these paths. Hence, vertices of H'_i are “close” to D'_i and their number is “small”. The crux, therefore, is to find D'_i . To this end, we devise a dynamic-programming

algorithm that maintains a small part of D'_i and of the solution that we are building. We show that this is indeed possible by relying on the AT-freeness of F_i .

To formalize the preceding intuition, we require some definitions. Let u_1, \dots, u_p be the terminal vertices on D_i enumerated in path order, that is, $x_i = u_1$ and $y_i = u_p$. Recall that P_l is an (s_l, t_l) -path in H_i for $(s_l, t_l) \in S_i$; each P_l has a single internal vertex, its path vertex. For $j \in \{1, \dots, p\}$, let U_j be the set of terminal vertices in $(N_G(u_j) \setminus \{u_1, \dots, u_p\}) \cup \{u_j\}$. Observe that $\bigcup_{j=1}^p U_j = T_i$, because D_i is a dominating path of H_i and if D_i contains a path vertex (i.e. a non-terminal vertex of H_i), then D_i also contains both its neighbours, the vertices of the corresponding terminal pair (note that the ends of D_i are terminal vertices). From this observation and Condition ii), for any pair $(s_l, t_l) \in S_i$, we have that $s_l \in U_j$ and $t_l \in U_{j'}$ for some $j, j' \in \{1, \dots, p\}$ and $j \neq j'$. To simplify our arguments, we assume without loss of generality that if $s_l \in U_j$ and $t_l \in U_{j'}$, then $j < j'$. For a set R' of vertices, we let $U(R') = \bigcup_{j \in R'} U_j$, where the union is over all terminals in both R' and D_i . For a vertex z of D_i , we say that a terminal $u \in T_i$ is *behind* z if $u \in U_j$ for u_j in the (x_i, z) -subpath of D_i ; otherwise we say that a terminal $u \in T_i$ is *ahead* of z . By our earlier observation, the notions of behind and ahead of z induce a partition of T_i .

Now, as mentioned before, we would like to “trace” an (x_i, y_i) -path D'_i in F_i that is a subdivision of D_i . We will argue that, in the dynamic program, it suffices to maintain only the last five vertices of the path D'_i under construction. To this end, we require the following structural property.

Lemma 18. *Let $C = \{x_1, \dots, x_t\}$ be the ordered set of vertices of a cycle in an AT-free graph for some $t \geq 3$. Then there exist integers i, j with $1 \leq i < t, 2 \leq j \leq 4$ and $i + j \leq t$ such that x_i and x_{i+j} are adjacent.²*

Proof. Suppose C has a chord. Then there exist indices α, β with $1 \leq \alpha < \alpha + 1 < \beta \leq t$ such that $x_\alpha x_\beta$ is an edge. Choose the chord such that $\beta - \alpha$ is smallest. If $\beta - \alpha \leq 4$, then the lemma holds. Hence, we may assume otherwise. Note that for any integers γ, δ with $\alpha \leq \gamma < \gamma + 1 < \delta \leq \beta$ it holds that either $v_\gamma v_\delta$ is not an edge (this would contradict the choice of α, β) or $\gamma = \alpha$ and $\delta = \beta$. Hence, the vertices of v_α, \dots, v_β induce a cycle in the graph, which has length at least 6 since $\beta - \alpha > 4$. Such a chordless cycle trivially has an asteroidal triple for the graph (for example, take $v_\alpha, v_{\alpha+2}, v_{\alpha+4}$), a contradiction.

It follows that the cycle has no chord. If $|C| \leq 5$, then the lemma follows immediately. Otherwise, the chordless cycle trivially contains an asteroidal triple for the graph (for example, take v_1, v_3, v_5), a contradiction. \square

We can now proceed with the description of the algorithm but first need a definition. Namely, a pair of paths P'_a, P'_b that are subdivisions of P_a, P_b in G_a and G_b , respectively, *conflicts* if P'_a intersects P'_b in a vertex v or there is an edge $vw \in E_G$ such that v is an internal vertex of P'_a and w is an internal vertex of P'_b . In the former case, we call v an *intersection conflict vertex*, or simply a *conflict vertex*. In the latter case, we call vw a *conflict edge* and both v and w *conflict vertices*.

Note that it would appear another potential conflict between paths P'_a, P'_b in G_a and G_b can occur, namely if there is an edge between an endpoint v of (say) P'_a and an internal vertex w of P'_b . However, then $w \notin V_{G_b}$ by definition (a contradiction) or $v \in \{s_b, t_b\}$ (which is allowed). Hence, if we are able to find paths P'_l for all $(s_l, t_l) \in S_i$ such that P'_l is an induced path in G_l and the paths do not conflict, then these paths form a mutually induced set of paths.

Intuitively, ensuring there are no conflicts would be relatively straightforward to guarantee for pairs whose path vertices are not on D_i , because by combining that they should have length at most 3 by Lemma 11 and an argument that we need to consider few such pairs at any point in time, this limits the potential of conflict to polynomially many choices. However, for the other terminal pairs, this is much more complicated, because their paths P'_l can have arbitrary length.

We need two more definitions. Let $O_i \subseteq S_i$ be the set of pairs $(s_a, t_a) \in S_i$ for which the path vertex of P_a is on D_i . A *walk* W in G is an ordered set $\{v_1, \dots, v_{|W|}\}$ of vertices of G such that $v_i v_{i+1} \in E_G$ for $i \in \{1, \dots, |W| - 1\}$.

Our algorithm is built on the intuition that we only ensure that D'_i and the paths P'_l with $l \in O_i$ are walks. In AT-free graphs, however, Lemma 18 will help prove that by considering the last five vertices of D'_i , we avoid all conflicts.

Definition. A *realization* $(H'_i, D'_i, \{P'_l\})$ for $(H_i, D_i, \{P_l\})$ consists of:

- an induced subgraph H'_i of F_i such that no non-terminal vertices in H'_i are adjacent to X and all non-terminal vertices in $V_{H'_i} \cap N(Z_i)$ are in Y ;
- a walk D'_i in H'_i that is a subdivision of D_i ; in particular, it visits u_1, \dots, u_p in order (without duplication);
- for each $(s_l, t_l) \in S_i \setminus O_i$, there is an induced path P'_l in H'_i and in G_l from s_l to t_l of length 2 or 3;
- for each $(s_l, t_l) \in O_i$, there is a walk P'_l in H'_i and G_l of length at least 2 from s_l to t_l and P'_l is a subwalk of D'_i .

A realization is an *embedding* if the set $\{P'_l\}$ of paths is a set of mutually induced disjoint paths.

² Note that the lemma does *not* use indices module t . Although the lemma is certainly true in that setting, this is not how we will use it.

Note that if $(H'_i, D'_i, \{P'_l\})$ is an embedding for $(H_i, D_i, \{P_l\})$, then the paths $\{P'_l\}$ form a solution. Moreover, the conditions on the paths imply that D'_i is an induced path of F_i . Indeed, then the only possible edges between vertices of D'_i would be between non-consecutive terminals on D'_i . However, these edges are part of G and thus H_i as well. As D'_i is a subdivision of D_i by the definition of a realization, this contradicts that D_i is a shortest path in H_i .

Observe that the preceding argument also implies that if D'_i is not an induced path, then this is due to a conflicting pair or because D'_i is not a subdivision of D_i .

Clearly, our goal is to find an embedding, which is stronger than a realization. As mentioned, our algorithm aims to find a realization. The crux is how to ensure that we end up with an embedding. The following lemma is crucial.

Lemma 19. *Suppose $(H'_i, D'_i, \{P'_l\})$ is a realization for $(H_i, D_i, \{P_l\})$ in an AT-free graph, but not an embedding. Then there is a subwalk R' of D'_i with $|R'| \leq 5$ such that:*

- R' is not an induced path, or
- $U(R')$ contains at least one of $\{s_a, t_a\}$ and at least one of $\{s_b, t_b\}$ for some conflict pair P'_a, P'_b ; moreover, if $(s_a, t_a) \in O_i$ (or $(s_b, t_b) \in O_i$), then R' contains a conflict vertex of P'_a (or P'_b).

Proof. We start with the following claims.

(C1) D'_i is a path, or a subwalk R' as in the lemma statement exists.

Suppose that D'_i is not a path. Let x be any vertex on the walk D'_i such that the distance between two occurrences of x on D'_i is minimum. Let $x = x_1, \dots, x_t = x$ be the subwalk of D'_i between those two occurrences; that is, $x \notin \{x_2, \dots, x_{t-1}\}$, x_2, \dots, x_{t-1} are distinct, and t is minimum over all choices of x . (Note that any x_i might still occur multiple times on D'_i outside of the subwalk.) It follows that the subwalk forms a cycle. Then by Lemma 18, there exist integers i, j with $1 \leq i \leq t - 1$ and $2 \leq j \leq 4$ such that $x_i x_{i+j}$ is an edge. Then let $R' = \{x_i, \dots, x_{i+j}\}$. This subwalk of D'_i is not an induced path and has $|R'| \leq 5$.

(C2) D'_i is an induced path, or a subwalk R' as in the lemma statement exists.

By (C1), we may assume that D'_i is a path. Let x, y be two non-consecutive vertices of D'_i such that xy is an edge. Consider the cycle $x = x_1, \dots, x_t = y$ formed by the edge xy and the part of the path D'_i between x and y . Then by Lemma 18, there exist integers i, j with $1 \leq i \leq t - 1$ and $2 \leq j \leq 4$ such that $x_i x_{i+j}$ is an edge. Then let $R' = \{x_i, \dots, x_{i+j}\}$. This subwalk of D'_i is not an induced path and has $|R'| \leq 5$.

(C3) All paths in $\{P'_l\}$ are induced paths, or a subwalk R' as in the lemma statement exists.

Note that a path P'_l with $(s_l, t_l) \in S_i \setminus O_i$ is an induced path by the definition of a realization. Since a path P'_l with $(s_l, t_l) \in O_i$ is a subwalk of D'_i , the claim now follows from (C2).

By the preceding claims, we may assume that D'_i is an induced path and that all paths in $\{P'_l\}$ are induced paths. Hence, what is preventing the realization from being an embedding are conflicts. We now consider any conflicting pair and show that a subwalk R' as in the lemma statement exists.

Suppose that P'_a and P'_b with $(s_a, t_a), (s_b, t_b) \in O_i$ conflict. Since both are subwalks of D'_i by the definition of a realization, this implies that D'_i is either not a path or not induced. Then we obtain a contradiction, because (C1) or (C2) implies that a subwalk R' as in the lemma statement exists. Hence, we may now assume that no such conflicting pairs exist.

Suppose that P'_a and P'_b with $(s_a, t_a) \in O_i$ and $(s_b, t_b) \in S_i \setminus O_i$ conflict and have an intersection conflict vertex v . Since P'_b has length at most 3, v is a neighbour of s_b or t_b ; say $v \in N(s_b)$. Since P'_a is a path in G_a , but contains v , it follows that $s_a = s_b$ or $t_a = s_b$; say $s_a = s_b$. Then $v \in N(s_a)$ and thus v and s_a are consecutive vertices of P'_a and thus D'_i . Let $R' = \{v, s_a\}$. Note that $U(R')$ contains s_a and s_b and R' contains a conflict vertex of P'_a . Thus, a subwalk R' as in the lemma statement exists. Hence, we may now assume that no such conflicting pair exists. Since D'_i is a subdivision of D_i by the definition of a realization, this also implies that we may assume that no path P'_b with $(s_b, t_b) \in S_i \setminus O_i$ has an internal vertex that is a vertex of D'_i .

Suppose that P'_a and P'_b with $(s_a, t_a) \in O_i$ and $(s_b, t_b) \in S_i \setminus O_i$ conflict and have a conflict edge vw but no intersection conflict vertex. Assume $v \in P'_a$ and $w \in P'_b$. By Lemma 9, without loss of generality $s_b \in D_i$ and thus $s_b \in D'_i$. Consider the cycle C_{vw} that contains the edge vw and the vertex s_b of the graph induced by P'_b and D'_i . Note that such a cycle indeed exists and is uniquely defined, and consists of the edge vw , a part of P'_b and a part of D'_i . Now choose a, b and vw such that the length of C_{vw} is minimum.

If $|C_{vw}| > 5$, then by Lemma 18, C_{vw} has a chord xy . This chord cannot have both endpoints in D'_i or both in P'_b , or (C2) or (C3) implies that a subwalk R' as in the lemma statement exists. In particular, combined with the fact that P'_a and P'_b are paths in G_a and G_b respectively by the definition of a realization, this implies that the set of endpoints of the chord is disjoint from $\{s_a, t_a, s_b, t_b\}$. Hence, the chord xy is a conflict edge of P'_a and P'_b and the cycle C_{xy} induced by it is shorter than C_{vw} . This contradicts the choice of vw .

It follows that $|C_{vw}| \leq 5$. Let R' be the subpath of C_{vw} between s_b and w . Note that this is a subpath of D'_i . Moreover, $|R'| \leq 4$, because $|C_{vw}| \leq 5$ and v is not on R' . Since $s_b \in R'$, $s_b \in U(R')$. Moreover, since P'_a is a subpath of D'_i and s_b is a terminal on R' , it must be that s_a or t_a is in R' , and thus at least one of them is in $U(R')$. Finally, we note that $w \in R'$ by definition. Thus, a subwalk R' as in the lemma statement exists. Hence, we may now assume that no such conflicting pair exists.

Suppose that P'_a and P'_b with $(s_a, t_a) \in S_i \setminus O_i$ and $(s_b, t_b) \in S_i \setminus O_i$ conflict and have an intersection conflict vertex v . Since P'_b has length at most 3, v is a neighbour of s_b or t_b ; say $v \in N(s_b)$. Since P'_a is a path in G_a , but contains v , it follows that $s_a = s_b$ or $t_a = s_b$; say $s_a = s_b$. If s_a (and thus s_b) is on D_i , then let $R' = \{s_a\}$ and we obtain a subwalk R' as in the lemma statement. Otherwise, there is a terminal vertex u in D_i such that $s_a \in N(u)$. Note that $u \in D'_i$. Let $R' = \{u\}$; then $s_a \in U(R')$ and we obtain a subwalk R' as in the lemma statement. Hence, we may now assume that no such conflicting pair exists.

Suppose that P'_a and P'_b with $(s_a, t_a) \in S_i \setminus O_i$ and $(s_b, t_b) \in S_i \setminus O_i$ conflict and have a conflict edge vw but no intersection conflict vertex. Assume $v \in P'_a$ and $w \in P'_b$. By Lemma 9, we may assume without loss of generality that $s_a \in D_i$ and $s_b \in D_i$. Hence, $s_a \in D'_i$ and $s_b \in D'_i$. Consider the cycle C_{vw} formed by the edge vw , the subpath of P'_a between v and s_a , the subpath of D'_i between s_a and s_b and the subpath of P'_b between s_b and w . Choose a, b and vw such that the length of C_{vw} is minimum.

If $|C_{vw}| > 5$, then by Lemma 18, C_{vw} has a chord xy . This chord cannot have both endpoints in D'_i , both in P'_a , or both in P'_b , or (C2) or (C3) implies that a subwalk R' as in the lemma statement exists. Moreover, by our earlier assumption, the chord cannot have an endpoint in D'_i and an endpoint in one of P'_a or P'_b . Also, combined with the fact that P'_a and P'_b are paths in G_a and G_b respectively by the definition of a realization, this implies that the set of endpoints of the chord is disjoint from $\{s_a, t_a, s_b, t_b\}$. Hence, the chord xy is a conflict edge of P'_a and P'_b and the cycle C_{xy} induced by it is shorter than C_{vw} . This contradicts the choice of vw .

It follows that $|C_{vw}| \leq 5$. Let R' be the subpath of C_{vw} between s_a and s_b . Note that this is a subpath of D'_i . Moreover, $|R'| \leq 3$, because $|C_{vw}| \leq 5$ and v and w are not on R' (as v and w are internal vertices of pairs $(s_a, t_a) \in S_i \setminus O_i$ and $(s_b, t_b) \in S_i \setminus O_i$). Since $s_a, s_b \in R'$, $s_a, s_b \in U(R')$. Thus, a subwalk R' as in the lemma statement exists. This completes our proof. \square

3.3.2. Dynamic program

We now define the table used in the subroutine $\text{Component}(i, X, Y)$. Each entry uses the following fields:

- an integer $\ell \geq 0$;
- an ordered set of vertices $R' \subseteq V_{F_i}$ of at most five vertices of F_i such that $F_i[R']$ is a path;
- a vertex $z \in D_i$;
- a set of non-terminal vertices $N' \subseteq V_{F_i}$;

and stores whether there exists a subgraph H'_i of F_i with the following properties:

- no non-terminal vertices in H'_i are adjacent to X and all non-terminal vertices in $V_{H'_i} \cap N_{F_i}(Z_i)$ are in Y ;
- there is a walk D'_i in H'_i of length ℓ that is a subdivision of the subpath of D_i from x_i to z ;
- for each $(s_l, t_l) \in S_i \setminus O_i$ for which at least one terminal is behind z , there is an induced path P'_l in H'_i and G_l from s_l to t_l of length 2 or 3; if s_l or t_l is in $U(R')$, then all internal vertices of this path are in N' ;
- for each $(s_l, t_l) \in O_i$ for which both terminals are behind z , there is a walk P'_l in H'_i and G_l of length at least 2 from s_l to t_l and P'_l is a subwalk of D'_i ;
- if a single terminal of $(s_l, t_l) \in O_i$ is behind z and z is a path vertex, there is a walk P'_l in H'_i and G_l of length at least 2, starting in s_l but not ending in t_l , and P'_l is a subwalk of D'_i ;
- the path $F_i[R']$ and all paths for terminal pairs $(s_l, t_l) \in S_i \setminus O_i$ for which s_l or t_l is in $U(R')$ are mutually induced disjoint;
- R' is equal to the last $\min\{5, |D'_i|\}$ vertices of D'_i ;
- $N' \subseteq V_{H'_i}$;
- there is a subpath R of D_i , ending in z , and a map from R' to R such that any terminal vertex v on R' is mapped to the same vertex on R and any non-terminal vertex v on R' is mapped to a path vertex on R . Moreover, the map respects the order of the terminal and path vertices of R .

The last condition is called the *mappability condition* and is fully determined by R' and z . We can always pick the subpath R to have length at most the length of R' , and thus at most 5. The first five conditions effectively ensure that H'_i is a partial embedding of all terminal pairs with a terminal vertex behind z .

We use $\text{last}(R)$ respectively $\text{last}(R')$ to denote the last vertex of R respectively R' . We also use $\text{term}(z)$ to be equal to $(s_l, t_l) \in O_i$ for which either z is the path vertex or the vertex immediately following z on D_i is the path vertex; otherwise, $\text{term}(z)$ is undefined.

Table construction. The table is initialized by having an entry for $\ell = 0$, $z = x_i$, $R' = \{x_i\}$, and each set of non-terminal vertices N' for which $F_i[N' \cup \{x_i\}]$ satisfies the conditions of an entry. In particular, for each $(s_i, t_i) \in S_i \setminus O_i$ for which $\{s_i, t_i\} \cap U(R') \neq \emptyset$, there is an induced path P'_i in H'_i and G_i from s_i to t_i of length 2 or 3 for which all internal (non-terminal) vertices are in N' , and all such paths are mutually induced disjoint. Also, no non-terminal vertices in N' are adjacent to X and all non-terminal vertices in $N' \cap N_{F_i}(Z_i)$ are in Y .

Suppose the table has been fully filled for $\ell - 1$. Consider all vertices $u \in V_{F_i}$ for which the table contains an entry for $(\ell - 1, R', z, N')$ such that u is adjacent to $\text{last}(R')$ and is not adjacent to any other vertices in R' . For each such u and each entry $(\ell - 1, R', z, N')$ with these properties, we create (possibly) new entries for the table for ℓ . We call these new entries *successor entries*. We distinguish two cases.

Case 1. $u \notin T_i$.

Note that in this case, u will be an internal vertex of a path in the solution for a terminal pair in O_i . Hence, it cannot be adjacent to internal vertices of other paths and terminals except, possibly, the terminals joined by the path. We only construct new entries if the following conditions hold:

- u is not adjacent to any vertex in N' ;
- u is not adjacent to X and if it is in $N_{F_i}(Z_i)$, then it is in Y ;
- z is a path vertex or the vertex on D_i after z is a path vertex;
- u is a vertex in G_i for the terminal pair $(s_i, t_i) = \text{term}(z)$.

If the above holds, then entries (ℓ, R'_n, z_n, N'_n) are constructed as follows. If z is a path vertex, then $z_n = z$; otherwise, z_n is the vertex on D_i after z .

If $|R'| \leq 4$, then R'_n is equal to R' and u is added as the last vertex of the ordered set. Let R and R_n be the subpaths of D_i that follow from the mappability condition on R' and z respectively R'_n and z_n . Note that $R_n \supseteq R$ and the set of terminal vertices of R_n and R are equal. Hence, we can set $N'_n = N'$ by the above conditions.

If $|R'| = 5$, let x be the first vertex of R' . Then R'_n is equal to the last four vertices of R' and u is added as the last vertex of the ordered set. Let R and R_n be the subpaths of D_i that follow from the mappability condition on R' and z respectively R'_n and z_n . If x is a non-terminal vertex, then the sets of terminal vertices of R_n and R are equal and we can set $N'_n = N'$. Otherwise, obtain N'_n from N' by removing any vertices from N' belonging to a path for a terminal $(s_i, t_i) \in S_i \setminus O_i$ for which neither s_i nor t_i is in $U(R'_n)$. In either case, N'_n is correctly set by the above conditions.

Case 2. $u \in T_i$.

In this case, u can be adjacent to other terminals and u can be a terminal vertex for several terminal pairs. This makes the construction of entries slightly more involved. We only construct new entries if u is the vertex on D_i after z . In that case, entries (ℓ, R'_n, z_n, N'_n) are constructed as follows. Set $z_n = u$.

If $|R'| \leq 4$, then R'_n is equal to R' and u is added as the last vertex of the ordered set. If $|R'| = 5$, then R'_n is equal to the last four vertices of R' and u is added as the last vertex of the ordered set. Let R and R_n be the subpath of D_i that follow from the mappability condition on R' and z respectively R'_n and z_n .

Now construct N'_n from N' . If $|R'| = 5$, let x be the first vertex of R' . If x is a terminal vertex, then remove any vertices from N' belonging to a path for a terminal $(s_i, t_i) \in S_i \setminus O_i$ for which neither s_i nor t_i is in $U(R'_n)$.

We might also add several possible sets of paths to N'_n . For each terminal $(s_i, t_i) \in S_i \setminus O_i$ for which s_i or t_i is in $U(R'_n) \setminus U(R')$ while neither are in $U(R')$, consider all possible induced paths in G_i between s_i and t_i of which the internal vertices are not adjacent to any vertex in R' , R'_n , and N' . Consider all possible combinations of such paths for such terminal pairs, and add the internal vertices of them to N'_n . Note that this creates many different sets N'_n , each of which yields a new table entry.

Decision and analysis. We finish the description of the subroutine $\text{Component}(i, X, Y)$. It can be readily verified that the entries constructed as above satisfy all stated conditions. Finally, if there is an entry in the table for (ℓ, R', y_i, N') for any ℓ, R', N' , then return *yes*; otherwise, return *no*.

The crucial property of $\text{Component}(i, X, Y)$ is given in the following lemma.

Lemma 20. For each $i \in \{1, \dots, k\}$, the subroutine $\text{Component}(i, X, Y)$ tests in polynomial time whether $(H_i, D_i, \{P_i\})$ has an embedding.

Proof. Suppose that $(H_i, D_i, \{P_i\})$ has an embedding $(H'_i, D'_i, \{P'_i\})$. Let v_1, \dots, v_j be vertices of D'_i in order from $x_i = v_1$ to $y_i = v_j$. By induction, it can be readily shown from the construction of the entries that there is a sequence of successor entries $(\ell, \{v_{\ell-4}, \dots, v_\ell\}, z_\ell, N'_\ell)$ for $1 \leq \ell \leq j$ and appropriate z_ℓ, N'_ℓ that traces along D'_i for which there is an entry for $(j, \{v_{j-4}, \dots, v_j\}, y_i, N'_j)$ for some appropriate N'_j . Hence, the subroutine will return *yes*.

Suppose that the subroutine returns *yes*. Let $(\ell, R'_\ell, z_\ell, N'_\ell)$ be a sequence of successor entries for $1 \leq \ell \leq j$, R'_ℓ, z_ℓ, N'_ℓ for which there is an entry for (j, R'_j, y_i, N'_j) for some R'_j, N'_j . Let D'_i be the walk induced by the sets R'_ℓ . Let H'_i be the subgraph of F_i induced by the terminal vertices T_i , the walk D'_i , and $\bigcup_{1 \leq \ell \leq j} N'_\ell$. Let P'_i be the set of walks for terminal pairs (s_i, t_i) implied by the entries $(\ell, R'_\ell, z_\ell, N'_\ell)$.

Note that the walks for terminal pairs $(s_i, t_i) \in O_i$ are uniquely defined by D'_i . Then by Lemma 19, claims (C1) and (C2), and the fact that each set R'_ℓ is an induced path by definition, it follows that D'_i is an induced path.

For terminal pairs $(s_i, t_i) \in S_i \setminus O_i$, a small argument is needed, because potentially there is no R'_ℓ for which $s_i, t_i \in U(R'_\ell)$, meaning that a path for (s_i, t_i) appears, disappears, and then re-appears and is potentially different. Let ℓ_i be smallest such that $\{s_i, t_i\} \cap U(R'_{\ell_i})$. Without loss of generality, $s_i \in U(R'_{\ell_i})$. Suppose that $t_i \notin U(R'_{\ell_i}) \cup \dots \cup U(R'_{\ell_i+4})$. Note there is an induced path P'_i in G_i between s_i and t_i stored in $N'_{\ell_i}, \dots, N'_{\ell_i+4}$ that is the same path for each entry. Moreover, by construction, no vertices of this path are adjacent to vertices of $R'_{\ell_i}, \dots, R'_{\ell_i+4}$. Let v be the first vertex on D'_i after $U(R'_{\ell_i+4})$ that is adjacent to a vertex of P'_i . Since t_i is not in $U(R'_{\ell_i}), \dots, U(R'_{\ell_i+4})$ by definition and $t_i \in U(D'_i)$ by definition, v indeed exists. Then the subpath of D'_i from s_i to v plus the subpath of P'_i from s_i to the neighbour of v on P'_i induce a cycle in F_i . Since the subpath of D'_i from s_i to v contains at least six vertices, the length of the cycle is at least 6, a contradiction to the AT-freeness of F_i . Hence, the sequence of entries uniquely defines a path P'_i between s_i and t_i for each terminal pair $(s_i, t_i) \in S_i \setminus O_i$.

Consider $(H'_i, D'_i, \{P'_i\})$. By construction, we know that $(H'_i, D'_i, \{P'_i\})$ is a realization of $(H_i, D_i, \{P_i\})$. For sake of contradiction, suppose that $(H'_i, D'_i, \{P'_i\})$ is not an embedding of $(H_i, D_i, \{P_i\})$. Then by Lemma 19, there is a subwalk R' of D'_i with $|R'| \leq 5$ such that:

- R' is not an induced path, or
- $U(R')$ contains at least one of $\{s_a, t_a\}$ and at least one of $\{s_b, t_b\}$ for some conflict pair P'_a, P'_b ; moreover, if $(s_a, t_a) \in O_i$ (or $(s_b, t_b) \in O_i$), then R' contains a conflict vertex of P'_a (or P'_b).

However, the first case is excluded by the fact that each set R'_ℓ is an induced path by definition. The second case is also excluded. Let $(s_a, t_a) \in S_i$ and $(s_b, t_b) \in S_i$ be terminal pairs for which $U(R'_\ell)$ contains at least one of $\{s_a, t_a\}$ and at least one of $\{s_b, t_b\}$.

- If $(s_a, t_a) \in O_i$ and $(s_b, t_b) \in O_i$, then their paths cannot conflict, because R'_ℓ would contain conflict vertices of both pairs, but R'_ℓ is an induced path.
- If $(s_a, t_a) \in O_i$ and $(s_b, t_b) \in S_i \setminus O_i$, then their paths cannot conflict, because R'_ℓ would contain a conflict vertex and P'_b would be contained in N'_ℓ and thus is not be adjacent or intersect with R'_ℓ by construction.
- If $(s_a, t_a) \in S_i \setminus O_i$ and $(s_b, t_b) \in S_i \setminus O_i$, then their paths cannot conflict by the construction of N'_ℓ , because at least one of $\{s_a, t_a\}$ and at least one of $\{s_b, t_b\}$ and thus both P'_a and P'_b are contained in N'_ℓ .

It follows that $(H'_i, D'_i, \{P'_i\})$ is an embedding of $(H_i, D_i, \{P_i\})$.

To prove polynomial time complexity, it suffices to bound $|N'|$. We argue that if $(H_i, D_i, \{P_i\})$ has an embedding $(H'_i, D'_i, \{P'_i\})$ and $(\ell, R'_\ell = \{v_{\ell-4}, \dots, v_\ell\}, z_\ell, N'_\ell)$ for $1 \leq \ell \leq |D'_i|$ is the corresponding sequence of successor entries, then $|N'_\ell|$ is bounded by 47 for each ℓ . Hence, we may limit the size of N' to 47 during our dynamic program. By Lemma 11, D'_i will dominate all but at most two vertices of H'_i . We call these the special vertices. As argued in Lemma 11, the special vertices are internal vertices of some (s_i, t_i) paths P'_j .

We claim that all non-terminals in N'_ℓ that are not special are adjacent to a terminal in R' or the two terminal vertices preceding R' on D'_i or the two terminal vertices succeeding R' on D'_i . Call the latter two sets of terminal vertices K and L respectively. We consider two cases.

In the first case, consider a terminal pair (s_i, t_i) for which $t_i \in U(R'_\ell) \setminus D'_i$ (note that this implies that $(s_i, t_i) \in S_i \setminus O_i$). Then $s_i \in D'_i$ by Lemma 9. By Lemma 11, P'_i has length at most 3. If P'_i has length 2, then the single non-terminal vertex on P'_i is adjacent to s_i . If P'_i has length 3, then the vertex of P'_j that follows t_i is not dominated by D'_i , or the set of paths is not mutually induced. Note that this vertex is special. Moreover, the other internal vertex of P'_j is adjacent to s_i . Now let $v \in R'_\ell$ such that t_i is a neighbour of v in H'_i . Suppose that s_i precedes the vertices of K on D'_i (the case that s_i succeeds L on D'_i is similar). Then, as D'_i is a subdivision of D_i , the subpath of D_i between s_i and v has length at least 3, whereas the path in H_i between s_i and v via the path vertex of P_i and t_i has length at most 3. Then either we contradict that D_i is a shortest path or, since D_i was constructed by a breadth-first search that puts path vertices in its queue first, the former path would be preferred. Hence, $s_i \in K \cup R'_\ell \cup L$, implying the claim in the first case.

In the second case, consider a terminal pair (s_i, t_i) for which $t_i \in R'_\ell$ and $(s_i, t_i) \in S_i \setminus O_i$. If $s_i \notin D'_i$, then following the preceding argument, any internal vertices of P'_i that are not adjacent to t_i are not on D'_i , and thus special. Otherwise, $s_i \in D'_i$. By Lemma 11, P'_i has length at most 3, and thus its internal vertices are adjacent to s_i or t_i . Suppose that s_i precedes the vertices of K on D'_i (the case that s_i succeeds L on D'_i is similar). Again, the subpath of D_i between s_i and t_i has length at least 3, whereas the path in H_i between s_i and t_i via the path vertex of P_i has length at most 2, contradicting that D_i is a shortest path. Hence, $s_i \in K \cup R'_\ell \cup L$, implying the claim in the second case and overall.

Now following Lemma 10, any terminal vertex of D_i is adjacent to at most five path-vertices of H_i . This implies that any terminal vertex of D'_i is adjacent to at most five non-terminal vertices of H'_i . Following the claim, all non-terminals in N'_ℓ (except the special vertices) are adjacent to a terminal in $K \cup R' \cup L$. By definition, there are at most 9 such vertices, implying that $|N'_\ell| \leq 9 \cdot 5 + 2 = 47$. It follows that the size of the table is polynomially bounded. It can then be readily seen that the algorithm runs in polynomial time. \square

The full algorithm, its correctness, and its running time now follows from Lemma 20 and the preceding discussion.

Theorem 1. *The INDUCED DISJOINT PATHS problem can be solved in polynomial time for AT-free graphs.*

4. Induced topological minors

Recall that a graph G contains a graph H as an induced topological minor if G contains an induced subgraph isomorphic to a subdivision of H , that is, to a graph obtained from H by a number of edge subdivisions, and that the INDUCED TOPOLOGICAL MINOR problem is the corresponding decision problem. Recall also that a graph G contains a graph H on vertices x_1, \dots, x_k as an anchored induced topological minor if there exist k pre-specified vertices u_1, \dots, u_k in G , such that G contains an induced subgraph H' isomorphic to a subdivision of H and the isomorphism maps x_i to u_i for $i = 1, \dots, k$. In that case we say that H is anchored in $V_H \subseteq V_G$. The ANCHORED INDUCED TOPOLOGICAL MINOR is the corresponding decision problem.

As mentioned in Section 1 one can use an algorithm that solves INDUCED DISJOINT PATHS in polynomial time for any fixed integer k on some graph class \mathcal{G} as a subroutine for obtaining an algorithm that solves H -INDUCED TOPOLOGICAL MINOR in polynomial time for any fixed graph H on \mathcal{G} . This is because ANCHORED INDUCED TOPOLOGICAL MINOR can be reduced to INDUCED DISJOINT PATHS. This reduction is known, see for example Corollary 3 of the paper of Belmonte et al. [1], where \mathcal{G} is the class of chordal graphs. Due to Theorem 1, we can now also let \mathcal{G} be the class of AT-free graphs. Below we provide a short proof for this. In particular we do this in order to show that the reduction works when H has isolated vertices. This case is not yet captured as condition i) excludes the case when $s_i = t_i$ for some $1 \leq i \leq k$.

Corollary 1. *The ANCHORED INDUCED TOPOLOGICAL MINOR problem can be solved in polynomial time for AT-free graphs.*

Proof. Let G and H be graphs, such that H is anchored in $V_H \subseteq V_G$. If H has no isolated vertices, then we reduce ANCHORED INDUCED TOPOLOGICAL MINOR to INDUCED DISJOINT PATHS on G by constructing the pair of terminals (u, v) for each edge $uv \in E_H$.

Suppose now that H has at least one isolated vertex. Then for each isolated vertex u in H , we do the following. First, we check whether u is adjacent in G to some other vertex of H . If so, then we stop and return no . Otherwise, we remove u from H and remove u with its neighbourhood in G from the graph G . If we removed all vertices of H , then we return yes . Otherwise, denote by G' and H' the graphs obtained from G and H , respectively, by these removals. Now for each edge $uv \in E_{H'}$, we construct the pair of terminals (u, v) in G' and solve the obtained instance of INDUCED DISJOINT PATHS on G' . \square

If H is a fixed graph, that is, not part of the input, then INDUCED TOPOLOGICAL MINOR can be solved in polynomial time by a standard reduction, as mentioned in Section 1. To make our paper self-contained, we give a short proof.

Corollary 2. *The INDUCED TOPOLOGICAL MINOR problem can be solved in time $n^{k+O(1)}$ for pairs (G, H) where G is an n -vertex AT-free graph and H is a k -vertex graph.*

Proof. Let G be an AT-free graph on n vertices, and let H be a graph with $V_H = \{x_1, \dots, x_k\}$ for some $k \leq n$. First suppose that H has no isolated vertices. We guess k vertices u_1, \dots, u_k in G , and for each choice we check whether G has an anchored topological minor H' with the terminals u_1, \dots, u_k isomorphic to H , such that the isomorphism maps v_i to u_i for $i \in \{1, \dots, k\}$. Then H is an induced topological minor of G if and only if H' is an anchored topological minor of G for some choice of the vertices u_1, \dots, u_k . We check the latter statement by applying the $n^{O(1)}$ time algorithm of Corollary 1. As we have $O(n^k)$ possible choices, we only have to do this $O(n^k)$ times.

If H has one or more isolated vertices, we do the following. First, we guess a set of independent vertices in G that correspond to the isolated vertices of H . If we cannot find such a set of isolated vertices in G , then we stop and return no . Otherwise we remove the guessed vertices from G together with their neighbours. We check for each choice, whether the resulting graph G' contains the graph H' obtained from H by removing its isolated vertices as an induced topological minor. Note that such a new instance of INDUCED TOPOLOGICAL MINOR is trivial if G' or H' has no vertices. Since we have $O(n^k)$ new graphs G' , this only adds a $O(n^k)$ factor to the running time. We conclude that the total running time is $n^{k+O(1)}$, as desired. \square

If H is a part of the input, or if we parameterize the problem by the size of H , then the problem becomes hard even for a subclass of AT-free graphs. A graph is *cobipartite* if its vertex set can be partitioned into two cliques. Hence, the class of cobipartite graphs forms a subclass of the class of AT-free graphs.

Theorem 2. *The INDUCED TOPOLOGICAL MINOR problem is NP-complete for cobipartite graphs, and W[1]-hard for cobipartite graphs parameterized by $|V_H|$.*




Fig. 6. Construction of the graphs G' and H in the proof of Theorem 2 for $k = 4$.

Proof. We prove $W[1]$ -hardness. This proof immediately implies NP-hardness, as we reduce from the well-known $W[1]$ -complete CLIQUE problem [11]. For a graph G and a parameter k , this problem is to test whether G has a clique of size k .

Let (G, k) be an instance of CLIQUE. We assume without loss of generality that $k \geq 5$ and that G has minimum degree at least 4. We construct a cobipartite graph G' as follows:

- create a copy of V_G and construct a clique U on these vertices;
- for each edge $uv \in E_G$, create a vertex e_{uv} adjacent to u and v ;
- construct the clique $W = \{e_{uv} \mid uv \in E_G\}$.

Now we construct the graph H :

- create a clique X of size k with vertices x_1, \dots, x_k ;
- create a clique Y with $\frac{1}{2}k(k-1)$ vertices $y_{ij}, 1 \leq i < j \leq k$;
- for each pair i, j with $1 \leq i < j \leq k$, make x_i and x_j adjacent to y_{ij} .

The construction is shown in Fig. 6. We prove that G has a clique of size k if and only if G' contains H as an induced topological minor.

First suppose that G has a clique $\{u_1, \dots, u_k\}$. Then the subgraph of G' induced by this clique and the set of vertices $\{e_{u_i u_j} \mid 1 \leq i < j \leq k\}$ is isomorphic to H , i.e., H is an induced subgraph of G' , and therefore an induced topological minor of G' .

Now suppose that G' contains H as an induced topological minor. Then G' has an induced subgraph H' that is a subdivision of H . We call the vertices of H' that are vertices of H *branch vertices* and all other vertices of H' *subdivision vertices*.

Let x'_i for $1 \leq i \leq k$ and y'_{ij} for $1 \leq i < j \leq k$ be the branch vertices of H' , where x'_i and y'_{ij} correspond to the vertices x_i and y_{ij} of the graph H respectively. Let $X' = \{x'_i \mid 1 \leq i \leq k\}$ and $Y' = \{y'_{ij} \mid 1 \leq i < j \leq k\}$.

Suppose that H' has a subdivision vertex z . Assume that $z \in U$. Since $d_{H'}(z) = 2$, H' contains at most two other vertices of U . Hence, W contains at least $k-2 \geq 3$ vertices of X' and at least one vertex $y'_{ij} \in Y'$. Then y'_{ij} is adjacent to at least three vertices of X' , a contradiction. For the case $z \in W$ we use the same arguments, and conclude that H' has no subdivision vertices.

Assume that there are vertices $y'_{ij} \in U$ and $y'_{pq} \in W$. Since $k \geq 5$, U or W contains at least three vertices of X' . Note that y'_{ij} is adjacent to at least three vertices of X' in the first case and that y'_{pq} has at least three such neighbours in the second, a contradiction. Therefore, either $Y' \subseteq U$ or $Y' \subseteq W$.

Now we prove that either $X' \subseteq U$ or $X' \subseteq W$. To obtain a contradiction, assume that there are $x'_i \in U$ and $x'_j \in W$. If $Y' \subseteq U$, then all vertices of Y' are adjacent to x'_i , but since $k > 2$, at least one vertex of Y' is not adjacent to x'_i , a contradiction. We get the same conclusion if $Y' \subseteq W$.

The sets X' and Y' clearly cannot be subsets of the same clique U or W . Moreover, because each vertex of W is adjacent to exactly two vertices of U and each vertex of X' is adjacent to at least four vertices of Y' , $X' \subseteq U$ and $Y' \subseteq W$. By the construction of G' and H , this means that X' is a clique of size k in G . This concludes the proof of Theorem 2. \square

5. Induced paths and trees through specified vertices

In this section, we consider the k -IN-A-TREE and k -IN-A-PATH problem. Recall that this is the problem of detecting an induced tree respectively path containing a set S of k specified vertices (called terminals as well).

We use the following structural property of AT-free graphs.

Lemma 21. *Let G be an AT-free graph. For two vertices u and v of G let P be an induced (u, v) -path of length at least 4 in G . Let G' be the subgraph obtained from G after removing the vertices of $N[V_P \setminus \{u, v\}]$. If G_1 and G_2 are connected components of G' containing a neighbour of u and v in G , respectively, then $G_1 \neq G_2$.*

Proof. For contradiction, suppose that $G_1 = G_2$. Let u' and v' denote the neighbours of u and v in G that belong to $G_1 = G_2$. As G_1 is connected, G_1 contains an induced (u', v') -path P' . As G_1 contains no vertices from $N[V_P \setminus \{u, v\}]$, we then find that P' contains no such vertices either. As P and P' are induced and P has length at least 4, this means that the union C of P and P' is an induced cycle on at least six vertices. The latter implies that C has three mutually non-adjacent vertices that form an asteroidal triple in G , a contradiction. \square

For k -IN-A-TREE, we need the following definitions and well-known observation. A tree T is a *caterpillar* if there are two leaves x and y such that all other vertices are either vertices of the (x, y) -path P or are adjacent to the vertices of P . We say that P is a *central path* of T .

Lemma 22. *If T is an induced tree in an AT-free graph G then T is a caterpillar.*

We are now ready to prove the results of this section.

Theorem 3. *The k -IN-A-TREE problem can be solved in polynomial time for AT-free graphs even if k is part of the input.*

Proof. Using Lemma 22, we look for an induced path P between two terminals such that V_P and S induce a caterpillar with central path P , that is, all terminal vertices that are not in P are pairwise non-adjacent and each has exactly one neighbour in P .

We build this caterpillar step by step, using a dynamic programming algorithm that is somewhat similar to the algorithm used in the subroutine `Component(i, X, Y)` in Section 3.3.1. We use similar notations as well. The idea will be to “trace” the central path through the graph. However, in contrast to the subroutine `Component(i, X, Y)`, we have no “template” to go by and need to build the central path and the caterpillar step by step.

Before we show how the dynamic programming table is constructed and updated, we first give some additional terminology. For an ordered set (a sequence of distinct elements) $R = \{r_1, \dots, r_c\}$, we say that r_1 is the *first* element of R and r_c is the *last* element. The first and last elements of R are denoted `first(R)` and `last(R)` respectively. We also say that an ordered set of vertices $R = \{r_1, \dots, r_c\}$ of a graph G induces a path if $r_1 \dots r_c$ is an induced path in G . Let P be an induced (u, v) -path P in a graph G and let c be a positive integer. By $P(c)$ we denote the ordered set of the last c vertices of P in the path order. If P has length at most c , then $P(c)$ is the ordered set of the vertices of P . Observe that `last($P(c)$) = v`.

For an ordered set of vertices R that induces a path with end-vertices `first(R)` and `last(R)`, $A(R)$ denotes the set of terminals in the connected component of $G[(V_G \setminus N_G[V_R \setminus \{\text{first}(R), \text{last}(R)\}]) \cup \{\text{last}(R)\}]$ that contains `last(R)`. We can think of $A(R)$ as the terminals that are ahead of R .

Our dynamic programming algorithm creates tables that store entries with the following fields:

- a non-negative integer ℓ ;
- a terminal $x \in S$ (the starting terminal of the caterpillar);
- a non-negative integer r (the number of terminals in the caterpillar as constructed so far);
- a vertex z of G (the last vertex of the central path of the caterpillar as constructed so far; note that z might be but is not necessarily equal to x); and
- an ordered set R of at most five vertices of G with $z = \text{last}(R)$ that induces a path,

such that there is an induced (x, z) -path P in G and a partition A, B, N of S (some sets can be empty) with the following properties:

- $P(5) = R$;
- $N = S \cap (N_G[V_R \setminus \{\text{first}(R), z\}])$;
- $|B \cup N| = r$;
- $G[V_P \cup B \cup N]$ is a caterpillar; and
- for $A = A(R) \setminus \{z\}$, it holds that $A \cap V_P = \emptyset$ and vertices of A are not adjacent to the vertices of $(V_P \cup N) \setminus \{z\}$.

We can think of A as the set of terminals that are still ahead, of B as the set of terminals that are behind, and of N as the set of terminals that are currently in the neighbourhood of R . Also note that since $z \in R$, it would not be necessary to make z a separate field; however, we will do so for notational convenience.

The tables are constructed consecutively for $\ell = 0, \dots, |V_G| - 1$. We stop if one of the following occurs:

- the table contains a entry with $r = k$, and we return `yes` in this case; or
- we cannot construct any entry for the next value of ℓ , and we return `no` in this case.

The table for $\ell = 0$. This table only contains entries $(0, 0, x, z, R)$, where $x \in S$, $z = x$, and $R = \{x\}$.

The tables for $\ell > 0$. Assume that the table for $\ell - 1$ has been constructed. We consider all vertices $v \in V_G$ for which the table for $\ell - 1$ contains an entry $(\ell - 1, x, r, w, R)$ such that v is adjacent to w but is not adjacent to any other vertices of R . Then, for each v and $(\ell - 1, x, r, w, R)$, we create new entries in the table for ℓ . In order to do this, we construct the sets $A = A(R) \setminus \{w\}$ and $N = S \cap (N_G[V_R \setminus \{\text{first}(R), \text{last}(R)\}])$. Then we construct R' by adding v to R as the last element. Let $U = (N_G(w) \cap A) \setminus \{v\}$. Then we check whether the following holds:

- each vertex of U is adjacent to exactly one vertex of R' (i.e. only to the vertex w);
- each vertex of U is not adjacent to the vertices of $N \setminus \{w\}$;
- for the set $A' = A \setminus (\{v\} \cup U)$, it holds that $A' = A(R') \setminus \{v\}$.

The first two verify that we are indeed inducing a caterpillar, whereas the last verifies that we are expanding the caterpillar in the correct direction (towards the remaining terminals that are ahead). If these conditions are fulfilled, then we set $R'' = R' \setminus \{\text{first}(R)\}$ if $|R| = 5$, and $R'' = R'$ otherwise. Let $r' = r + |A| - |A'|$. Then we include the entry (ℓ, x, r', v, R'') in the table.

Observe that by Lemma 21, any vertex of G adjacent to $z = \text{last}(R)$ but non-adjacent to other vertices of R , is not adjacent to the vertices of the (x, z) -path P , except the last vertex z , and is not adjacent to the vertices of B . Then the correctness of our algorithm follows from the description and this observation.

It remains to prove that the algorithm is polynomial. Let $n = |V_G|$. Since $r \leq k$, there are k possibilities for the terminal x and at most n^5 ways to choose z and R . Hence, the total size of the table for any ℓ is at most $k^2 \cdot n^5$. Since $\ell \leq n - 1$, the running time is polynomial. \square

Theorem 4. *The k -IN-A-PATH problem can be solved in polynomial time for AT-free graphs even if k is part of the input.*

Proof. We use a dynamic programming algorithm that is constructed along the same lines as for the k -IN-A-TREE problem.

We have to find an induced path P between two terminals such that $S \subseteq V_P$. As in the proof of Theorem 3, for an ordered set of vertices R that induces a path with the end-vertices $\text{first}(R)$ and $\text{last}(R)$, $A(R)$ is the set of terminals in the connected component $G[(V_G \setminus N_G[V_R \setminus \{\text{first}(R), \text{last}(R)\}]) \cup \{\text{last}(R)\}]$ that contains $\text{last}(R)$.

Our dynamic programming algorithm creates tables that store entries with the following fields:

- a non-negative integer ℓ ;
- a terminal $x \in S$ (the starting terminal of the path);
- a non-negative integer r (the number of terminals on the path as constructed so far);
- a vertex z of G (the last vertex of the path as constructed so far; note that z might be but is not necessarily equal to x); and
- an ordered set R of at most five vertices of G with $z = \text{last}(R)$ that induces a path,

such that there is an induced (x, z) -path P in G , and a partition A, B, N of S (some sets can be empty) with the following properties:

- $P(5) = R$;
- $N = S \cap R$;
- $B \subseteq V_P \setminus R$;
- $|B \cup N| = r$; and
- for $A = A(R) \setminus \{z\}$, it holds that $A \cap V_P = \emptyset$ and vertices of A are not adjacent to the vertices of $V_P \setminus \{z\}$.

The intuition for A, B , and N are similar as before.

The tables are constructed consecutively for $\ell = 0, \dots, |V_G| - 1$. We stop if one of the following occurs:

- the table contains a entry with $r = k$, and we return `yes` in this case;
- we cannot construct any entry for the next value of ℓ , and we return `no` in this case.

The tables are constructed and updated as follows.

The table for $\ell = 0$. This table only contains entries $(0, 0, u, z, R)$, where $u \in S$, $z = u$, and $R = \{u\}$.

The tables for $\ell > 0$. Assume that the table for $\ell - 1$ has been constructed. We consider all vertices $v \in V_G$ for which the table for $\ell - 1$ contains a entry $(\ell - 1, x, r, w, R)$ such that v is adjacent to w but is not adjacent to any other vertices of R . Then, for each v and $(\ell - 1, x, r, w, R)$, we create new entries for the table for ℓ . In order to do this, we construct the set $A = A(R) \setminus \{w\}$. Then we construct R' by adding v to R as the last element. Let $A' = A \setminus \{v\}$. Then we check whether $A' = A(R') \setminus \{v\}$. This verifies that we are expanding the path in the correct direction (towards the remaining terminals that are ahead). If this holds, then we set $R'' = R' \setminus \{\text{first}(R)\}$ if $|R| = 5$, and $R'' = R'$ otherwise. Let $r' = r + |A| - |A'|$. Then we include the entry (ℓ, x, r', v, R'') in the table.

Observe that by Lemma 21, any vertex of G adjacent to $z = \text{last}(R)$ but non-adjacent to other vertices of R , is not adjacent to the vertices of the (x, z) -path P , except the last vertex z . Then the correctness of our algorithm follows from the description and the preceding observation.

It remains to prove that the algorithm is polynomial. Let $n = |V_G|$. Since $r \leq k$, there are k possibilities for the terminal x and at most n^5 ways to choose z and R . Hence, the total size of the table for any ℓ is at most $k^2 \cdot n^5$. Since $\ell \leq n - 1$, the running time is polynomial. \square

6. Concluding remarks

We have presented a polynomial-time algorithm that solves INDUCED DISJOINT PATHS for AT-free graphs, and we used this algorithm for a polynomial-time algorithm that solves H -INDUCED TOPOLOGICAL MINOR on this graph class for every fixed graph H . We complemented the latter result by proving that INDUCED TOPOLOGICAL MINOR, restricted to cobipartite graphs, is NP-complete and W[1]-hard when parameterized by $|V_H|$. We also showed that the problems k -IN-A-TREE and k -IN-A-PATH can be solved in polynomial time for AT-free graphs even if k is part of the input.

6.1. Open problem 1

Motivated by our application on testing for induced topological minors, we assumed that all terminal pairs in an instance (G, S) of INDUCED DISJOINT PATHS are distinct. For general graphs, we can easily drop this assumption by replacing a vertex u representing $\ell \geq 2$ terminals by ℓ new mutually non-adjacent vertices, each connected to all neighbours of u via subdivided edges. This yields an equivalent instance of INDUCED DISJOINT PATHS, in which all terminal pairs are distinct. However, we cannot apply this reduction for AT-free graphs, because it may create asteroidal triples. Hence, the complexity of this more general problem remains an open question.

We note that the special case in which all terminal pairs coincide, that is, in which $(s_1, t_1) = \dots = (s_k, t_k)$ is already NP-complete for $k = 2$ [2,12] and solvable in $O(n^2)$ time on n -vertex planar graphs for arbitrary k [36]. We can solve this case in polynomial time for AT-free graphs as follows.

Let G be an AT-free graph with k terminal pairs that are copies of the same terminal pair (s, t) . If $k = 1$, then the problem is trivial: we have to find an (s, t) -path in G . Hence, we assume that $k \geq 2$. Suppose that (s, t) -paths P_1, \dots, P_k form a solution. Then each path P_i has length at most 4. Otherwise, if there is a path P_i of length at least 5, this path and any other path P_j would induce a cycle with at least six vertices, but an AT-free graph has no induced cycles of length at least 6. Moreover, at most two paths P_i, P_j have length at least 3. To obtain a contradiction, assume that three paths P_i, P_j, P_r have length at least 3. Let u, v, w be the vertices adjacent to t in P_i, P_j, P_r respectively. Then u, v, w form an asteroidal triple, a contradiction. By these observations, we consider cases when a solution has $r = 0, 1, 2$ paths of length at least 3 and at most 4. For each r , we guess these paths by brute force. Then we construct the graph G' by the deletion of all internal vertices of these paths, together with their neighbours different from s, t . Now we have to find k' (s, t) -paths in G' of length 2, where $k' = k - r$ if s, t are non-adjacent, and $k' = k - r - 1$ if s, t are neighbours. These paths exist if and only if the graph $H = G'[N_{G'}(s) \cap N_{G'}(t)]$ has an independent set of size at least k' . It remains to observe that H is an AT-free graph, and that the INDEPENDENT SET problem can be solved in polynomial time for AT-free graphs [3].

6.2. Open problem 2

Our three algorithms for solving INDUCED DISJOINT PATHS, k -IN-A-TREE and k -IN-A-PATH on AT-free graphs are polynomial-time algorithms, but the polynomials in the running time bounds have high degree. It may be interesting to construct more efficient algorithms even for subclasses of the AT-free graphs, such as the class of permutation graphs or the class of cocomparability graphs. In contrast to the class of AT-free graphs, both these graph classes are known to have a geometric intersection model.

6.3. Open problem 3

Our last open problem is whether it is possible to extend our result for INDUCED DISJOINT PATHS to some superclass of AT-free graphs, such as graphs of bounded asteroidal number. An *asteroidal set* in a graph G is an independent set $S \subseteq V_G$, such that every triple of vertices of S is asteroidal. The *asteroidal number* introduced by Walter [41] (see also Kloks, Kratsch and Müller [22]) is the size of a largest asteroidal set in G . Note that complete graphs are exactly those graphs that have asteroidal number at most 1, and that AT-free graphs are exactly those graphs that have asteroidal number at most 2. There exist problems for which polynomial-time algorithms for AT-free graphs could be extended to polynomial-time algorithms for graphs of bounded asteroidal number (see, e.g., [3,27,28]). Is this also possible for INDUCED DISJOINT PATHS, or will this problem become NP-complete for some constant value of the asteroidal number?

CRedit authorship contribution statement

We confirm that all three authors contributed to the manuscript.

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.

References

- [1] 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 (2014) 501–521.
- [2] D. Bienstock, On the complexity of testing for odd holes and induced odd paths, *Discrete Math.* 90 (1991) 85–92, see also Corrigendum, *Discrete Math.* 102 (1992) 109.
- [3] H.J. Broersma, T. Kloks, D. Kratsch, H. Müller, Independent sets in asteroidal triple-free graphs, *SIAM J. Discrete Math.* 12 (1999) 276–287.
- [4] K. Cameron, Induced matchings, *Discrete Appl. Math.* 24 (1989) 97–102.
- [5] M. Chudnovsky, P.D. Seymour, The structure of claw-free graphs, in: *London Mathematical Society Lecture Notes Series*, vol. 327, 2005, pp. 153–171.
- [6] M. Chudnovsky, P.D. Seymour, The three-in-a-tree problem, *Combinatorica* 30 (2010) 387–417.
- [7] N. Derhy, C. Picouleau, Finding induced trees, *Discrete Appl. Math.* 157 (2009) 3552–3557.
- [8] N. Dehry, C. Picouleau, N. Trotignon, The four-in-a-tree problem in triangle-free graphs, *Graphs Comb.* 25 (2009) 489–502.
- [9] D.G. Corneil, S. Olariu, L. Stewart, Asteroidal triple-free graphs, *SIAM J. Discrete Math.* 10 (1997) 299–430.
- [10] D.G. Corneil, S. Olariu, L. Stewart, Linear time algorithms for dominating pairs in asteroidal triple-free graphs, *SIAM J. Comput.* 28 (1999) 1284–1297.
- [11] R.G. Downey, M.R. Fellows, *Parameterized Complexity*, Monographs in Computer Science, Springer-Verlag, New York, 1999.
- [12] 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.
- [13] J. Fiala, M. Kamiński, B. Lidický, D. Paulusma, The k -in-a-path problem for claw-free graphs, *Algorithmica* 62 (2012) 499–519.
- [14] P.A. Golovach, D. Paulusma, E.J. van Leeuwen, Induced disjoint paths in AT-free graphs, in: *Proc. SWAT 2012*, in: *Lecture Notes in Computer Science*, vol. 7357, 2012, pp. 153–164.
- [15] 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.
- [16] P.A. Golovach, D. Paulusma, E.J. van Leeuwen, Induced disjoint paths in claw-free graphs, *SIAM J. Discrete Math.* 29 (2015) 348–375.
- [17] G.C.M. Gomes, V.F. dos Santos, M.V.G. da Silva, J.L. Szwarcfiter, FPT and kernelization algorithms for the induced tree problem, in: *Proc. CIAC 2021*, in: *Lecture Notes in Computer Science*, vol. 12701, 2021, pp. 158–172.
- [18] M. Grohe, K. Kawarabayashi, D. Marx, P. Wollan, Finding topological subgraphs is fixed-parameter tractable, in: *Proc. STOC 2011*, 2011, pp. 479–488.
- [19] P. Heggernes, P. van 't Hof, R. Saei, E.J. van Leeuwen, Finding disjoint paths in split graphs, *Theory Comput. Syst.* 57 (2015) 140–159.
- [20] L. Jaffke, O. Kwon, J.A. Telle, Mim-width I. Induced path problems, *Discrete Appl. Math.* 278 (2020) 153–168.
- [21] R.M. Karp, On the complexity of combinatorial problems, *Networks* 5 (1975) 45–68.
- [22] T. Kloks, D. Kratsch, H. Müller, On the structure of graphs with bounded asteroidal number, *Graphs Comb.* 17 (2001) 295–306.
- [23] K. Kawarabayashi, Y. Kobayashi, B.A. Reed, The disjoint paths problem in quadratic time, *J. Comb. Theory, Ser. B* 102 (2012) 424–435.
- [24] 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.
- [25] Y. Kobayashi, K. Kawarabayashi, A linear time algorithm for the induced disjoint paths problem in planar graphs, *J. Comput. Syst. Sci.* 78 (2012) 670–680.
- [26] D. Kratsch, Domination and total domination on asteroidal triple-free graphs, *Discrete Appl. Math.* 99 (2000) 111–123.
- [27] D. Kratsch, H. Müller, Colouring AT-free graphs, in: *Proc. ESA 2012*, in: *Lecture Notes in Computer Science*, vol. 7501, 2012, pp. 707–718.
- [28] D. Kratsch, H. Müller, I. Todinca, Feedback vertex set and longest induced path on AT-free graphs, in: *Proc. WG 2003*, in: *Lecture Notes in Computer Science*, vol. 2880, 2003, pp. 309–321.
- [29] K.-Y. Lai, H.-I. Lu, M. Thorup, Three-in-a-tree in near linear time, in: *Proc. STOC 2020*, 2020, pp. 1279–1292.
- [30] N.K. Le, Detecting induced subdivision of K_4 , *J. Graph Theory* 90 (2019) 160–171.
- [31] C.G. Lekkerkerker, J.Ch. Boland, Representation of a finite graph by a set of intervals on the real line, *Fundam. Math.* 51 (1962) 45–64.
- [32] B. Lévêque, D.Y. Lin, F. Maffray, N. Trotignon, Detecting induced subgraphs, *Discrete Appl. Math.* 157 (2009) 3540–3551.
- [33] B. Lévêque, F. Maffray, N. Trotignon, On graphs with no induced subdivision of K_4 , *J. Comb. Theory, Ser. B* 102 (2012) 924–947.
- [34] W. Liu, N. Trotignon, The k -in-a-tree problem for graphs of girth at least k , *Discrete Appl. Math.* 158 (2010) 1644–1649.
- [35] J.F. Lynch, The equivalence of theorem proving and the interconnection problem, *SIGDA News.* 5 (1975) 31–36.
- [36] C.J.H. McDiarmid, B.A. Reed, A. Schrijver, F.B. Shepherd, Induced circuits in planar graphs, *J. Comb. Theory, Ser. B* 60 (1994) 169–176.
- [37] S. Natarajan, A.P. Sprague, Disjoint paths in circular arc graphs, *Nord. J. Comput.* 3 (1996) 256–270.
- [38] M. Radovanović, N. Trotignon, K. Vučković, The (theta, wheel)-free graphs part IV: induced paths and cycles, *J. Comb. Theory, Ser. B* 146 (2021) 495–531.
- [39] N. Robertson, P.D. Seymour, Graph minors XIII. The disjoint paths problem, *J. Comb. Theory, Ser. B* 63 (1995) 65–110.
- [40] J. Stacho, 3-Colouring AT-free graphs in polynomial time, *Algorithmica* 64 (2012) 384–399.
- [41] J.R. Walter, Representations of chordal graphs as subtrees of a tree, *J. Graph Theory* 2 (1978) 265–267.