Induced Disjoint Paths Without an Induced Minor

Pierre Aboulker ✉ 🏠 🆔
DIENS, École normale supérieure, CNRS, PSL University, Paris, France

Édouard Bonnet ✉ 🏠 🆔
CNRS, ENS de Lyon, Université Claude Bernard Lyon 1, LIP UMR 5668, Lyon, France

Timothé Picavet ✉ 🆔
LaBRI, Université de Bordeaux, France

Nicolas Trotignon ✉ 🏠 🆔
CNRS, ENS de Lyon, Université Claude Bernard Lyon 1, LIP UMR 5668, Lyon, France

──── **Abstract** ────

We exhibit a new obstacle to the nascent algorithmic theory for classes excluding an induced minor. We indeed show that on the class of string graphs – which avoids the 1-subdivision of, say, K_5 as an induced minor – INDUCED 2-DISJOINT PATHS is NP-complete. So, while k-DISJOINT PATHS, for a fixed k, is polynomial-time solvable in general graphs, the absence of a graph as an induced minor does not make its induced variant tractable, even for $k = 2$. This answers a question of Korhonen and Lokshtanov [SODA '24], and complements a polynomial-time algorithm for INDUCED k-DISJOINT PATHS in classes of bounded genus by Kobayashi and Kawarabayashi [SODA '09]. In addition to being string graphs, our produced hard instances are subgraphs of a constant power of bounded-degree planar graphs, hence have bounded twin-width and bounded maximum degree.

We also leverage our new result to show that there is a fixed subcubic graph H such that deciding if an input graph contains H as an induced subdivision is NP-complete. Until now, all the graphs H for which such a statement was known had a vertex of degree at least 4. This answers a question by Chudnovsky, Seymour, and Trotignon [JCTB '13], and by Le [JGT '19]. Finally we resolve another question of Korhonen and Lokshtanov by exhibiting a subcubic graph H without two adjacent degree-3 vertices and such that deciding if an input n-vertex graph contains H as an induced minor is NP-complete, and unless the Exponential-Time Hypothesis fails, requires time $2^{\Omega(\sqrt{n})}$. This complements an algorithm running in subexponential time $2^{\tilde{O}(n^{2/3})}$ by these authors [SODA '24] under the same technical condition.

1 Introduction

In k-DISJOINT PATHS, one is given a graph G, together with k pairs of vertices, often called *terminals*, $(s_1, t_1), \ldots, (s_k, t_k)$, and has to decide if G admits k vertex-disjoint paths $P_1, \ldots, P_k$ such that for every $i \in \{1, \ldots, k\}$, the endpoints of P_i are s_i and t_i. This problem

is also called k-LINKAGE as the pairs to connect are prescribed. In the "flow variant" of this problem, DISJOINT S–T PATHS, the input consists of a graph G and two disjoint vertex subsets $S, T \subset V(G)$ with $|S| = |T|$, and the question is whether there are $|S|$ vertex-disjoint paths, each with one endpoint in S and the other endpoint in T. These problems are polynomial-time solvable in general graphs, respectively by the work of Robertson and Seymour [27] (for k-DISJOINT PATHS), and simply by Menger's theorem [24] (for DISJOINT S–T PATHS).

In this paper, we are interested in their induced variants INDUCED k-DISJOINT PATHS and INDUCED DISJOINT S–T PATHS, where the paths are further requested to be mutually induced (i.e., no edge of the graph is incident to two of these paths). These problems are NP-complete for $k = 2$ and for $|S| = |T| = 2$, respectively, in general graphs [8, 1]. Note that INDUCED DISJOINT S–T PATHS with $|S| = |T| = k$ can be solved with $k!$ calls to INDUCED k-DISJOINT PATHS. So when k is constant, the former problem admits a polynomial-time Turing reduction to the latter, and thus cannot be harder.

Kawarabayashi and Kobayashi give a linear-time algorithm for INDUCED k-DISJOINT PATHS (for any fixed k) in planar graphs [18], and a polynomial-time algorithm in classes of bounded genus [19]. The latter result is lifted to the model checking of the FO+SDP logic (i.e., first-order with *scattered* disjoint paths predicates, which can natively express the existence of a constant number of induced disjoint paths between some specified pairs of terminals) in fixed-parameter tractable (FPT) time in classes of bounded genus [13]. It is believed that the tractability of INDUCED k-DISJOINT PATHS (and perhaps even of FO+SDP model checking) holds more generally in classes excluding a fixed minor, and could be shown by combining the irrelevant-vertex techniques (developed for the planar and bounded-genus cases) with the graph structure theorem of Robertson and Seymour [28]. However, to our knowledge, this has not been proven yet (provided it indeed holds).

INDUCED k-DISJOINT PATHS is polynomial-time solvable on classes of bounded mim-width on which mim-width can be efficiently approximated [17], which includes for instance interval graphs and permutation graphs. In claw-free graphs (i.e., graphs excluding $K_{1,3}$ as an induced subgraph), INDUCED k-DISJOINT PATHS is solvable in polynomial time for any fixed k [10], and even in FPT time in parameter k [11]. In graphs without asteroidal triple, this problem is polynomial-time solvable even if k is part of the input [12]. Finally, in (theta, wheel)-free graphs, there is a polynomial-time algorithm for INDUCED k-DISJOINT PATHS [26], while its complexity in theta-free graphs is open.

To understand better the tractability frontier of INDUCED k-DISJOINT PATHS, Korhonen and Lokshtanov [20] ask if this problem is NP-hard in H-induced-minor-free graphs for some fixed k and H. We resolve this question already in the case when $k = 2$ and H is the 1-subdivision of K_5 (or of $K_{3,3}$). Indeed string graphs (i.e., intersection graphs of curves in the plane) exclude the 1-subdivision of any non-planar graph as an induced minor.

▶ **Theorem 1.** *INDUCED 2-DISJOINT PATHS is NP-complete in string graphs that are subgraphs of a constant power of bounded-degree planar graphs.*

We actually show the stronger result that INDUCED DISJOINT S–T PATHS with $|S| = |T| = 2$ is NP-complete in this subclass of string graphs. Note that subgraphs of constant powers of bounded-degree planar graphs both have bounded twin-width [3, 2] and bounded maximum degree. Thus Theorem 1 considerably limits the scope within which the existing polynomial algorithms could be extended.

The Exponential-Time Hypothesis (ETH for short), a stronger assumption than $P \neq NP$ but still widely believed, asserts that there is a real $\lambda > 1$ such that n-variable 3-SAT cannot be solved in time $O(\lambda^n)$ [15]. More quantitatively, the proof of Theorem 1, by providing a linear reduction from a variant of PLANAR 3-SAT (see [23]), implies the following.

▶ **Theorem 2.** INDUCED DISJOINT S–T PATHS *with* $|S| = |T| = 2$ *is NP-complete in string graphs of bounded maximum degree and twin-width, and requires time* $2^{\Omega(\sqrt{n})}$ *on n-vertex such graphs, unless the Exponential-Time Hypothesis fails.*

Our result has some consequences for the detection of induced subdivisions, which we now turn our attention to.

Detecting induced subdivisions and induced minors

Let H-INDUCED SUBDIVISION CONTAINMENT (H-ISC for short) input a graph G and ask whether H is an induced subdivision of G. This problem has attracted some attention. Chudnovsky and Seymour introduced the THREE-IN-A-TREE problem – whether there is an induced subtree containing three given vertices – and showed how to solve it in polynomial time via the so-called extended strip decompositions, in order to obtain a polynomial algorithm for $K_{2,3}$-ISC [4]. The first graphs H for which H-ISC is NP-complete came from [22] where several examples are given: notably the complete graph K_5 and some trees, among other graphs. In the same paper, some other examples of tractable H-ISC were given, all relying on the polynomial-time algorithm for THREE-IN-A-TREE. There are also some ad hoc algorithms for H-ISC when H is the *net* (i.e., the graph obtained by adding a pendant neighbor to each vertex of a triangle) [5], when H is K_4 [21], or when H is the disjoint union of a fixed number of triangles [25].

Despite this line of work, no subcubic graph H was known to make H-ISC NP-hard. Actually, Chudnovsky, Seymour, and Trotignon [5] and Le [21] asked whether there is a polynomial-time algorithm for H-ISC for any subcubic graph H. As a consequence of Theorem 1, we answer this question by the negative by exhibiting a subcubic graph H (the graph of Figure 1) for which H-ISC is NP-hard.
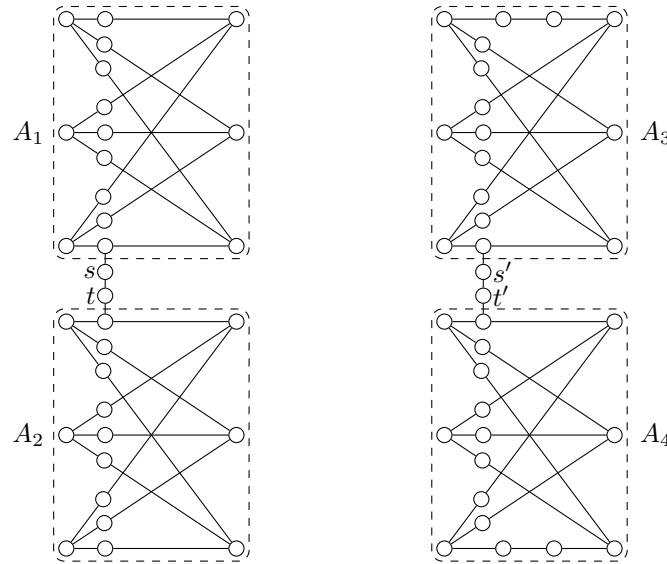
Figure 1 The subcubic graph H. Both $H[A_1]$ and $H[A_2]$ are the 1-subdivision of $K_{3,3}$. Both $H[A_3]$ and $H[A_4]$ are obtained from $K_{3,3}$ by subdividing every edge once but one edge that is subdivided twice. The vertices of H which are not in $\bigcup_{i \in [4]} A_i$ are labeled s, t, s', t'. In total H has 66 vertices.

▶ **Theorem 3.** H-INDUCED SUBDIVISION CONTAINMENT *is NP-complete for the subcubic graph H of Figure 1, and requires time* $2^{\Omega(\sqrt{n})}$ *on n-vertex graphs, unless the ETH fails.*

Theorem 3 holds within subgraphs of a constant power of bounded-degree planar graphs that are string graphs plus a constant number of (bounded-degree) apices, which are graphs of bounded twin-width [3, 2] excluding a fixed induced minor. In other words, Theorem 3 holds in the graphs for which Theorem 1 holds augmented by a constant number of vertices of bounded degree.

Theorem 1 also has consequences for the detection of induced minors, topic that we now briefly survey. Let H-INDUCED MINOR CONTAINMENT (H-IMC for short) input a graph G and ask whether H is an induced minor of G. It was first shown by Fellows et al. [9] that H-IMC can be NP-hard for a fixed graph H (unlike the minor containment). The latter graph H was not planar, which prompted the authors to ask if there is a planar graph H for which H-IMC is NP-hard, and whether H-IMC is always polynomial-time solvable when H is a tree. Recently, Korhonen and Lokshtanov [20] answered both questions by showing that H-IMC is NP-hard for some fixed tree H (whose number of vertices is not explicit, and estimated to be larger than 2^{300} in [7]). Dallard et al. [6] show that $K_{2,3}$-IMC can be solved in polynomial time. Finding the disjoint union of a constant number of triangles as an induced minor or as an induced subdivision is all the same, so for any natural t, tK_3-IMC is also in P [25]. For a more complete survey on the complexity of H-IMC, we refer the reader to the introduction of [7]; the paper provides some polynomial-time algorithms for three infinite families of graphs H, and notes that H-IMC is in P for every graph H on at most 5 vertices, except for three remaining open cases.

As a similar consequence to Theorem 2, we show a $2^{\Omega(\sqrt{n})}$ lower bound for H-IMC on n-vertex graphs, under the ETH. This can be put in perspective with a $2^{\tilde{O}(n^{2/3})}$-time algorithm for H-IMC when every edge of H is incident to a vertex of degree at most 2 [20]. The authors ask if the mere NP-hardness of H-IMC can be shown for some graph H with this property. Actually by subdividing eight edges in the graph of Figure 1, we obtain a graph H' satisfying the property, and for which we can show the same lower bound.

▶ **Theorem 4.** *There is a subcubic graph H' such that every edge of H' is incident to a vertex of degree 2, and H'-INDUCED MINOR CONTAINMENT is NP-complete, and requires time $2^{\Omega(\sqrt{n})}$ on n-vertex graphs, unless the ETH fails.*

We observe that it is not too difficult to show Theorems 3 and 4 with *connected* graphs H and H' (having the extra properties of their respective statement). As this complicates a bit their proofs without being a significantly stronger result, we opted against doing it explicitly.

Open questions

We suggest the following open questions, which come more or less directly from our work and the literature. In light of the surveyed polynomial algorithms applying more generally to INDUCED k-DISJOINT PATHS and the hardness proofs applying more generally to INDUCED DISJOINT S–T PATHS with $|S| = |T| = k$, we wonder if there is a hereditary graph class in which INDUCED k-DISJOINT PATHS is NP-complete but INDUCED DISJOINT S–T PATHS with $|S| = |T| = k$ is polynomial-time solvable; the case $k = 2$ is of particular interest.

The string graphs that our proof of Theorem 1 produces do not seem to be 1-string graphs (i.e., realizable with strings every pair of which intersects at most once). We leave open the complexity of INDUCED k-DISJOINT PATHS in 1-string graphs, and also in segment intersection graphs (a further restriction to 1-string graphs).

A natural conjecture stems from our paper and existing algorithms, under P $\neq$ NP.

▶ **Conjecture 5.** *For any subcubic graph H, H-ISC is in P if and only if H is planar.*

2 Preliminaries

If i is a positive integer, we denote by $[i]$ the set of integers $\{1, 2, \ldots, i\}$.

2.1 Subgraphs, induced subgraphs, neighborhoods, and some graphs

We denote by $V(G)$ and $E(G)$ the set of vertices and edges of a graph G, respectively. A graph H is a *subgraph* of a graph G if H can be obtained from G by vertex and edge deletions. Graph H is an *induced subgraph* of G if H is obtained from G by vertex deletions only. A graph G is H-*free* if G does not contain H as an induced subgraph. For $S \subseteq V(G)$, the *subgraph of G induced by S*, denoted $G[S]$, is obtained by removing from G all the vertices that are not in S (together with their incident edges). Then $G - S$ is a short-hand for $G[V(G) \setminus S]$. A collection of paths $P_1, \ldots, P_h$ in a graph G is said *mutually induced* if $G[\bigcup_{i \in [h]} V(P_i)]$ is exactly the disjoint union of paths $P_1, \ldots, P_h$.

A set $X \subseteq V(G)$ is connected (in G) if $G[X]$ has a single connected component. A vertex whose removal increases the number of connected components is called a *cutvertex*. Similarly, a *bridge* is an edge whose removal increases the number of connected components. We denote by G^r the *r-th power of G*, that is, the graph with vertex set $V(G)$ and an edge between any two vertices at (shortest-path) distance at most r in G. A *constant power* of G is G^r for some constant r.

A *subdivision* of a graph G is any graph obtained from G by replacing each edge of G by a path of at least one edge. In a subdivision of G the vertices of G are called *branching vertices*, and the created vertices are called *subdivision vertices*. The *s-subdivision* of G is the graph obtained from G by replacing each edge of G by a path of $s + 1$ edges.

We denote by $N_G(v)$ and $N_G[v]$, the open, respectively closed, neighborhood of v in G. For $S \subseteq V(G)$, we set $N_G(S) := (\bigcup_{v \in S} N_G(v)) \setminus S$ and $N_G[S] := N_G(S) \cup S$. The *degree* $d_G(v)$ of a vertex $v \in V(G)$ is the cardinality of $N_G(v)$, and the *maximum degree* of G is defined as $\max_{v \in V(G)} d_G(v)$. A *subcubic graph* is a graph of maximum degree at most 3.

The *t-clique*, denoted by K_t, is obtained by making adjacent every pair of two distinct vertices among t vertices, and the *biclique* $K_{t,t}$ with bipartition (A, B) such that $|A| = |B| = t$ is obtained by making every vertex of A adjacent to every vertex of B. A *theta* is any subdivision of $K_{2,3}$. A *wheel* is any graph obtained by adding to a cycle of length at least 4, a vertex with at least three neighbors on the cycle.

A *string graph* is the intersection graph of some collection of (non-self-intersecting) curves in the plane (usually called strings), or equivalently the intersection graph of a collection of connected sets of some planar graph. The collection of strings is called *string representation*. We may see the string representation as a planar diagram with one vertex at each string endpoint and at each intersection of two strings. For instance, any string representation defines an infinite face (the infinite face of this planar diagram).

It is known (and easy to see) that the 1-subdivision of any non-planar graph is not a string graph [29].

2.2 Induced subdivisions and induced minors

A graph H is an induced subdivision of a graph G if a subdivision of H is isomorphic to an induced subgraph of G. An *induced subdivision model* of H in G is given by an injective map $\phi : V(H) \to V(G)$ and a collection of paths $(P_e)_{e \in E(H)}$ in G such that for every $uv \in E(H)$, P_{uv} is a $\phi(u)$–$\phi(v)$ path, and $G[\bigcup_{e \in E(H)} V(P_e)]$ has no more edges than the paths $(P_e)_{e \in E(H)}$ themselves.

An *induced minor model* of H in G is a collection $\mathcal{M} := \{X_1, \ldots, X_{|V(H)|}\}$ of pairwise-disjoint connected subsets of $V(G)$, called *branch sets*, together with a bijective map $\phi : V(H) \to \mathcal{M}$ such that $uv \in E(H)$ if and only if there is at least one edge in G between $\phi(u)$ and $\phi(v)$. In which case, we may say that the branch sets $\phi(u)$ and $\phi(v)$ are *adjacent*. We then say that H is an *induced minor* of G (or otherwise that G is H-induced-minor-free). Or equivalently, H can be obtained from G after a series of vertex deletions and edge contractions.

An induced minor model $(\{X_1, \ldots, X_h\}, \phi)$ of an h-vertex graph H is *minimal* if for every $X_1' \subseteq X_1, \ldots, X_h' \subseteq X_h$, the fact that $(\{X_1', \ldots, X_h'\}, \phi')$ is an induced minor model of H with $\phi'(u) = X_i' \Leftrightarrow \phi(u) = X_i$ for every $u \in V(H)$, implies that for every $i \in [h]$, $X_i' = X_i$.

With the second given definition of string graphs (see end of Section 2.1), it is easy to see that the class of string graphs is closed under taking induced minors. Thus no string graph admits the 1-subdivision of a non-planar graph as an induced minor.

2.3 Useful facts on twin-width

As we only mention twin-width in side remarks, we refrain from giving a definition. We list the theorems useful in Section 3. This can be read in a black-box fashion.

It was first proven in [3] that the class of planar graphs has bounded twin-width. The current best upper bound is 8 [14].

▶ **Theorem 6** (Theorem 6.3 in [3], [14]). *Planar graphs have bounded twin-width, more precisely upper bounded by 8.*

The constant powers of bounded twin-width graphs have bounded twin-width, as a special case of so-called first-order transductions.

▶ **Theorem 7** (Theorem 8.1 in [3]). *There is a function f such that for any graph G of twin-width d and for any positive integer r, G^r has twin-width at most $f(d, r)$.*

Among weakly sparse classes (excluding a $K_{t,t}$ as a subgraph), the subgraph closure of any class of bounded twin-width has bounded twin-width.

▶ **Theorem 8** ([2]). *There is a function g such that for any graph G of twin-width d excluding $K_{t,t}$ as a subgraph (or in particular, of maximum degree at most $t - 1$), then any subgraph of G has twin-width at most $g(d, t)$.*

3 Hardness of Induced 2-Disjoint Paths in string graphs

In this section, we show that INDUCED 2-DISJOINT PATHS is NP-complete in (a proper subclass of) string graphs.

▶ **Theorem 1.** INDUCED 2-DISJOINT PATHS *is NP-complete in string graphs that are subgraphs of a constant power of bounded-degree planar graphs.*

Proof. We reduce from the NP-complete problem CLAUSE-LINKED PLANAR E3-OCC 3-SAT, with variables $x_1, \ldots, x_n$ and clauses $c_1, \ldots, c_m$ where every clause c_j is on two or three variables, each variable appears in exactly three clauses, and the variable-clause incidence graph is planar even when the cycle $c_1 c_2 \ldots c_m c_1$ is added. This problem has been proven NP-complete by Fellows et al. [9], even in a restricted form with some additional constraint on the clauses, but we will not need this restriction. To get a unique representation of the variable gadget, we assume that every variable has exactly three occurrences: two positive

and one negative, or two negative and one positive. Indeed variables appearing only positively (or only negatively) can be set to true (or to false). We refer the reader to [30] for the complexity of variants of PLANAR 3-SAT.

We present the variable gadgets, clause gadgets, and variable-clause incidences, which assembled together form the graph G input to INDUCED 2-DISJOINT PATHS. The reader can also refer to Figure 2 where these three elements are depicted.

Variable gadget. For each variable x_i, we now describe the variable gadget $\mathcal{G}(x_i)$. For each clause c_j containing x_i, we add to $\mathcal{G}(x_i)$ two vertices $w_j^N(x_i), w_j^S(x_i)$ if x_i appears positively in c_j, and we add two vertices $w_j^N(\neg x_i), w_j^S(\neg x_i)$ if, instead, x_i appears negatively in c_j. We finally turn $\mathcal{G}(x_i)$ into an induced biclique by adding every edge between pairs of vertices $w_j^d(x_i), w_{j'}^{d'}(\neg x_i)$ for $d, d' \in \{N, S\}$ and j, j' may be equal or distinct. Observe that every variable gadget is isomorphic to the bipartite complete graph $K_{2,4}$.

Clause gadget. We describe the clause gadget $\mathcal{G}(c_j)$ for each clause c_j. Let ℓ_1, ℓ_2, ℓ_3 be the three literals of c_j (or ℓ_1, ℓ_2 if c_j is a 2-clause). The gadget $\mathcal{G}(c_j)$ has 16 vertices (or 12 if c_j is a 2-clause): four *entry points* $u_j^N, u_j^S, u_{j+1}^N, u_{j+1}^S$, and four vertices $v_j^{NW}(\ell_a), v_j^{NE}(\ell_a), v_j^{SW}(\ell_a), v_j^{SE}(\ell_a)$ for each $a \in \{1, 2, 3\}$ (for each $a \in \{1, 2\}$). We add an edge between a pair of vertices $v_j^d(\ell_a), v_j^{d'}(\ell_{a'})$, with $d, d' \in \{NW, NE, SW, SE\}$, whenever $a \neq a'$. Thus those 12 vertices (8 vertices) form a tripartite complete graph $K_{4,4,4}$ (bipartite complete graph $K_{4,4}$). For every $d \in \{N, S\}$ and $a \in \{1, 2, 3\}$ ($a \in \{1, 2\}$), we also add the six (four) edges $u_j^d v_j^{dW}(\ell_a)$, and the six (four) edges $v_j^{dE}(\ell_a)u_{j+1}^d$. This finishes the description of $\mathcal{G}(c_j)$. Note that $\mathcal{G}(c_j)$ and $\mathcal{G}(c_{j+1})$ share two vertices: u_{j+1}^N and u_{j+1}^S.

Variable-clause incidence. Finally for every clause c_j and literal ℓ in c_j, and each $d \in \{N, S\}$, we add the edges $v_j^{dW}(\ell)w_j^d(\ell)$ and $w_j^{dE}(\ell)v_j^d(\ell)$.

The graph G, input of INDUCED 2-DISJOINT PATHS, is obtained by having a gadget for each variable and each clause, and adding the edges as described in the previous sentence. We set the first terminal pair to u_1^N, u_{m+1}^N and the second terminal pair to u_1^S, u_{m+1}^S. This finishes the construction; see Figure 2. The next two lemmas show the correctness of the reduction.

▷ **Claim 9.** If $c_1 \wedge \ldots \wedge c_m$ is satisfiable, then G admits two u_1^N–u_{m+1}^N and u_1^S–u_{m+1}^S mutually induced paths.

Proof. We fix a truth assignment $\ell_1, \ldots, \ell_n$ of the variables satisfying $c_1 \wedge \ldots \wedge c_m$, with $\ell_i \in \{x_i, \neg x_i\}$ for every $i \in [n]$. For each clause c_j, we choose any literal $\ell(c_j)$ of c_j such that $\ell(c_j) = \ell_i$ for some $i \in [n]$. For each $d \in \{N, S\}$, we build the path P^d:

$$u_1^d, v_1^{dW}(\ell(c_1)), w_1^d(\ell(c_1)), v_1^{dE}(\ell(c_1)),\ u_2^d, v_2^{dW}(\ell(c_2)), w_2^d(\ell(c_2)), v_2^{dE}(\ell(c_2)),\ \ldots,$$

$$u_m^d, v_m^{dW}(\ell(c_m)), w_m^d(\ell(c_m)), v_m^{dE}(\ell(c_m)), u_{m+1}^d.$$

Note that P^N is a u_1^N–u_{m+1}^N path in G, P^S is a u_1^S–u_{m+1}^S path, and P^N, P^S are vertex-disjoint. We claim there is no chord between P^N and P^S. This is essentially because

- within every $\mathcal{G}(c_j)$, $V(P^N)$ and $V(P^S)$ induce a $4K_2$ comprising four edges of $P^N \uplus P^S$,
- apart from edges incident to entry points, there is no edge between two clause gadgets,
- P^N and P^S enter each variable gadget on the same side of the induced biclique,

which concludes the proof of the claim.   ◁

▷ **Claim 10.** If G admits two u_1^N–u_{m+1}^N and u_1^S–u_{m+1}^S mutually induced paths, then $c_1 \wedge \ldots \wedge c_m$ is satisfiable.
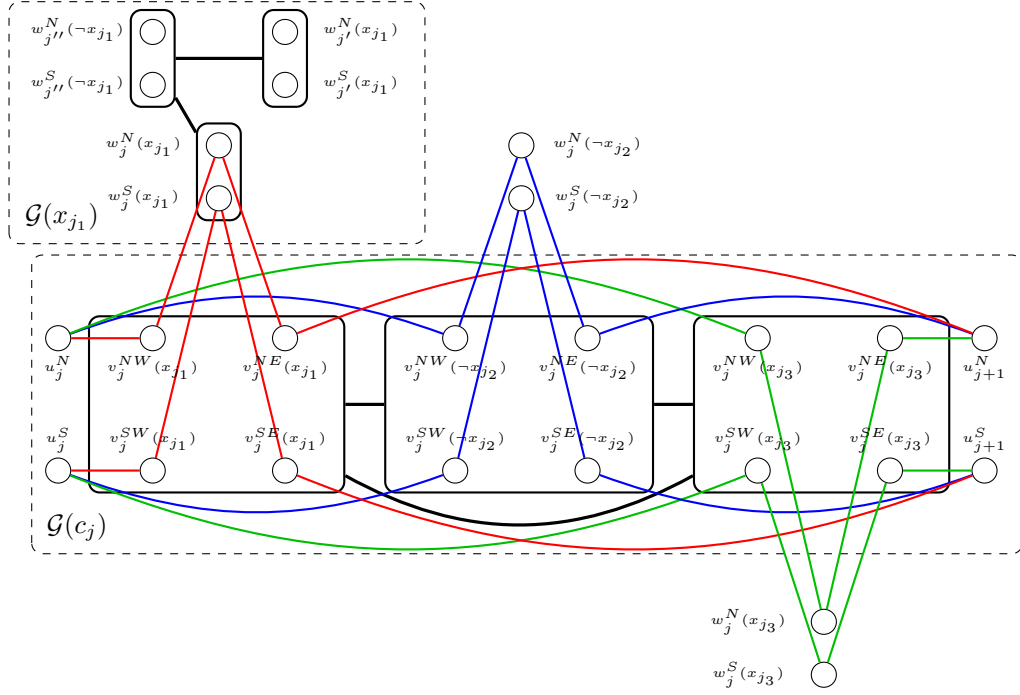
Figure 2 The clause gadget $\mathcal{G}(c_j)$ with $c_j = x_{j_1} \vee \neg x_{j_2} \vee x_{j_3}$, variables x_{j_1}, x_{j_2} in one face defined by cycle $c_1 c_2 \ldots c_m c_1$ (above), and x_{j_3} in the other face (below). We also drew the variable gadget $\mathcal{G}(x_{j_1})$ when x_{j_1} appears positively in c_j and $c_{j'}$, and negatively in $c_{j''}$. Edges linking two rounded boxes represent bicliques.

Proof. Let P^N, P^S be two mutually induced paths in G such that P^d is a u_1^d–u_{m+1}^d path for each $d \in \{N, S\}$. We think of P^d as going from u_1^d, its start, to u_{m+1}^d, its end. We first show the following invariant. For every $j \in [m]$, there exists a literal ℓ of c_j such that for each $d \in \{N, S\}$, the path P^d goes from u_j^d to u_{j+1}^d via the 4-edge subpath $u_j^d, v_j^{dW}(\ell), w_j^d(\ell), v_j^{dE}(\ell), u_{j+1}^d$.

Fix some $j \in [m]$, and assume that the invariant holds for every $j' \in [j-1]$ (with $[0] = \emptyset$). The vertex following u_j^N along P^N cannot be in $\mathcal{G}(c_{j-1})$ as either $j = 1$ (and this gadget does not exist) or it would create a chord between P^N and P^S, by the induction hypothesis. Thus it has to be $v_j^{NW}(\ell)$ for some literal ℓ of c_j. To avoid a chord between P^N and P^S, the vertex following u_j^S along P^S has to be $v_j^{SW}(\ell)$. Then the next vertex along P^N (resp. P^S) has to be $w_j^N(\ell)$ (resp. $w_j^S(\ell)$). As the variable gadget of literal ℓ is an induced biclique, the next vertices have to be back to the clause gadget $\mathcal{G}(c_j)$: $v_j^{NE}(\ell)$ along P^N, and $v_j^{SE}(\ell)$ along P^S. Finally the next vertices have to be u_{j+1}^N along P^N, and u_{j+1}^S along P^S. This completes the proof that the invariant holds for every $j \in [m]$.

Let $\mathcal{A}$ be the truth assignment setting x_i to true if P^N (resp. P^S) does not contain any vertex $w_j^N(\neg x_i)$ (resp. $w_j^S(\neg x_i)$), and to false otherwise. As P^N, P^S are mutually induced, in the latter case P^N (resp. P^S) does not contain any vertex $w_j^N(x_i)$ (resp. $w_j^S(x_i)$). Thus, for every clause c_j, the vertex $v_j^{NW}(\ell) \in P^N$ is such that literal ℓ is satisfied by $\mathcal{A}$. Hence $\mathcal{A}$ is a satisfying assignment.   ◁

Observe that the invariant in the first paragraph of the proof of Claim 10 still holds under the weaker assumption that the endpoints of P^N, P^S are in $\{u_{m+1}^N, u_{m+1}^S\}$ (but are unspecified). Therefore finding two mutually induced paths between $\{u_1^N, u_1^S\}$ and $\{u_{m+1}^N, u_{m+1}^S\}$ in G (the induced flow problem) is equivalent to the induced linkage problem.

We now show that G is as advertised by the theorem statement. We start by giving a string representation for G.

▷ **Claim 11.** G is a string graph.

Proof. First, we draw a planar embedding $\mathcal{P}$ of the variable-clause incidence graph augmented with the cycle $c_1 c_2 \ldots c_m c_1$. We refer to one face delimited by this cycle as the *upper face* (drawn in the figures "above" the path $c_1 c_2 \ldots c_m$), and the other face as the *lower face*. Second, for every clause c_j, draw $\mathcal{G}(c_j)$, and the six (or four, if c_j is a 2-clause) vertices in variable gadgets that are adjacent to $\mathcal{G}(c_j)$, as the intersection graph of strings such that

- the strings of $u_j^N, u_j^S, u_{j+1}^N, u_{j+1}^S$ protrude in the infinite face,[1]
- as well as $w_j^N(\ell)$ (resp. $w_j^S(\ell)$) if ℓ is a literal of c_j whose variable is drawn in the upper face (resp. lower face) in $\mathcal{P}$,
- these strings appear along the infinite face in the order prescribed by $\mathcal{P}$, and
- for each literal ℓ of c_j, there is a face that contains a substring of $w_j^N(\ell)$ and a substring of $w_j^S(\ell)$, one of which bounding the infinite face (as requested by the second item).

Such a string representation is given in Figure 3.
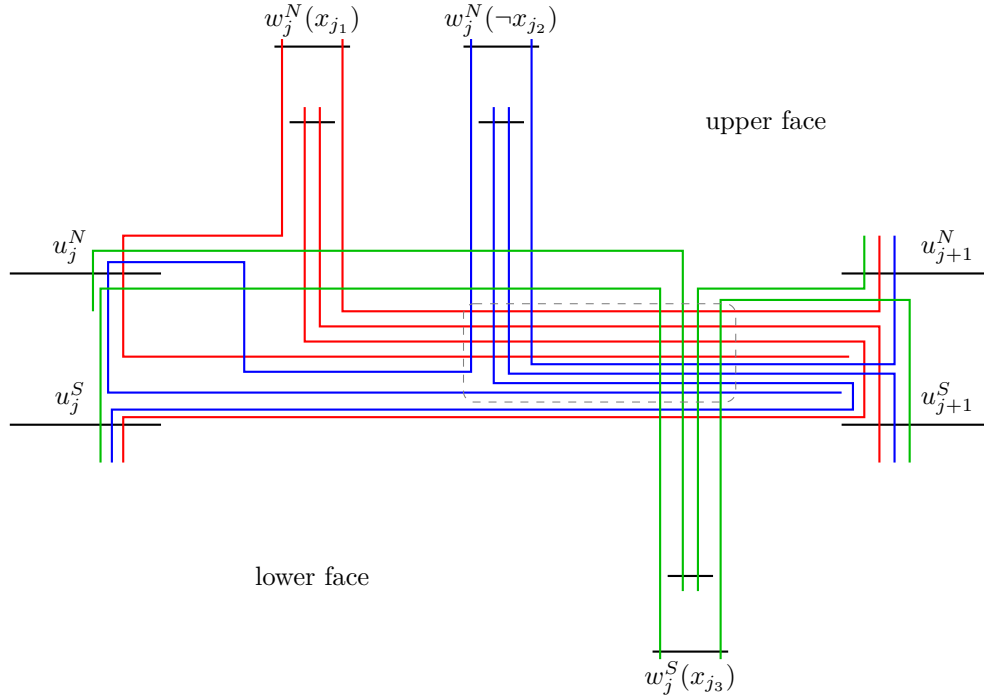
Figure 3 String representation of the clause gadget of Figure 2 satisfying all four items. The string intersections for the $3 \cdot 4^2 = 48$ adjacencies of the $K_{4,4,4}$ occur in the dashed box with rounded corners. It is also easy to check that no two strings of the same color intersect, so the $K_{4,4,4}$ remains induced. We kept "half" of the strings of $u_j^N, u_j^S, u_{j+1}^N, u_{j+1}^S$ uncrossed for the representation of clause gadgets $\mathcal{G}(c_{j-1})$ and $\mathcal{G}(c_{j+1})$.

The figure illustrates the case when no three variables of the clause lie on the same face. However it is easy to mirror the green strings to draw the other case. For 2-clauses, one can simply remove four strings of the same color (red, blue, or green) from Figure 3.

---

[1] In all these items, the faces are those of the planar diagram of the string representation of $\mathcal{G}(c_j)$, except the upper and lower faces, which are defined above.

At this point, the string representation of G has all the desired edges (intersections) except for those of the $K_{2,4}$ in variable gadgets. It has no extra edge, due to the planarity of $\mathcal{P}$. Figure 4 finally shows how to edit the strings $w_\bullet^N(\bullet)$ and $w_\bullet^S(\bullet)$ to add these edges, without incurring any other edges.
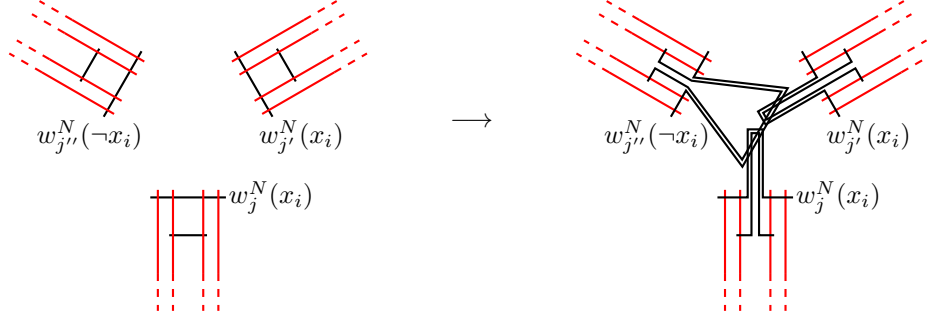
■ **Figure 4** How to realize the $K_{2,4}$ of the variable gadget $\mathcal{G}(x_i)$ without creating any other string intersections. The black strings represent vertices of $\mathcal{G}(x_i)$, and the red strings, vertices in clause gadgets. In this example, x_i is in the upper face, and has two positive occurrences and one negative occurrence.

After editing every string of the variable gadgets, we obtain a string representation for G. ◁

We finally show the following property of G.

▷ **Claim 12.** There is a subcubic planar graph H such that G is a subgraph of H^{16}.

Proof. In the planar drawing $\mathcal{P}$ (see proof of Claim 11), replace every clause vertex by a path of length 16, and every variable vertex by a cycle of length 12, in such a way that the obtained graph H is subcubic, still planar, and the variable-clause incidences are preserved. It is then easy to see that G is a subgraph of H^{16}. ◁

This finishes the proof of the main theorem. ◀

We then derive the following.

▶ **Theorem 2.** INDUCED DISJOINT S–T PATHS with $|S| = |T| = 2$ *is NP-complete in string graphs of bounded maximum degree and twin-width, and requires time* $2^{\Omega(\sqrt{n})}$ *on n-vertex such graphs, unless the Exponential-Time Hypothesis fails.*

Proof. By Theorem 6, planar graphs have bounded twin-width. By Theorem 7, constant powers of planar graphs have bounded twin-width. Constant powers of bounded-degree graphs have themselves bounded maximum degree, and in particular, no $K_{t,t}$ subgraph for some constant t. Thus by Theorem 8, the graphs produced by the previous reduction have bounded twin-width. Besides, they are bounded-degree string graphs.

The previous reduction also works for INDUCED DISJOINT S–T PATHS as the invariant shown in Claim 10 can be established in the same way under the weaker assumption that G admits two mutually induced paths between u_1^N and $\{u_{m+1}^N, u_{m+1}^S\}$, and between u_1^S and $\{u_{m+1}^N, u_{m+1}^S\}$, respectively. In particular, the invariant shows that (if there is at least one clause) there cannot be u_1^N–u_{m+1}^S and u_1^S–u_{m+1}^N paths that are mutually induced in G.

The reduction from CLAUSE-LINKED PLANAR E3-OCC 3-SAT of the proof of Theorem 1 is linear. Indeed it creates a constant number of vertices for each variable and for each clause. By the Sparsification Lemma of Impagliazzo, Paturi, and Zane [16] and classic linear

reductions (see [9, 30]), n-variable Clause-Linked Planar E3-OCC 3-SAT requires time $2^{\Omega(\sqrt{n})}$, under the ETH. We conclude that Induced Disjoint S–T Paths with $|S| = |T| = 2$ requires time $2^{\Omega(\sqrt{n})}$ on n-variable string graphs of bounded twin-width and maximum degree. ◀

4   Hardness of detecting a subcubic graph as an induced subdivision

As a consequence of the previous section, we can prove Theorem 3, which we recall.

▶ **Theorem 3.** *H-Induced Subdivision Containment is NP-complete for the subcubic graph H of Figure 1, and requires time $2^{\Omega(\sqrt{n})}$ on n-vertex graphs, unless the ETH fails.*

Proof. We reduce from Induced 2-Disjoint Paths in string graphs. Given any input (G', s_1, t_1, s_2, t_2), we construct a graph G as follows. Take the disjoint union of G' and H, remove the edges st and $s't'$ (of H), and identify s_1 with s, t_1 with t, s_2 with s', and t_2 with t'. For the sake of clarity, for any $i \in [4]$, vertex set A_i (in H) is renamed B_i in G. Hence $V(G) = V(G') \cup \bigcup_{i \in [4]} B_i$. One can observe that if (G', s_1, t_1, s_2, t_2) is a positive instance of Induced 2-Disjoint Paths, then H is an induced subdivision of G. We show that the converse also holds, hence this linear reduction is correct.

Let $(\phi : V(H) \to V(G), (P_e)_{e \in E(H)})$ be an induced subdivision model of H in G. We say that a vertex of H is *mapped* to its image by ϕ.

▷ **Claim 13.**   For any $i \in [4]$, no vertex of A_i can be mapped to a vertex of $V(G')$.

Proof.   No subdivision of $K_{3,3}$ has a bridge. However, any induced subdivision in G whose branching vertices intersect both $V(G')$ and $\bigcup_{i \in [4]} B_i$ contains a bridge. This is because any of s, t, s', t' is a cutvertex in G. Thus for every $i \in [4]$, $\phi(A_i) \subset \bigcup_{i \in [4]} B_i$ or $\phi(A_i) \subset V(G')$. We now simply have to rule out the latter. Observe that there is no path in G between two vertices in $V(G')$ that exits $V(G')$. Thus, if $\phi(A_i) \subset V(G')$, G' would contain an induced subdivision of the 1-subdivision of $K_{3,3}$, which does not hold as G' is a string graph.   ◁

Claim 13 implies that $\phi(\bigcup_{i \in [4]} A_i) = \bigcup_{i \in [4]} B_i$. Again, the absence of cutvertices in each $H[A_i]$ and $G[B_i]$ implies that for every $i \in [4]$ there is some $j \in [4]$ such that $\phi(A_i) = B_j$. As $|A_1| = |A_2| \neq |A_3| = |A_4|$ and ϕ is injective, $\{\phi(A_1), \phi(A_2)\} = \{B_1, B_2\}$ and $\{\phi(A_3), \phi(A_4)\} = \{B_3, B_4\}$. Now for the induced subdivision model of H to be completed, there has to be in G' two mutually induced paths between s_1 and t_1, and between s_2 and t_2. Thus (G', s_1, t_1, s_2, t_2) is indeed a positive instance.   ◀

The previous reduction also works for H-Induced Minor Containment. However, the proof is slightly more involved. We tune H a little bit such that it is still subcubic but has no two adjacent vertices of degree 3 (so that the forthcoming result answers a question of Korhonen and Lokshtanov). We call the resulting graph H'; see Figure 5.

▶ **Theorem 4.** *There is a subcubic graph H' such that every edge of H' is incident to a vertex of degree 2, and H'-Induced Minor Containment is NP-complete, and requires time $2^{\Omega(\sqrt{n})}$ on n-vertex graphs, unless the ETH fails.*

Proof. Again, we reduce from Induced 2-Disjoint Paths in string graphs and build from any input (G', s_1, t_1, s_2, t_2), a graph G in the following way: Take the disjoint union of G' and H', remove the edges st and $s't'$ (of H'), and identify s_1 with s, t_1 with t, s_2 with s', and t_2 with t'. For any $i \in [4]$, vertex set A_i (in H') is renamed B_i in G. Thus
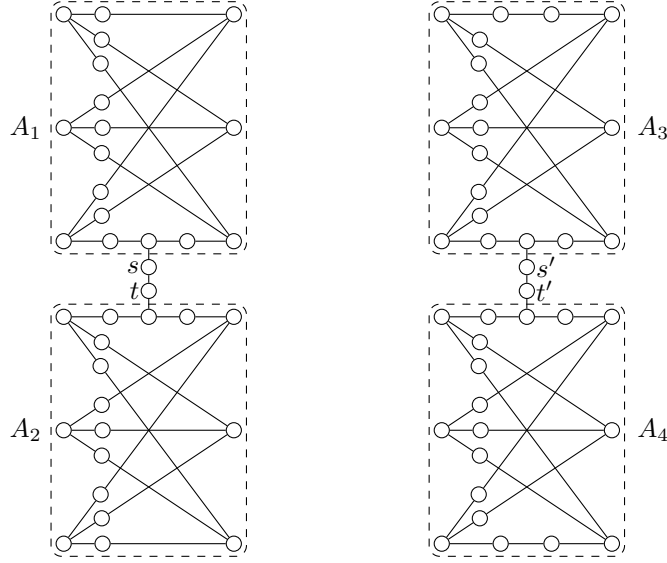
Figure 5 The graph H' obtained from H of Figure 1 by subdividing in each A_i the two edges incident to the vertex with a neighbor in $\{s, t, s', t'\}$. In total H' has 74 vertices.

$V(G) = V(G') \cup \bigcup_{i \in [4]} B_i$. If (G', s_1, t_1, s_2, t_2) is a positive instance of INDUCED 2-DISJOINT PATHS, then H' is an induced subdivision, and hence an induced minor of G. Again, we show that the converse also holds.

Let $h := |V(H')| = 74$, and $(\mathcal{M} := \{X_1, \ldots, X_h\}, \phi : V(H') \to \mathcal{M})$ be a minimal induced minor model of H' in G. First, assume that there is an $i \in [4]$ and $x \in A_i$ such that $\phi(x) \cap V(G') \neq \emptyset$. As G' is a string graph (but $H'[A_i]$ is not), there is also some $y \in A_i$ and $j \in [4]$ such that $\phi(y) \cap B_j \neq \emptyset$. Let u be the vertex of $\{s_1, t_1, s_2, t_2\}$ with a neighbor in B_j. As u is a cutvertex in G that disconnects B_j from the rest of G, there is a $z \in A_i$ such that $u \in \phi(z)$.

As every branch set is connected, for every $z' \in A_i \setminus \{z\}$, either $\phi(z') \subset B_j$ or $\phi(z') \subset V(G - B_j)$. Furthermore, as $H'[A_i]$ has no cutvertex (in particular z is not a cutvertex of $H'[A_i]$), either for every $z' \in A_i \setminus \{z\}$, $\phi(z') \subset B_j$, or for every $z' \in A_i \setminus \{z\}$, $\phi(z') \subset V(G - B_j)$. The latter would contradict the minimality of $(\mathcal{M}, \phi)$ as the (non-empty subset of) vertices of B_j could then be removed from every branch set; and in particular $\phi(y)$ would strictly decrease.

Therefore, for every $z' \in A_i \setminus \{z\}$, $\phi(z') \subset B_j$. Moreover, it should hold that z has a neighbor in $\{s, t, s', t'\}$, say s. Now, one can change the induced minor model by moving $\phi(z) \cap V(G - B_j)$ to the adjacent branch set $\phi(s)$. This indeed still works as an induced minor model of H' in G, and can greedily be made minimal (if need be).

After at most three more similar steps, we obtain a (minimal) induced minor model $(\mathcal{M}', \phi')$ such that for every $x \in \bigcup_{i \in [4]} A_i$, $\phi'(x) \subseteq \bigcup_{i \in [4]} B_i$. We may then conclude as in the proof of Theorem 3. The ETH lower bound is a direct consequence of Theorem 2 and of the fact that the present reduction is linear.   ◄

References

1   Daniel Bienstock. On the complexity of testing for odd holes and induced odd paths. *Discret. Math.*, 90(1):85–92, 1991. See also Corrigendum by B. Reed, *Discrete Mathematics*, 102:109–109, 1992. `doi:10.1016/0012-365X(91)90098-M`.

2   Édouard Bonnet, Colin Geniet, Eun Jung Kim, Stéphan Thomassé, and Rémi Watrigant. Twin-width II: small classes. *Combinatorial Theory*, 2(2), 2022. `doi:10.5070/C62257876`.

3   Édouard Bonnet, Eun Jung Kim, Stéphan Thomassé, and Rémi Watrigant. Twin-width I: tractable FO model checking. *J. ACM*, 69(1):3:1–3:46, 2022. `doi:10.1145/3486655`.

4   Maria Chudnovsky and Paul Seymour. The three-in-a-tree problem. *Combinatorica*, 30(4):387–417, 2010. `doi:10.1007/S00493-010-2334-4`.

5   Maria Chudnovsky, Paul Seymour, and Nicolas Trotignon. Detecting an induced net subdivision. *Journal of Combinatorial Theory, series B*, 103(5):630–641, 2013. `doi:10.1016/J.JCTB.2013.07.005`.

6   Clément Dallard, Maël Dumas, Claire Hilaire, Martin Milanic, Anthony Perez, and Nicolas Trotignon. Detecting $K_{2,3}$ as an Induced Minor. In Adele Anna Rescigno and Ugo Vaccaro, editors, *Combinatorial Algorithms - 35th International Workshop, IWOCA 2024, Ischia, Italy, July 1-3, 2024, Proceedings*, volume 14764 of *Lecture Notes in Computer Science*, pages 151–164. Springer, 2024. `doi:10.1007/978-3-031-63021-7_12`.

7   Clément Dallard, Maël Dumas, Claire Hilaire, and Anthony Perez. Sufficient conditions for polynomial-time detection of induced minors, 2025. `doi:10.48550/arXiv.2501.00161`.

8   Michael R. Fellows. The Robertson-Seymour theorems: A survey of applications. *Contemporary Mathematics*, 89:1–18, 1989.

9   Michael R. Fellows, Jan Kratochvíl, Matthias Middendorf, and Frank Pfeiffer. The complexity of induced minors and related problems. *Algorithmica*, 13(3):266–282, 1995. `doi:10.1007/BF01190507`.

10 Jirí Fiala, Marcin Kaminski, Bernard Lidický, and Daniël Paulusma. The k-in-a-Path Problem for Claw-free Graphs. *Algorithmica*, 62(1-2):499–519, 2012. `doi:10.1007/S00453-010-9468-Z`.

11 Petr A. Golovach, Daniël Paulusma, and Erik Jan van Leeuwen. Induced Disjoint Paths in Claw-Free Graphs. *SIAM J. Discret. Math.*, 29(1):348–375, 2015. `doi:10.1137/140963200`.

12 Petr A. Golovach, Daniël Paulusma, and Erik Jan van Leeuwen. Induced Disjoint Paths in AT-free graphs. *J. Comput. Syst. Sci.*, 124:170–191, 2022. `doi:10.1016/J.JCSS.2021.10.003`.

13 Petr A. Golovach, Giannos Stamoulis, and Dimitrios M. Thilikos. Model-checking for first-order logic with disjoint paths predicates in proper minor-closed graph classes. In Nikhil Bansal and Viswanath Nagarajan, editors, *Proceedings of the 2023 ACM-SIAM Symposium on Discrete Algorithms, SODA 2023, Florence, Italy, January 22-25, 2023*, pages 3684–3699. SIAM, 2023. `doi:10.1137/1.9781611977554.CH141`.

14 Petr Hlinený and Jan Jedelský. Twin-width of planar graphs is at most 8, and at most 6 when bipartite planar. In Kousha Etessami, Uriel Feige, and Gabriele Puppis, editors, *50th International Colloquium on Automata, Languages, and Programming, ICALP 2023, July 10-14, 2023, Paderborn, Germany*, volume 261 of *LIPIcs*, pages 75:1–75:18. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023. `doi:10.4230/LIPICS.ICALP.2023.75`.

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

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

17 Lars Jaffke, O-joung Kwon, and Jan Arne Telle. Mim-Width I. Induced path problems. *Discret. Appl. Math.*, 278:153–168, 2020. `doi:10.1016/J.DAM.2019.06.026`.

18 Ken-ichi Kawarabayashi and Yusuke Kobayashi. A linear time algorithm for the induced disjoint paths problem in planar graphs. *J. Comput. Syst. Sci.*, 78(2):670–680, 2012. `doi:10.1016/J.JCSS.2011.10.004`.

19 Yusuke Kobayashi and Ken-ichi Kawarabayashi. Algorithms for finding an induced cycle in planar graphs and bounded genus graphs. In Claire Mathieu, editor, *Proceedings of the Twentieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2009, New York, NY, USA, January 4-6, 2009*, pages 1146–1155. SIAM, 2009. `doi:10.1137/1.9781611973068.124`.

20   Tuukka Korhonen and Daniel Lokshtanov. Induced-minor-free graphs: Separator theorem, subexponential algorithms, and improved hardness of recognition. In David P. Woodruff, editor, *Proceedings of the 2024 ACM-SIAM Symposium on Discrete Algorithms, SODA 2024, Alexandria, VA, USA, January 7-10, 2024*, pages 5249–5275. SIAM, 2024. `doi:10.1137/1.9781611977912.188`.

21   Ngoc-Khang Le. Detecting an induced subdivision of K_4. *Journal of Graph Theory*, 90(2):160–171, 2019. `doi:10.1002/JGT.22374`.

22   Benjamin Lévêque, David Y. Lin, Frédéric Maffray, and Nicolas Trotignon. Detecting induced subgraphs. *Discret. Appl. Math.*, 157(17):3540–3551, 2009. `doi:10.1016/J.DAM.2009.02.015`.

23   David Lichtenstein. Planar formulae and their uses. *SIAM J. Comput.*, 11(2):329–343, 1982. `doi:10.1137/0211025`.

24   Karl Menger. Zur allgemeinen kurventheorie. *Fundamenta Mathematicae*, 10(1):96–115, 1927.

25   Tung Nguyen, Alex D. Scott, and Paul Seymour. Induced paths in graphs without anticomplete cycles. *J. Comb. Theory B*, 164:321–339, 2024. `doi:10.1016/J.JCTB.2023.10.003`.

26   Marko Radovanovic, Nicolas Trotignon, and Kristina Vuskovic. The (theta, wheel)-free graphs part IV: induced paths and cycles. *J. Comb. Theory B*, 146:495–531, 2021. `doi:10.1016/J.JCTB.2020.06.002`.

27   Neil Robertson and Paul Seymour. Graph Minors. XIII. The Disjoint Paths Problem. *Journal of Combinatorial Theory, Series B*, 63(1):65–110, 1995. `doi:10.1006/jctb.1995.1006`.

28   Neil Robertson and Paul Seymour. Graph minors. XVI. excluding a non-planar graph. *J. Comb. Theory B*, 89(1):43–76, 2003. `doi:10.1016/S0095-8956(03)00042-X`.

29   Frank W. Sinden. Topology of thin film RC circuits. *Bell System Technical Journal*, 45(9):1639–1662, 1966.

30   Simon Tippenhauer and Wolfgang Muzler. On planar 3-SAT and its variants. *Fachbereich Mathematik und Informatik der Freien Universitat Berlin*, 2016.