Interval Deletion is Fixed-Parameter Tractable*

Yixin Cao† Dániel Marx†

Abstract

We study the minimum *interval deletion* problem, which asks for the removal of a set of at most k vertices to make a graph on n vertices into an interval graph. We present a parameterized algorithm of runtime $10^k \cdot n^{O(1)}$ for this problem, thereby showing its fixed-parameter tractability.

1 Introduction

A graph is an interval graph if its vertices can be assigned to intervals of the real line such that there is an edge between two vertices if and only if their corresponding intervals intersect. Interval graphs are the natural models for DNA chains in biology and many other applications, among which the most cited ones include jobs scheduling in industrial engineering [2] and seriation in archeology [20]. Motivated by pure contemplation of combinatorics and practical problems of biology respectively, Hajós [16] and Benzer [3] independently initiated the study of interval graphs.

Interval graphs are a proper subset of chordal graphs. After more than half century of intensive investigation, the properties and the recognition of interval and chordal graphs are well understood (e.g., [4]). More generally, many NP-hard problems (coloring, maximum independent set, etc.) are known to be polynomial-time solvable when restricted to interval and chordal graphs. Therefore, one would like to generalize these results to graphs that do not belong to these classes, but close to them in the sense that they have only a few "erroneous"/"missing" edges or vertices. As a first step in understanding such generalizations, one would like to know how far the given graph is from the class and to find the erroneous/missing elements. This leads us naturally to the area of graph modification problems, where given a graph G, the task is to apply a minimum number of operations on G to make it a member of some prescribed graph class $\mathcal{F}$. Depending on the operations allowed, we can consider, e.g., completion (edge addition), edge deletion, and vertex deletion versions of these problems. Recall that a graph class $\mathcal{F}$ is *hereditary* if any induced subgraph of a graph G in $\mathcal{F}$ is also in $\mathcal{F}$; interval graphs and chordal graphs, among others, are hereditary. For such graph classes, the vertex deletion version can be considered as the most robust variant, which in some sense encompasses both edge addition and edge deletion: if G can be made a member of $\mathcal{F}$ by k_1 edge additions and k_2 edge deletions, then it can be also made a member of $\mathcal{F}$ by deleting at most $k_1 + k_2$ vertices (e.g., by deleting one endvertex of each added/deleted edge).

Unfortunately, most of these graph modification problems are computationally hard: for example, a classical result of Lewis and Yannakakis [23] shows that the vertex deletion problem is NP-hard for *every* nontrivial and hereditary class $\mathcal{F}$, and according to Lund and Yannakakis [24], they are also MAX SNP-hard. Therefore, early work of Kaplan et al. [17] and Cai [6] focused on the fixed-parameter tractability of graph modification problems. Recall that a problem, parameterized by k, is *fixed-parameter tractable (FPT)* if there is an algorithm with runtime $f(k) \cdot n^{O(1)}$, where f is a computable function depending only on k [12]. In the special case when the desired graph class $\mathcal{F}$ can be characterized by a finite number of forbidden (induced) subgraphs, then fixed-parameter tractability of such a problem follows from a basic bounded search tree algorithm [17, 6]. However, many important graph classes, such as forests, bipartite graphs, and chordal graphs have minimal obstructions of arbitrary large size (cycles, odd cycles, and holes, respectively). It is much more challenging to obtain fixed-parameter tractability results for such classes, see results, e.g., on bipartite graphs [29, 19], planar graphs [26, 18], acyclic graphs [8, 10], and minor-closed classes [1, 13].

For interval graphs, the fixed-parameter tractability of the completion problem was raised as an open question by Kaplan et al. [17] in 1994, to which a positive answer with a $k^{2k} \cdot n^{O(1)}$ time algorithm was given by Villanger et al. [31] in 2007. In this paper, we answer the complementary question on vertex deletion:

Theorem 1.1. (**Main result**) *There is a $10^k \cdot n^{O(1)}$ time algorithm for deciding whether or not there is a set of at most k vertices whose deletion makes an n-vertex graph G an interval graph.*

Related work. Let us put our result into context.

*Supported by ERC grant 280152 and OTKA grant NK105645.

†Institute for Computer Science and Control, Hungarian Academy of Sciences, yixin@sztaki.hu, dmarx@cs.bme.hu.

Interval graphs form a subclass of chordal graphs, which are graphs containing no induced cycle of length greater than 3 (also called *holes*). In other words, the minimal obstruction for being a chordal graph might be holes of arbitrary length, hence infinitely many of them. Even so, CHORDAL COMPLETION (to make a graph chordal by the addition of at most k edges) can still be solved by a bounded search tree algorithm by observing that a large hole immediately implies a negative answer to the problem [17, 6]. No such simple argument works for CHORDAL DELETION (to make the graph chordal by removing at most k edges/vertices) and its fixed-parameter tractability was procured by a completely different and much more complicated approach [25, 9].

It is known that a graph is an interval graph if and only if it is chordal and does not contain a structure called "asteroidal triple" (AT for short), i.e., three vertices such that each pair of them is connected by a path avoiding neighbors of the third one [22]. Therefore, in the graph modification problems related to interval graphs, one has to destroy not only all holes, but all ATs as well. The INTERVAL COMPLETION problem was shown to be FPT by Villanger et al. [31]. Their algorithm first destroys all holes by the same bounded search tree technique as in CHORDAL COMPLETION. This step is followed by a delicate analysis of the ATs and a complicated branching step to break them in the resulting chordal graph.

A subclass of interval graphs that received attention is the class of unit interval graphs: graphs that can be represented by intervals of unit length. Interestingly, this class coincides with proper interval graphs, which are those graphs that have a representation with no interval containing another one. It is known that unit interval graphs can be characterized as not having holes and other three specific forbidden subgraphs, thus graph modification problems related to unit interval graphs are very different from those related to interval graphs, where the minimal obstructions include an infinite family of ATs [30].

Our techniques. Even though both CHORDAL DELETION and INTERVAL COMPLETION seem related to INTERVAL DELETION, our algorithm is completely different from the published algorithms for these two problems. The algorithm of Marx [25] for CHORDAL DELETION is based on iterative compression, identifying irrelevant vertices in large cliques, and the use of Courcelle's Theorem on a bounded treewidth graph; none of these techniques appears in the present paper.

Villanger et al. [31] used a simple bounded search tree algorithm to try every minimal way of completing all the holes; therefore, one can assume that the input graph is chordal. ATs in a chordal graph are known to have the property of being *shallow*, and in a minimal witness of an AT, every vertex of the triple is *simplicial*. This means that the algorithm of [31] can focus on completing such ATs (see also [7]). On the other hand, there is no similar upper bound known on the number of minimal ways of breaking all holes by removing vertices, and it is unlikely to exist. Therefore, in a sense, INTERVAL DELETION is inherently harder than INTERVAL COMPLETION: in the former problem, we have to deal with two types of forbidden structures, holes and shallow ATs, while in the second problem, only shallow ATs concern us. Indeed, we spend significant effort in the present paper to make the graph chordal; the main part of the proof is understanding how holes interact and what the minimal ways of breaking them are.

The main technical idea to handle holes is developing a reduction rule based on the modular decomposition of the graph and analyzing the structural properties of reduced graphs. It turns out that the holes remaining in a reduced graph interact in a very special way (each hole is fully contained in the closed neighborhood of any other hole). This property allows us to prove that the number of minimal ways of breaking the holes is polynomially bounded, and thus a simple branching step can reduce the problem to the case when the graph is chordal. As another consequence of our reduction rule, we can prove that this chordal graph already has a structure close to interval graphs (it has a clique tree that is a caterpillar). We can show that in such a chordal graph, ATs interact in a well-behaved way and we can find a set of 10 vertices such that there always exists a minimum solution that contains at least one of these 10 vertices. Therefore, we can complete our algorithm by branching on the deletion of one of these vertices.

Motivation. Many classical graph-theoretic problems can be formulated as graph deletion to special graph classes. For instance, VERTEX COVER, FEEDBACK VERTEX SET, CLUSTER VERTEX DELETION, and ODD CYCLE TRANSVERSAL can be viewed as vertex deletion problems where the class $\mathcal{F}$ is the class of all empty graphs, forests, cluster graphs (i.e., disjoint union of cliques), and bipartite graphs, respectively. Thus the study of graph modification problems related to important graph classes can be seen as a natural extension of the study of classical combinatorial problems. In light of the importance of interval graphs, it is not surprising that there are natural combinatorial problems that can be formulated as, or computationally reduced to INTERVAL DELETION, and then our algorithm for INTERVAL DELETION can be applied. For instance, Thm. 1.1 has recently been used as a subroutine to solve the maximum CONSECUTIVE ONES SUB-MATRIX problem and the minimum CONVEX BIPARTITE DELETION problem [27].

2 Outline

Before embarking upon a presentation of our algorithm in full details, let us describe the main steps at a high level. We say that a set $Q \subset V(G)$ is an *interval deletion set* to a graph G if $G - Q$ is an interval graph. An interval deletion set Q is *minimum* if there is no interval deletion set strictly smaller than $|Q|$, and it is *minimal* if no proper subset $Q' \subset Q$ is an interval deletion set. A set X of vertices is called a *minimal forbidden set* if X does not induce an interval graph but every proper subset $X' \subset X$ does; the subgraph $G[X]$ is called a *minimal forbidden induced subgraph*. Clearly, set Q is an interval deletion set if and only if it intersects every minimal forbidden set. Our goal is to find an interval deletion set of size at most k. For technical reasons, it will be convenient to define the problem as follows:

> Given a graph G and an integer k, return
> - if an interval deletion set of size $\leq k$ exists, a *minimum* interval deletion set $Q \subset V(G)$;
> - otherwise, "NO."

PHASE 1: Preprocessing. The first phase of the algorithm applies two reduction rules exhaustively. They either simplify the instance or branch into a constant number of instances with strictly smaller parameter value. The first reduction rule is straightforward: we destroy every forbidden set of size at most 10.

Reduction 1. **[Small forbidden sets]** *Given an instance (G, k) and a minimal forbidden set X of no more than 10 vertices, we branch into $|X|$ instances, $(G - v, k - 1)$ for each $v \in X$.*

A graph on which Reduction 1 cannot be applied is called *prereduced*. It can be checked in polynomial time whether Reduction 1 is applicable by enumerating every set X of size at most 10 (as discussed in §4, it is possible to do this more efficiently, but optimizing the exponent of n in the running time is not the focus of the paper).

The second reduction rule is less obvious and more involved. Recall that a subset M of vertices forms a *module* if each vertex in M has the same neighbors outside M [15]. A module M of G is *nontrivial* if $1 < |M| < |V(G)|$. We observe (see §4.2) that a minimal forbidden set X of at least 5 vertices is either fully contained in a module M or contains at most one vertex of M. Moreover, if $X \cap M = \{x\}$, then replacing x by any other vertex $x' \in M \setminus \{x\}$ in X results in another minimal forbidden set. This permits us to branch on modules, as described in the following reduction rule.

Reduction 2. **[Main]** *Let $I = (G, k)$ be an instance where the graph G is prereduced, and a nontrivial module M that does not induce a clique.*

1. *If every minimal forbidden set is contained in M, then return the instance $(G[M], k)$.*
2. *If no minimal forbidden set is contained in M, then return the instance (G_M, k), where G_M is obtained from G by inserting edges to make $G[M]$ a clique.*
3. *Otherwise, we solve three instances: $I_1 = (G - M, k - |M|)$, $I_2 = (G[M], k - 1)$, and $I_3 = (G', k - 1)$, where G' is obtained from G by adding a clique M' of $(k + 1)$ vertices, connecting every pair of vertices $u \in M'$ and $v \in N(M)$, and deleting M; letting Q_1, Q_2, and Q_3 be the solutions of these instances respectively, we return the smaller of $Q_1 \cup M$ and $Q_2 \cup Q_3$ ("NO" when $|Q_2 \cup Q_3| > k$).*

That is, in the third case we branch into two directions: the solution is obtained either as the union of M and the solution of I_1, or as the union of solutions of I_2 and I_3. The two branches correspond to the two cases where the solution fully contains M or only a minimum interval deletion set to $G[M]$ (i.e., Q_2), respectively. Note that in the second branch, it can be shown that Q_3 is disjoint from M'; hence $Q_2 \cup Q_3$ is indeed a subset of $V(G)$. Moreover, we have to clarify what the behavior of the reduction is if one or more of Q_1, Q_2, and Q_3 are "NO." If Q_2 or Q_3 is "NO," then we define $Q_2 \cup Q_3$ to be "NO" as well. If one of Q_1 and $Q_2 \cup Q_3$ is "NO," we return the other one; if both of them are "NO," we return "NO" as well.

A graph on which neither reduction rule applies is called *reduced*; in such a graph, every nontrivial module induces a clique. In §4, we prove the correctness of the reductions rules and that it can be checked in polynomial time if a reduction rule is applicable. Hence after exhaustive application of the reductions, we may assume that the graph is reduced.

The reductions are followed by a comprehensive study on reduced graphs that yields two crucial combinatorial statements. The first statement is on ATs that are witnessed by a minimal forbidden induced subgraph different from a hole. Of such an AT $\{x, y, z\}$, we say that x is the *shallow terminal* if the defining path (for the AT) between y and z is strictly longer than the other two defining paths. We prove the shallow terminal x is simplicial in G, i.e., $N(x)$ induces a clique.

Theorem 2.1. **[Shallow terminals]** *All shallow terminals in a reduced graph are simplicial.*

We say that two holes are *congenial* to each other if each vertex of one hole is a neighbor of the other hole. It turns out that

THEOREM 2.2. **[Congenial holes]** *All holes in a reduced graph are congenial to each other.*

PHASE 2: Breaking holes. A consequence of Thm. 2.2 is that if a vertex v is in a hole, then $N[v]$ intersects every hole and thus makes a *hole cover*. Intuitively, this suggests that a minimal hole cover has to be very local in a certain sense. Indeed, by relating minimal hole covers in the reduced graph to minimal separators in the subgraph $G - N[v]$, we are able to establish a quadratic bound on the number of minimal hole covers, and more importantly, a cubic time algorithm for constructing them.

THEOREM 2.3. **[Hole covers]** *Every reduced graph of n vertices contains at most n^2 minimal hole covers, and they can be enumerated in $O(n^3)$ time.*

Any interval deletion set must be a hole cover, and thus contains a minimal hole cover. This allows us to branch into at most n^2 instances, in each of which the input graph is chordal. Note that this branching step is applied only once; hence only a polynomial factor will be induced in the running time.

PHASE 3: Breaking ATs. As all the holes have been broken, the graph is already chordal at the onset of the third phase. It should be noted that, however, the graph might not be reduced, as new nontrivial non-clique modules can be introduced with the deletion of a hole cover in Phase 2. In principle, we could rerun the reductions of Phase 1 to obtain a reduced instance, but there is no need to do so at this point. The properties that we need in this phase are that graph is prereduced, chordal, and every shallow terminal is simplicial (Thm. 2.1). We give a name to such graphs and compare it with previously defined notions here.

- A graph is *prereduced* if Reduction 1 does not apply.
- A prereduced graph is *reduced* if Reduction 2 does not apply.
- A prereduced graph is *nice* if it is chordal and every shallow terminal in it is simplicial.

While both reduced graphs and nice graphs are prereduced, they are incomparable to each other. As only vertex deletions are applied after Phase 1, in the remainder of this algorithm the graph is an induced subgraph of that in a previous step. In other words, once a hereditary property is obtained after Phase 1, it remains true thereafter. It is easy to verify that the three defining properties of nice graphs are all hereditary. On the one hand, after the end of Phase 1, a reduced graph is prereduced by definition, and according to Thm. 2.1, every shallow terminal in it is simplicial. On the other hand, Phase 2 destroys all holes and the chordal property is obtained. Therefore, the graph becomes nice after Phase 2 and will remain nice till the end of our algorithm.

By definition, the removal of all simplicial vertices from a nice graph breaks all ATs, thereby yielding an interval graph. This implies that a nice graph has a very special structure: It has a clique tree decomposition where the tree is a caterpillar, i.e., a path with degree-1 vertices attached to it. In other words, all vertices other than the shallow terminals can be arranged in a linear way, which greatly simplifies the examination of interactions between ATs. As a consequence, we can select an AT that is minimal in a certain sense, and single out 10 vertices such that there must exist a minimum interval deletion set destroying this AT with one of these 10 vertices. We can thereby safely branch on removing one of these 10 vertices.

THEOREM 2.4. **[Nice graphs]** *There is a $10^k \cdot n^{O(1)}$ time algorithm for* INTERVAL DELETION *on nice graphs.*

Putting together these steps (see Fig. 1), the fixed-parameter tractability of INTERVAL DELETION follows.

Proof. (of Thm. 1.1) The algorithm described in Fig. 1 solves the problem by making recursive calls to itself, or calling the algorithm of Thm. 2.4 $O(n^2)$ times. In the former case, at most 10 recursive calls are made, all with parameter value at most $k - 1$. In the latter case, the running time is $10^k \cdot n^{O(1)}$. It follows that the total running time of the algorithm is $10^k \cdot n^{O(1)}$. □

We point out that in a straightforward implementation, the constant hidden behind the big-Oh in the exponent of n is 9. We proclaim that we have no intention of optimizing this part, as it will make the algorithm more complicated and hence blur the focus, which is unnecessary.

The paper is organized as follows. §3 recalls some basic facts. §4 presents the details of the first phase. The next four sections are devoted to the proofs of Thms. 2.1–2.4. §§5 and 6 put shallow terminals and congenial holes under thorough examination, and prove Thms. 2.1 and 2.2, respectively. §7 fully characterizes minimal hole covers in reduced graphs and proves Thm. 2.3. §8 presents the algorithm that destroys ATs in nice graphs and proves Thm. 2.4. §9 closes this paper by some possible improvement and new directions.

3 Preliminaries

We write $u \sim y$ ($u \not\sim y$) as a shorthand for the fact that a pair of vertices x and y is adjacent (nonadjacent). By $v \sim X$ we mean v is adjacent to at least one vertex of

Algorithm Interval-Deletion(G, k)

INPUT: a non-interval graph G and a positive integer k

OUTPUT: a minimum interval deletion set $Q \subset V(G)$ of size $\leq k$ or "NO."

1 Reduction 1: Let U be a minimal forbidden set of at most 10 vertices;
  branch on deleting one vertex of U.

\\ *the graph will then be prereduced and remains so hereafter;*

2 Reduction 2: Let M be a nontrivial module of G not inducing a clique.

2.1 **if** all minimal forbidden sets of G are contained in M **then**
  return Interval-Deletion$(G[M], k)$.

2.2 **else if** no minimal forbidden set is contained in M **then**
  return Interval-Deletion(G_M, k), where edges are inserted to make $G[M]$ a clique.

2.3 **else branch** into three instances I_1, I_2, I_3;

\\ *now the graph is reduced;*

3 use the algorithm of Thm. 2.3 to enumerate the at most n^2 minimal hole covers of G.

\\ *the graph will then be nice and remains so hereafter;*

4 **for each** minimal hole cover HC **do**
  use the algorithm of Thm. 2.4 to solve $(G - HC, k - |HC|)$;

5 **return** the smallest solution obtained, or "NO" if all solutions are "NO."

Figure 1: Outline of algorithm for INTERVAL DELETION

the set X, and we say $X \sim Y$ if $v \sim X$ for at least one vertex $v \in Y$. Two vertex sets X and Y are *completely connected* if $x \sim y$ for each pair of $x \in X$ and $y \in Y$. The notation $N_U(v)$ ($N_U[v]$) stands for the (closed) neighborhood of v in the set U, i.e., $N_U(v) = N(v) \cap U$ ($N_U[v] = N[v] \cap U$), regardless of whether $v \in U$ or not. For a graph G, we denote by $|G|$ the cardinality of $V(G)$, and sometimes it is customary to write $v \in G$ rather than $v \in V(G)$.

Chordal graphs admit several important and related characterizations. A set S of vertices *separates* x and y, or called an *x-y separator* if there is no x-y path in the subgraph $G - S$, and *minimal x-y separator* if no proper subset of S separates x and y. For any pair of vertices x and y, a minimal x-y separator is also called a *minimal separator*. A graph is chordal if and only if each minimal separator in it induces a clique [11]. A vertex is *simplicial* if its neighbors induce a clique. A nontrivial chordal graph contains at least two simplicial vertices, and there is at least one simplicial vertex in each connected component after the removal of any separator.

A tree $\mathcal{T}$ whose nodes are the maximal cliques of a graph G is a *(maximal) clique tree* of G if it satisfies the following conditions: any pair of adjacent nodes K_i and K_j defines a minimal separator that is $K_i \cap K_j$; for any vertex $x \in V$, the maximal cliques containing x correspond to a subtree of $\mathcal{T}$. A graph is chordal if and only if it has such a clique tree. A clique tree of a graph G will be denoted by $\mathcal{T}(G)$, or $\mathcal{T}$ when the graph G is clear from the context. Without distinguishing the node in a clique tree and the maximal clique in the graph G corresponding to it, we use K to denote both. A set of vertices is a minimal separator of G if and only if it is the intersection of K_i and K_j, denoted by $S_{i,j}$, for some edge K_iK_j in $\mathcal{T}$ [5]. To be precise, $S_{i,j}$ is a minimal x-y separator for any pair of vertices $x \in K_i \setminus K_j$ and $y \in K_j \setminus K_i$. Since there are at most n maximal cliques in a chordal graph of n vertices [11], a clique tree $\mathcal{T}$ is simpler than G, and commonly considered as a compact representation of G.

All aforementioned properties also apply to interval graphs, where are chordal. Moreover, Fulkerson and Gross [14] showed that each interval graph has a clique tree that is a path.

4 Reduction rules and branching

This section discusses the reduction rules described in §2 in more details.

4.1 Forbidden induced subgraphs Three vertices form an *asteroidal triple*, AT for short, if each pair of them is connected by a path that avoids the neighborhood of the third one. We use *asteroidal witness* (AW) to refer to a minimal induced subgraph that is not a hole and contains an AT but none of its proper induced subgraphs does. It should be easy to check that an AW contains precisely one AT, and its vertices are the union of these three defining paths for this triple; the three defining vertices will be called *terminals* of this AW. It can be observed from Fig. 2 that the three terminals are the only simplicial vertices of this AW and they are nonadjacent to each other. Lekkerkerker and Boland [22] observed that a graph is an interval graph if and only if it is chordal and contains no AW, and more importantly, proved the following characterization.

THEOREM 4.1. ([22]) *A minimal non-interval graph is either a hole or an AW depicted in Fig. 2.*

Some remarks are in order. First, it is easy to verify that a hole of 6 or more vertices witnesses an AT (specifically, any three nonadjacent vertices from it) and is minimal, but following convention, we only refer to it as a hole, while reserve the term AW for graphs listed in Fig. 2. Second, the set of AWs depicted in Fig. 2 are not a literal copy of the original list in [22], which contains neither net nor tent. We single out nets and tents, which can be viewed as †-AWs with $d = 2$ and ‡-AWs with $d = 1$, respectively, for the convenience of later presentation. To avoid ambiguities, in this paper we explicitly require the length of the longest defining path of a †-AW and a ‡-AW to be at least 4 (i.e., $d \geq 3$) and 3 (i.e., $d \geq 2$) respectively. Third, each of the four subgraphs in the first row of Fig. 2 consists of a constant number, 6 or 7, of vertices, and thus can be easily located and disposed of by standard enumeration. For the purpose of the current paper, we are mainly concerned with the two kinds of AWs in the second row, whose sizes are unbounded. In the three paths defining a †-AW (‡-AW resp.), two of them have length exactly 3 (2 resp.), and the third strictly larger than 3 (2 resp.). Among the three terminals, the one at distance 3 in a †-AW or 2 in a ‡-AW to both other terminals is called the *shallow terminal*, whose neighbor(s) are the *center(s)*. The other two terminals are called *base terminals*, and other vertices are called *base vertices*. The whole set of base vertices is called the *base*; we point out that base terminals are not a part of the base. Note that the defining path between base terminals is strictly longer than the other two defining paths; and base vertices are the inner vertices of this path. We use $(s : c : l, B, r)$ ($(s : c_1, c_2 : l, B, r)$ resp.) to denote the †-AW (‡-AW resp.) with shallow terminal s, center c (centers c_1 and c_2 resp.), base terminals l, r, and base $B = \{b_1, \ldots, b_d\}$. For the sake of notational convenience, we will also use b_0 and b_{d+1} to refer to the base terminals l and r, respectively. The center(s) and base vertices are called non-terminal vertices.

In time $O(n^5)$, we can find a small minimal forbidden set of at most 10 vertices or assert its nonexistence as follows. For a hole, we guess three consecutive vertices $\{h_1, h_2, h_3\}$, and then search for a shortest h_1-h_3 path in $G - (N[h_2] \setminus \{h_1, h_3\})$. For an AW, we guess three pairwise nonadjacent vertices $\{t_1, t_2, t_3\}$, and for $i = 1, 2, 3$, search for a shortest path between other two terminals in $G - N[t_i]$. As such Reduction 1 can be applied in polynomial time, and after its exhaustive application, the graph is *prereduced*. By definition, any AW in a prereduced graph contains at least 11 vertices, which rules out long claws, whipping tops, nets, and tents. Furthermore, the base of a †-AW (‡-AW resp.) in a prereduced graph contains at least 7 (6 resp.) vertices.

The following structural observations are immediate from the definition of prereduced graphs. They arise frequently in what follows, and hence we collect them here for later reference. Proofs of marked propositions are left for the full version.

PROPOSITION 4.1. ($\star$) *Let $P = (v_0 \ldots v_p)$ be a chordless path of length p in a prereduced graph, and u be adjacent to every inner vertex of P.*

(1) If $p \geq 4$ and u is also adjacent to v_0 and v_p, then $N[v_\ell] \subseteq N[u]$ for every $2 \leq \ell \leq p-2$.

(2) If $p \geq 3$ and u is also adjacent to v_0 and v_p, then $N[v_\ell] \cap N[v_{\ell+1}] \subseteq N[u]$ for every $1 \leq \ell \leq p-2$.

(3) If $p \geq 4$, then $N[v_\ell] \setminus (N(v_1) \cup N(v_{p-1})) \subseteq N[u]$ for every $2 \leq \ell \leq p-2$.

Let X be a nonempty set of vertices. A vertex v is a *common neighbor* of X if it is adjacent to every vertex $x \in X$. We denote by $\widehat{N}(X)$ the set of all common neighbors of X. It is easy to verify that in a prereduced graph, at least one of X and $\widehat{N}(X)$ induces a clique, as otherwise two nonadjacent vertices in $\widehat{N}(X)$, together with two nonadjacent vertices in X, will induce a 4-hole. In particular, we have the following proposition.

PROPOSITION 4.2. *Let X be a set of vertices of a prereduced graph that induces either a hole, an AW, or a path of length at least 2. Then $\widehat{N}(X)$ induces a clique.*

4.2 Modular decomposition A subset M of vertices forms a *module* of G if all vertices in M have the same neighborhood outside M. In other words, for any pair of vertices $u, v \in M$ and vertex $x \notin M$, $u \sim x$ if and only if $v \sim x$. The set $V(G)$ and all singleton vertex sets are modules, called *trivial*. A brief inspection shows that no graph in Fig. 2 has any nontrivial modules and this is true also for holes of length greater than 4:

PROPOSITION 4.3. *Let M be a module. If a minimal forbidden set X contains more than 4 vertices, then either $X \subseteq M$ or $|M \cap X| \leq 1$.*

Indeed, the only minimal forbidden set of size no more than 4 is a 4-hole, of which the pair of nonadjacent vertices might belong to a module. This observation allows us to prove the following statement, which is the main combinatorial reason behind the correctness of the branching in Reduction 2:

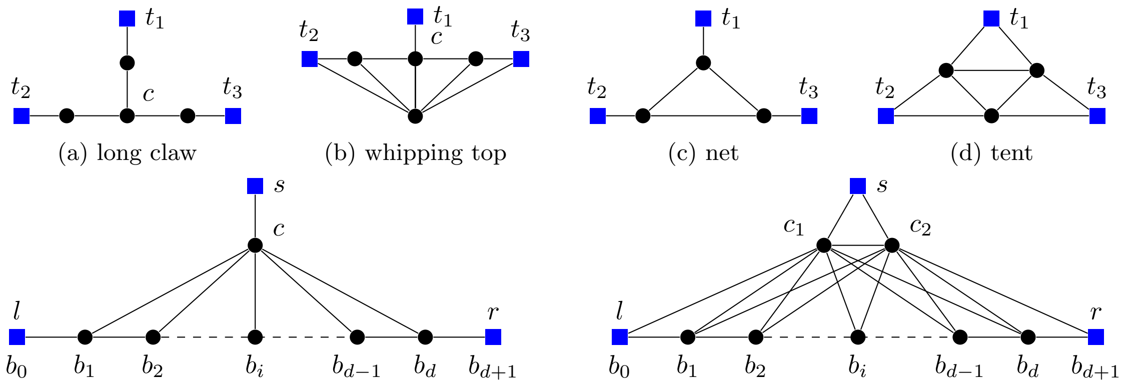

(a) long claw (b) whipping top (c) net (d) tent

(e) †-AW $(s : c : l, B, r)$ $(d = |B| \geq 3)$ (f) ‡-AW $(s : c_1, c_2 : l, B, r)$ $(d = |B| \geq 2)$

Figure 2: Minimal asteroidal witnesses in a chordal graph (terminals are marked as squares).

THEOREM 4.2. (⋆) *Let M be a module of a prereduced graph G. A minimum interval deletion set to G contains either all vertices of M, or only a minimum interval deletion set to $G[M]$.*

We are now ready to prove the correctness of Reduction 2 and explain its application.

LEMMA 4.1. *Reduction 2 is correct, and it can be checked in polynomial time whether Reduction 2 (and which case of it) is applicable.*

Proof. The correctness of the reduction is clear in case 1: removing the vertices of $V(G) \setminus M$ does not make the problem any easier, as these vertices do not participate any minimal forbidden set.

In case 2, the correctness of the reduction follows from the fact that G and G_M have the same set of minimal forbidden sets. Note that a clique is an interval graph, and more importantly, the insertion of edges to make M a clique neither breaks the modularity of M nor introduces any new 4-hole; thus Prop. 4.3 is applicable for G_M. As M induces an interval graph in both G and G_M, if X is a minimal forbidden set in G or G_M, then Prop. 4.3 implies that X contains at most one vertex of M. In other words, the insertion of edges has no effect on any minimal forbidden set, which means that Q is an interval deletion set to G if and only if it is an interval deletion set to G_M.

The correctness of case 3 can be argued using Thm. 4.2, which states the two possibilities of any interval deletion set to G with respect to M. In particular, the two branches of case 3 correspond to these two cases. The first branch is straightforward: we simply remove all vertices of M from the graph and solve the instance $I_1 = (G - M, k - |M|)$. It is the second branch (where we assume $M \not\subseteq Q$) that needs more explanation. Recall that by construction of I_3, the set M' is a module of G' and induces an interval graph. It is clear that either solution Q_2 or Q_3 is "NO" will rule out the existence of an interval deletion set of G that does not fully contain M. Hence we may assume Q_2 and Q_3 are minimum interval deletion sets of I_2 and I_3, respectively; and $Q = Q_2 \cup Q_3$. Note that both $|Q_2|$ and $|Q_3|$ are no more than $k - 1$.

CLAIM 1. *Set Q is an interval deletion set of G.*

Proof. Suppose that there is a minimal forbidden set X disjoint from Q. It cannot be fully contained in M, as $Q_2 \subseteq Q$ is an interval deletion set of $G[M]$. Then by Prop. 4.3, X contains exactly one vertex x of M and $X' = X \setminus \{x\} \cup \{x'\}$ is also a minimal forbidden set of G' for any $x' \in M'$. As Q_3 is an interval deletion set of G' disjoint from M', it has to contain a vertex of $X' \setminus \{x'\} = X \setminus \{x\}$. Now it remains to show $Q \subset V(G)$, which is equivalent to $Q_3 \cap M' = \emptyset$. According to Thm. 4.2, if Q_3 intersects M', then it must contain all $(k+1)$ vertices in M',[1] and then has size strictly larger than k; a contradiction. ⌟

CLAIM 2. *Set Q is not larger than the smallest interval deletion set Q' satisfying $M \not\subseteq Q'$.*

Proof. Suppose that Q' is an interval deletion set of G of size at most k with $M \not\subseteq Q'$; let $Q'_2 = Q' \cap M$ and $Q'_3 = Q' \setminus M$. We claim that Q'_2 and Q'_3 are interval deletion sets of I_2 and I_3, respectively. First, we argue that Q'_2 and Q'_3 are not empty; hence both of them has size at most $k-1$. The assumption that $G[M]$ is not an interval graph implies $Q'_2 \neq \emptyset$. By assumption, $M \not\subseteq Q'$, thus there is a vertex $x \in M \setminus Q'$. Now $Q'_3 = \emptyset$ would imply that $G - (M \setminus \{x\})$ is an interval graph, that is, there is no minimal forbidden set containing only one

[1] Indeed, $\min(k+1, |N(M)|)$ vertices will suffice for our bookkeeping purpose, and an alternative way to this is to add only one vertex but mark it as "forbidden."

vertex of M, and it follows that we should have been in Case 1. Since $|Q_2'| \leq k-1$, it is clear that Q_2' is a solution of instance $I_2 = (G[M], k-1)$. The only way Q_3' is not a solution of I_3 is that there is a minimal forbidden set X containing a vertex of the $(k+1)$-clique introduced to replace M. As this $(k+1)$-clique is a module, Prop. 4.3 implies that X contains exactly one vertex y of this clique. But in this case $X' = X \setminus \{y\} \cup \{x\}$ (where x is a vertex of $M \setminus Q'$) is a minimal forbidden set disjoint from Q', a contradiction. Thus $|Q| \leq |Q'|$ follows from the fact that both Q_2 and Q_3 are minimum. ⌟

As a consequence of Claim 2, if $|Q| > k$, then there cannot be an interval deletion set of size no more than k that does not fully include M. This finishes the proof of the correctness of Reduction 2.

On the applicability of Reduction 2, we first search for a nontrivial module that does not induce a clique. If such a module M is found, then Reduction 2 is applicable, and it remains to figure out which case should apply by checking the conditions in order. To check whether case 1 holds, we need to check if there is a minimal forbidden set X not contained in M. By Prop. 4.3, such an X, if exists, contains at most one vertex x from M; and x can be replaced by any other vertex of M. Therefore, it suffices to pick any vertex $x \in M$, and test in linear time whether $G - (M \setminus \{x\})$ is an interval graph. If it is not an interval graph, then there is an minimal forbidden set X not contained in M (as it contains at most one vertex of M). Otherwise, $G - (M \setminus \{x\})$ is an interval graph for every $x \in M$, and there is no such X; hence case 1 holds. To check whether case 2 holds, observe that the condition "there is no minimal forbidden set contained in M" is equivalent to saying that $G[M]$ is an interval graph, which can be checked in linear time. In all remaining cases, we are in case 3. □

5 Shallow terminals

This section proves Thm. 2.1 by showing that each shallow terminal is contained in a nontrivial module whose neighborhood induces a clique. As Reduction 2 cannot be applied, this module induces a clique, which means that all vertices in this module are simplicial. Recall that an AW in a prereduced graph G has to be a †- or ‡-AW. Let us start from a thorough scrutiny of neighbors of its shallow terminal, which, by definition, is disjoint from the base and base terminals.

Lemma 5.1. *Let W be an AW in a prereduced graph. Every common neighbor x of the base B is adjacent to the shallow terminal s.*

Proof. The center(s) of W are also common neighbors of B, and hence according to Prop. 4.2, they are adjacent to x. Suppose, for contradiction, $x \in \widehat{N}(B) \setminus N(s)$. If W is a †-AW, then there is (see the first row of Fig. 6) • a whipping top $\{s, c, l, b_1, x, b_d, r\}$ centered at c when $x \sim l, r$; • a net $\{s, c, l, b_1, r, x\}$ when $x \sim r$ but $x \not\sim l$ (similarly for $x \sim l$ but $x \not\sim r$); or • a †-AW $\{s : c : l, b_1 x b_d, r\}$ when $x \not\sim l, r$. If W is a ‡-AW, then there is (see the second row of Fig. 6) • a tent $\{x, c_1, b_1, s, b_d, c_2\}$ when $x \sim l, r$; • a ‡-AW $(s : c_1, c_2 : l, b_1 x, r)$ when $x \sim r$ but $x \not\sim l$ (similarly for $x \sim l$ but $x \not\sim r$); or • a ‡-AW $(s : c_1, c_2 : l, b_1 x b_d, r)$ when $x \not\sim l, r$. As none of these structures can exist in a prereduced graph, this proves this lemma. □

Lemma 5.2. *Let W be an AW in a prereduced graph G, and x is adjacent to the shallow terminal s of W.*

(1) Then x is also adjacent to the center(s) of W (different from x).

(2) Classifying x with respect to its adjacency to the base B of W, we have the following categories:

(full) *x is adjacent to every base vertex.*
Then x is also adjacent to every vertex in $N(s) \setminus \{x\}$.

(partial) *x is adjacent to some, but not all base vertices.*
Then there is an AW whose shallow terminal is s, one center is x, and base is a proper subpath of B.

(none) *x is adjacent to no base vertex.*
Then x is adjacent to neither base terminals, and thus replacing the shallow terminal of W by x makes another AW.

Proof. Assume to the contrary of statement (1), $x \not\sim c$ if W is a †-AW or (without loss of generality) $x \not\sim c_2$ if W is a ‡-AW. If $x \sim b_i$ for some $1 \leq i \leq d$ then there is a 4-hole $(xscb_ix)$ or (xsc_2b_ix) (See Fig. 7(a)). Hence we may assume $x \not\sim B$. (See Fig. 7(b,c,d,e).) There is • a 5-hole $(xscb_1lx)$ or $(xscb_drx)$ if W is a †-AW, and $x \sim l$ or $x \sim r$, respectively; • a 5-hole (xsc_2b_1lx) or 4-hole (xsc_2rx) if W is a ‡-AW, and $x \sim l$ or $x \sim r$, respectively; • a long-claw $\{x, s, c, b_1, l, b_d, r\}$ if W is a †-AW and $x \not\sim l, r$; • a net $\{x, s, l, c_1, r, c_2\}$ if W is a ‡-AW and $x \not\sim c_1, l, r$; or • a whipping top $\{r, c_2, s, x, c_1, l, b_1\}$ centered at c_2 if W is a ‡-AW and $x \not\sim l, r$, but $x \sim c_1$. Neither of these cases is possible, and thus statement (1) is proved.

For statement (2), let us handle category "none" first. Note that x, nonadjacent to B, cannot be a center of W. If $x \sim l$, then there is a 4-hole (xcb_1lx) or (xc_2b_1lx) when W is a †-AW or ‡-AW, respectively. A symmetric argument will rule out $x \sim r$. Now that x is adjacent to the center(s) but neither base terminals nor base vertices of W, then $(x : c : l, B, r)$

		$q = p + 1$	$q = p + 2$	$q > p + 2$
†-AW	$p = 0$ (Fig.8a)	4-hole $(xcb_1lx)^*$	tent $\{l, x, s, c, b_2, b_1\}$	‡-AW $(s : x, c : l, b_1 \dots b_{q-1}, b_q)^{**}$
	$p = 1$ (Fig.8b8c)	whipping top $\{l, b_1, x, s, c, b_3, b_2\}^{***}$	net $\{l, b_1, s, x, b_3, b_2\}$	†-AW $(s : x : l, b_1 \dots b_{q-1}, b_q)^{**}$
	$p > 1$ (Fig.8d8e)	long-claw[1] $\{b_{p-2}, b_{p-1}, b_p, s, x, b_{p+2}, b_{p+1}\}$	net $\{b_{p-1}, b_p, s, x, b_q, b_{q-1}\}$	†-AW $(s : x : b_{p-1}, b_p \dots b_{q-1}, b_q)$
‡-AW	$p = 0$	4-hole $(xc_2b_1lx)^*$	tent $\{l, x, s, c_2, b_2, b_1\}$	‡-AW $(s : x, c_2 : l, b_1 \dots b_{q-1}, b_q)^{**}$
	$p = 1$	whipping top $\{l, b_1, x, s, c_2, b_3, b_2\}^{***}$	net $\{l, b_1, s, x, b_3, b_2\}$	†-AW $(s : x : l, b_1 \dots b_{q-1}, b_q)^{**}$
	$p > 1$	long-claw $\{b_{p-2}, b_{p-1}, b_p, s, x, b_{p+2}, b_{p+1}\}$	net $\{b_{p-1}, b_p, s, x, b_q, b_{q-1}\}$	†-AW $(s : x : b_{p-1}, b_p \dots b_{q-1}, b_q)$

$*$: The vertex x is in category "none."
$**$: The vertex x would be in category "full" if $q = d + 1$.
$***$: A 4-hole $(xb_pb_{p+1}b_{p+2}x)$ would be introduced if $x \sim b_{p+2}$;

Table 1: Structures used in the proof of Lem. 5.2 (category "partial")

($(x : c_1, c_2 : l, B, r)$ resp.) makes another †-AW (‡-AW resp.).

Assume now that x is in category "full." Suppose the contrary and $x \not\sim v$ for some $v \in N(s) \setminus \{x\}$. We have already proved in statement (1) that v and x are adjacent to the center(s) of W (different from them). In particular, if one of v and x is a center, then they are adjacent. Therefore, we can assume that v and x are not centers. If $v \sim b_i$ for some $1 \leq i \leq d$, then there is a 4-hole $(xsvb_ix)$. Otherwise, $v \not\sim B$, and it is in category "none." Let W' be the AW obtained by replacing s in W by v; then $x \sim v$ follows from Lem. 5.1.

Finally, assume that x is in category "partial," that is, $x \sim B$, but $x \not\sim b_i$ for some $1 \leq i \leq d$. In this case, we construct the claimed AW as follows. As the case $x \not\sim l$ but $x \sim r$ is symmetric to $x \sim l$ but $x \not\sim r$, it is ignored in the following, i.e., we assume that $x \sim r$ only if $x \sim l$. Let p be the smallest index such that $x \sim b_p$, and q be the smallest index such that $p < q \leq d + 1$ and $x \not\sim b_q$ (q exists by assumptions). See Table 1 for the structures for †-AW and ‡-AW respectively (see also Fig. 8).[2]

As the graph is prereduced and contains no small forbidden induced subgraph, it is immediate from Table 1 that the case $q > p + 2$ holds; otherwise there always exists a small forbidden induced subgraphṪhis completes the categorization of vertices in $N(s) \setminus T$. □

[2]We omit the figure for ‡-AWs: For a ‡-AW $(s : c_1, c_2 : l, B, r)$, we are only concerned with the relation between center c_2 and $B \cup \{l\}$, which is the same as the relation between c and $B \cup \{l\}$ in a †-AW.

The proof of our main result of this section is an inductive application of Lem. 5.2. To avoid the repetition of the essentially same argument for †-AWs and ‡-AWs, especially for the interaction between AWs, we use a generalized notation to denote both. We will uniformly use c_1, c_2 to denote center(s) of an AW, and while the AW under discussion is a †-AW, both c_1 and c_2 refer to the only center. As long as we do not use the adjacency of c_1 and l, c_2 and r, or c_1 and c_2 in any of the arguments, this unified (abused) notation will not introduce inconsistencies.

Theorem 5.1. *Let W be a †- or ‡-AW in a prereduced graph G with shallow terminal s and base B. Let $C = N(s) \cap N(B)$ and let M be the vertex set of the connected component of $G - C$ containing s. Then M is completely connected to C, and $G[C]$ is a clique.*

Proof. Denote by $W = (s : c_1, c_2 : l, B, r)$, where $c_1 = c_2$ when W is a †-AW. Let x and y be any pair of vertices such that $x \in C$ and $y \in M$. By definition, $G[M]$ is connected, and there is a chordless path $P = (v_0 \dots v_p)$ from $v_0 = s$ to $v_p = y$ in $G[M]$. We claim that $P \not\sim B$. It holds vacuously if $p = 1$ and then $y \sim s$; hence we assume $p > 1$. Suppose the contrary and let q be the smallest index such that $v_q \sim B$. This means that every v_i with $i < q$ is in category "none" of Lem. 5.2(2). Therefore, applying Lem. 5.2(1,2) on v_i and AW $(v_{i-1} : c_1, c_2 : l, B, r)$ inductively for $i = 1, \dots, q - 1$, we conclude that there is an AW $W_i = (v_i : c_1, c_2 : l, B, r)$ for each $i < q$. One more application of Lem. 5.2(1) shows that v_q is adjacent to the center(s) of W_{q-1} as well. If v_q is adjacent to all vertices of B, i.e., in the category "full" with respect to every W_i, then Lem. 5.2(2) on v_q and W_{q-1} implies that v_q is adjacent to $v_{q-2} \in N(v_{q-1})$, contradicting the assumption that P is chordless. Otherwise (the category "partial"), according to Lem. 5.2(2), there is another AW $W' = (v_{q-1} : c'_1, c'_2 : l', B', r')$, where $B' \subset B$, and $v_q \in \{c'_1, c'_2\}$. Now an application of Lem. 5.2(1) on v_q and W' shows that v_q is adjacent to $v_{q-2} \in N(v_{q-1})$, again a contradiction. From these contradictions we can conclude $P \not\sim B$. Applying Lem. 5.2 inductively on v_{i+1} and $W_i = (v_i : c_1, c_2 : l, B, r)$, we get an AW with the same centers for every $0 \leq i \leq p$.

As x is adjacent to both s and B, it cannot be in category "none" with respect to W. We now separate the discussion based on whether x is in the category "full" or "partial." Suppose first that x is in the category "full"; as $x \in N(s)$, Lem. 5.2(1) implies that $x \sim c_1, c_2$. Then applying Lem. 5.2(2) inductively, where $i = 1, \ldots, p$, on vertex x and W_{i-1} we get that $x \sim v_i$ for every $i \leq p$; in particular, $x \sim v_p$ $(= y)$. Suppose now that x is in in category "partial." Then by Lem. 5.2(2), there is an AW $W'_0 = (v_0 : c'_1, c'_2 : l', B', r')$, where $B' \subset B$, and $x \in \{c'_1, c'_2\}$. As $P \not\sim B$, we have that $v_i \not\sim B'$ for any $0 \leq i \leq p$, i.e., v_i is in category "none" with respect to W'_0. Therefore, by an inductive application of Lem. 5.2(2) on the vertex v_i and AW $W'_{i-1} = (v_{i-1} : c'_1, c'_2 : l', B', r')$ for $i = 1, \ldots, p$, we conclude that there is an AW $W'_p = (v_p : c'_1, c'_2 : l', B', r')$, from which $x \sim y$ follows immediately.

Now we show the second assertion. For any pair of vertices x and y in C, we apply Lem. 5.2 on x and W; by definition, $x \sim B$ and thus cannot be in category "none." If x is in category "full" with respect to W, then Lem. 5.2(2) implies that x is adjacent to $y \in N(s)$. Otherwise, if x is in category "partial" with respect to W, then Lem. 5.2(2) implies that there is an AW $W' = (s : c'_1, c'_2 : l', B', r')$ where $B' \subset B$ and $x \in \{c'_1, c'_2\}$. Therefore, by Lem. 5.2(1) on the vertex $y \in N(s)$ and W', we get that $y \sim c'_1, c'_2$ and hence $x \sim y$. This completes the proof. □

We remark that the set C is an M-B separator. Now Thm. 2.1 follows from Thm. 5.1: the set M containing s is in a module whose neighborhood is a clique, hence every vertex in M is simplicial.

6 Long holes

This section proves Thm. 2.2 by showing that the holes in a reduced graph are pairwise congenial. During the study of vertices of a hole, their indices become very subtle. To simplify the presentation, we will frequently apply a common technique, that is, to number the vertices of a hole starting from a vertex of special interest for the property at hand. Needless to say, indexing two adjacent vertices in a hole will determine the indices of all the vertices in the hole, as well as the ordering used to transverse the hole.

We start from two simple facts on the relations between vertices and holes, from which we derive the relations between two holes, and finally generalize them to multiple holes.

Proposition 6.1. ($\star$) *For any vertex v and hole H of a prereduced graph, $N_H[v]$ are consecutive in H. Moreover, either $N_H[v] = H$ or $|N_H[v]| < |H| - 7$.*

Recall that $\widehat{N}(H)$ is the set of all common neighbors of the hole H. If $3 < |N_H[v]| < |H|$, then we can use v as a shortcut for the inner vertices of the path induced by $N_H[v]$ to obtain another hole that is strictly shorter than H.

Corollary 6.1. *Let H be a shortest hole. If $v \notin \widehat{N}(H)$, then $N_H[v] \leq 3$.*

Note that each hole H in a prereduced graph contains at least 11 vertices. If $v \in \widehat{N}(H)$, then on any five consecutive vertices of the hole H and v, Prop. 4.1(1) applies, which implies that v is dominating in the closed neighborhood of H.

Corollary 6.2. *Let H be a hole in a prereduced graph. If $v \in \widehat{N}(H)$, then v is adjacent to all vertices in $N[H] \setminus \{v\}$.*

So far we characterized neighbors of holes in a prereduced graph: Any vertex v is adjacent to a (possibly empty) set of consecutive vertices of a hole H; if v is adjacent to all vertices of H, then it is also adjacent to every neighbor of H. From these facts we now derive the relations between holes. Following is the most crucial concept of the section:

Definition 6.1. *Two holes H_1 and H_2 are called* congenial *(to each other) if each vertex of one hole is a neighbor of the other hole, that is, $H_1 \subseteq N[H_2]$ and $H_2 \subseteq N[H_1]$.*

We remark that every hole is congenial to itself by definition. The definition is partially motivated by:

Proposition 6.2. *Let $\mathcal{H}$ be a set of holes all congenial to H. For each $v \in H$, every hole in $\mathcal{H}$ intersects $N[v]$.*

Since a vertex in a hole cannot be a common neighbor of it, Cor. 6.2 and the definition of congenial holes immediately imply:

Corollary 6.3. *For any pair of congenial holes H_1 and H_2 in a prereduced graph, $\widehat{N}(H_1) = \widehat{N}(H_2)$. Moreover, no vertex of H_1 (H_2 resp.) is a common neighbor of H_2 (H_1 resp.).*

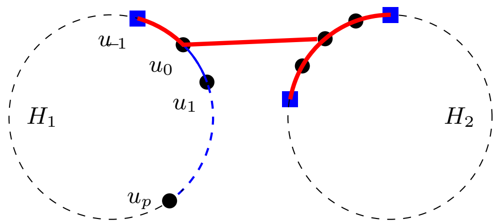

(a) $|N_{H_2}(u_0)| = 1$.

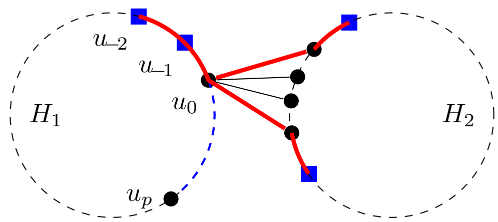

(b) $|N_{H_2}(u_0)| > 1$.

Figure 3: Adjacency of non-congenial holes ($u_{-1} \not\sim H_2$)

We analyze next the relation between two non-congenial holes. It turns out that if not all vertices of a hole H_1 are adjacent to another hole H_2, then, as shown in the following lemma, every vertex of H_1 is adjacent to either all or none of the vertices of H_2.

Lemma 6.1. *Let H_1 and H_2 be two adjacent holes in a prereduced graph. If $H_1 \not\subseteq N[H_2]$, then each neighbor of H_2 in H_1 is a common neighbor of H_2, i.e., $N_{H_1}[H_2] \subseteq \widehat{N}(H_2)$; in particular, H_1 and H_2 are disjoint.*

Proof. Let u be any vertex in $N_{H_1}[H_2]$, which is nonempty by assumption, and let P be the maximal path in H_1 with the property that $u \in P \subseteq N_{H_1}[H_2]$; denote by p the number of vertices of P. Note that some vertices of P can belong to H_2 (in particular, u can be in H_2). Observe that $p < |H_1|$, as by assumption, H_1 is not contained in $N[H_2]$. Numbering the vertices in H_1 such that $P = u_0 \dots u_{p-1}$ (the ordering of H_1 is immaterial when $p = 1$ and then u_1 can be either neighbor of u_0 in H_1), the selection of P means $u_i \sim H_2$ for each $0 \le i < p$, and $u_{-1}, u_p \not\sim H_2$ (it is immaterial whether $u_{-1} = u_p$ or not). In the following, we show that both ends of P belong to $\widehat{N}(H_2)$, which induces a clique (Prop. 4.2). Thus either $u_0 = u_{p-1}$ (i.e., $p = 1$) or u_0 and u_{p-1} are adjacent (i.e., $p = 2$); in either case, we have $u \in \{u_0, u_{p-1}\} \subseteq \widehat{N}(H_2)$. This proves the first assertion, and the second assertion ensues, as otherwise their common vertices will be common neighbors of H_2, which is not possible.

Note that $u_0 \notin H_2$, as otherwise u_{-1} is also adjacent to H_2, contradicting the maximality of P. Similarly, $u_{-1}, u_{-2} \notin H_2$. If u_0 has a unique neighbor v in H_2, then the subgraph induced by u_{-1}, u_0 and five consecutive H_2 vertices centered at v is a long claw (see Fig. 3a). Now we consider the case $2 \le |N_{H_2}[u_0]| \le |H_2| - 7$ (Prop. 6.1), and number the vertices of H_2 such that $N_{H_2}[u_0] = \{v_1, v_2 \dots, v_q\}$. Note that $|N_{H_2}[u_0]| \le |H_2| - 7$ implies that $v_0 \ne v_{q+1}$. If u_{-2} is adjacent to v_0, v_1, v_q, or v_{q+1}, then there is a hole $(u_{-2}u_{-1}u_0v_1v_0u_{-2})$, $(u_{-2}u_{-1}u_0v_1u_{-2})$, $(u_{-2}u_{-1}u_0v_qu_{-2})$, or $(u_{-2}u_{-1}u_0v_qv_{q+1}u_{-2})$, respectively. Otherwise, $u_{-2} \not\sim \{v_0, v_1, v_q, v_{q+1}\}$, then there is a net $\{u_{-1}, u_0, v_0, v_1, v_{q+1}, v_q\}$ when $|N_{H_2}(u_0)| = 2$, or long claw $\{u_{-2}, u_{-1}, u_0, v_0, v_1, v_{q+1}, v_q\}$ when $|N_{H_2}(u_0)| > 2$ (see Fig. 3b). This proves $u_0 \in \widehat{N}(H_2)$, and with a symmetric argument we can also prove $u_{p-1} \in \widehat{N}(H_2)$. □

We are now ready to establish the transitivity of the congenial relation. The reflexivity and symmetry of this relation are clear from definition; therefore congenial holes form an equivalence class.

Lemma 6.2. *Let H, H_1, and H_2 be three holes in a prereduced graph G. If both H_1 and H_2 are congenial to H, then H_1 and H_2 are congenial.*

Proof. According to Cor. 6.3, $\widehat{N}(H_1) = \widehat{N}(H) = \widehat{N}(H_2)$. If H_1 and H_2 are adjacent, then they have to be congenial, as otherwise Lem. 6.1 implies that one of them contains a common neighbor of the other, hence a common neighbor of all three holes, which is impossible. Assume hence $H_1 \not\sim H_2$. Let h be any vertex in H, and we number the vertices of H_1 and H_2 such that $N_{H_1}[h] = \{u_1, \dots, u_p\}$ and $N_{H_2}[h] = \{v_1, \dots, v_q\}$. Prop. 6.1 implies that $u_0 \ne u_{p+1}$ and $v_0 \ne v_{p+1}$. Note that h is adjacent to some but not all vertices of both H_1 and H_2. There is • a long-claw $\{v_1, h, u_{-1}, u_0, u_1, u_2, u_3\}$ when $p = 1$; • a net $\{v_1, h, u_0, u_1, u_3, u_2\}$ when $p = 2$; or • a long-claw $\{v_0, v_1, u_0, u_1, h, u_p, u_{p+1}\}$ when $p \ge 3$. □

To prove Thm. 2.2, we show that if there are two holes that are not congenial, then one of them is contained in a nontrivial module. This is impossible in a reduced graph, where every nontrivial module induces a clique. We construct this nontrivial module with the help of the following lemma, which shows that the common neighbors form a separator.

Lemma 6.3. *Let H be a hole that is the shortest among all the holes congenial to it in a prereduced graph G. Then $\widehat{N}(H)$separates $N[H] \setminus \widehat{N}(H)$ from $V(G) \setminus N[H]$.*

Proof. Suppose to the contrary, $N[H] \setminus \widehat{N}(H)$ and $V(G) \setminus N[H]$ are still connected in $G - \widehat{N}(H)$, then there are two adjacent vertices u and v such that $u \in N[H] \setminus \widehat{N}(H)$ and $v \in V(G) \setminus N[H]$. Note that $u \notin H$, and we have two adjacent vertices only one of which is adjacent to part of the hole H. Depending on the number of neighbors of u in H, we have either a long claw (when $|N_H(u)| = 1$), a net (when $|N_H(u)| = 2$), or a †-AW of size 7 (when $|N_H(u)| = 3$), none of which cannot exist in a prereduced graph. On the other hand, if $|N_H(u)| > 3$ then we can use u to find another hole H' that is strictly shorter than H; it is surely congenial to H, which contradicts the assumption. □

Proof. (of Thm. 2.2) Suppose, for contradicts, that not all holes are congenial to each other. By Lem. 6.2, being congenial is an equivalence relation. Hence there are two equivalence classes of holes, from each of which we pick a shortest one; let them be H_1 and H_2. Assume without loss of generality that H_2 has a vertex v not in $N[H_1]$. Lem. 6.3 implies that $\widehat{N}(H_1)$ separates $N[H_1] \setminus \widehat{N}(H_1)$ and $V(G) \setminus N[H_1]$. Either $\widehat{N}(H_1) = \emptyset$ and then G is disconnected where $N[H_1]$ induces a connected component ($v \notin N[H_1]$); or $\widehat{N}(H_1)$ is the neighbor of $N[H_1]$ and they are completely connected (Cor. 6.2). In either case, the set $N[H_1] \setminus \widehat{N}(H_1)$ is a nontrivial module that does not induce a clique. Reduction 2 is thus applicable and the graph is not reduced. □

7 Hole covers

A set of vertices is called a *hole cover* of a graph G if it intersects every hole in G, and the removal of any hole cover makes the graph chordal. A hole cover is minimal if any proper subset of it is not a hole cover. Any interval deletion set makes a hole cover of the input graph, and thus contains a minimal hole cover. The goal of this section is to prove Thm. 2.3, that is, to provide a polynomial bound on the number of minimal hole covers in a reduced graph and give a polynomial time algorithm to find all of them.

To simplify the task, observe that no minimal hole cover contains a vertex that is not in any hole.

Proposition 7.1. *Let $\mathcal{H}$ be the set of all holes in a reduced graph G, and G_0 be the subgraph induced by $\bigcup_{H\in\mathcal{H}} H$. A set HC of vertices is a minimal hole cover of G if and only if it is a minimal hole cover of G_0.*

In this section we will focus on the subgraph G_0 induced by the union of all holes in the reduced graph G. The subgraph G_0 has the same set of holes as G, and they remain pairwise congenial. Moreover, each vertex of G_0 is in the closed neighborhood of each hole H of G_0, which means G_0 is connected.

Proposition 7.2. ($\star$) *The subgraph $G_0 - HC$ is an interval graph for each hole cover HC of G_0.*

In particular, by Prop. 6.2, the subgraph $G_0 - N[v]$ is an interval graph for each vertex v of G_0. This suggests that G_0 might be a circular-arc graph. Recall that *circular-arc graphs* are a natural generalization of interval graphs, and they can be represented by arcs of a circle. In other words, G_0 can be obtained by gluing the ends of an interval graph together, whereupon minimal hole covers of G_0, except those containing ends vertices glued, coincide with minimal separators of the underlying interval graph.

In what follows we prove a series of claims on how the neighborhood of a vertex v of a hole H_1 looks like in another hole H_2. The first statement is a paraphrase of Cor. 6.3:

Corollary 7.1. *No vertex v of G_0 can be a common neighbor of any hole in G_0.*

Therefore, by definition of congenial holes and Prop. 6.1, we can assume that for every $v \in V(G_0)$ and hole H, we have that $N_H[v]$ is a proper nonempty subset of H and its vertices induce a path in H. Fixing any ordering of the vertices in H, we can denote two ends of the path as $\texttt{begin}_H(v)$ and $\texttt{end}_H(v)$ respectively; when $N_H[v]$ contains both h_0 and $h_{|H|-1}$, we number vertices of $N_H[v]$ as $\{h_{-p}, \dots, h_0, \dots, h_q\}$ where both p and q are nonnegative, and then $\texttt{begin}_H(v) = -p$, $\texttt{end}_H(v) = q$.

Proposition 7.3. *Let u, v be a pair of adjacent vertices of G_0. Their closed neighborhoods in any hole H intersect, and $N_H[u] \cup N_H[v] \neq H$.*

Proof. If either or both of u and v belong to H, then the first assertion holds vacuously and the second assertion follows from Prop. 6.1. Hence we assume $u, v \notin H$. We number vertices of H such that $N_H[u] = \{h_0, \dots, h_{\ell_1}\}$; the order can be either way if $|N_H[u]| = 1$, i.e., $\ell_1 = 0$. Let then $\{h_{\ell_2}, \dots, h_{\ell_3}\} = N_H[v]$.

Suppose first, for contradiction, $N_H[u] \cap N_H[v] = \emptyset$; we may assume then both ℓ_2 and ℓ_3 are positive, i.e., $\ell_1 < \ell_2 \leq \ell_3 < |H|$. If $\ell_3 \geq |H| - 3$, then $(uvh_{\ell_3} \dots h_{|H|}u)$ is a hole of length at most 6. Otherwise, $(uu_1h_{\ell_2} \dots h_{\ell_1}u)$ is a hole not congenial to H: in particular, the vertex $h_{|H|-2}$ in H is nonadjacent to it. In either case, we end with a contradiction; hence $N_H[u]$ and $N_H[v]$ must intersect.

Suppose to the contrary of the second assertion, $N_H[u] \cup N_H[v] = H$. Then v is adjacent to every vertex in $(h_{\ell_1+1}h_{\ell_1+2} \cdots h_{|H|-1})$. Prop. 6.1 and Cor. 7.1 imply $6 < \ell_1 < |H| - 6$. If $v \not\sim h_{\ell_1}$, then $(uh_{\ell_1}h_{\ell_1+1}vu)$ is a 4-hole. A symmetric argument applies when $v \not\sim h_0$. Now suppose v is adjacent to both h_0 and h_{ℓ_1}, then

$(uh_{\ell_1}h_{\ell_1+1}\cdots h_{|H|-1}h_0u)$ is a hole and v is a common neighbor of it (contradicting Cor. 7.1). None of the cases is possible, which proves this assertion. The proof is now completed. □

LEMMA 7.1. *Let H and H_1 be two holes in G_0. For any vertex $u_i \in H_1$, both $N_H[u_{i-1}]$ and $N_H[u_{i+1}]$, where u_{i-1} and u_{i+1} are neighbors of u_i in H_1, contains at least one end of $N_H[u_i]$.*

Proof. By symmetry, it suffices to show that it holds for $N_H[u_{i+1}]$. If $N_H[u_i]$ does not contain $N_H[u_{i+1}]$ as a proper subset, then it follows from Prop. 7.3. Hence we may assume $N_H[u_{i+1}] \subset N_H[u_i]$. Here we show a stronger statement, that is, for any pair of vertices u, v such that $N_H[v] = \{h_1, \cdots, h_\ell\}$ and $\{h_0, \cdots, h_{\ell+1}\} \subseteq N_H[u]$, it always holds that $N[v] \subset N[u]$. Noting that u_{i+1} necessarily has a neighbor that is nonadjacent to u_i (they are both in H_1), this lemma ensues.

Note that $N_H[v]$ is nonempty and thus $|N_H[u]| \geq 3$. Consider first that u or v is in H. If $u \in H$, then $N_h[v] = \{u\}$, and it follows from Prop. 7.3. Assume now $u \notin H$; the argument below holds regardless of whether $v \in H$ or not. Let x be any vertex in $N[v]$ different from $\{h_0, h_{\ell+1}\}$. According to Prop. 7.3, $N_H[x]$ must intersect $\{h_1, \cdots, h_\ell\}$. On the other hand, it is nonadjacent to $\{h_0, h_{\ell+1}\}$; otherwise (h_0uvxh_0) or $(h_{\ell+1}uvxh_0)$ is a 4-hole, which is impossible. Therefore, $N_H[x]$ is also a subset of $\{h_1, \cdots, h_\ell\}$, and the statement follows from Prop. 4.1 (the path is taken as $N_H[u]$ as well as its processor and successor). □

Minimal hole covers of G_0 are captured by

LEMMA 7.2. *Any minimal hole cover of G_0 induces a clique.*

Proof. Suppose to the contrary, there is a minimal hole cover HC that contains two nonadjacent vertices u and v. By the minimality of HC, there are two holes H_1 and H_2 such that $HC \cap H_1 = \{u\}$ and $HC \cap H_2 = \{v\}$. In particular, $u \notin H_2$ and $v \notin H_1$. We number the vertices of H_1 such that $N_{H_1}[v] = \{u_1, u_2, \cdots, u_p\}$. The union of $N_{H_2}[u_1]$ and $N_{H_2}[u_p]$ is consecutive set of vertices in H_2: they both contain v, and, by Prop. 6.1, are consecutive in H_2. We number the vertices of H_2 such that $u_1 \sim v_1$ and $N_{H_2}[u_1] \cup N_{H_2}[u_p] = \{v_1, \ldots, v_q\}$.

CLAIM 3. *At least one vertex of H_2 is adjacent to neither u_1 nor u_p.*

Proof. The claim follows from Prop. 7.3 when $p = 2$; hence we may assume $p > 2$, which means $u_1 \not\sim u_p$ (note that $u_0 \neq u_{p+1}$). Suppose $N_{H_2}[u_1] \cup N_{H_2}[u_p] = H_2$, then by Prop. 6.1, we have $7 < |N_{H_2}[u_1]| < |H_2| - 7$, which means at least one end of the path induced by $N_{H_2}[u_1]$ is not adjacent to v. Without loss of generality, let it be v_i where $i = \texttt{end}_{H_2}[u_1]$; noting that by assumption $v_{i+1} \sim u_p$, there is either a 4-hole $(vu_1v_iu_pv)$ (if $v_i \sim u_p$) or a 5-hole $(vu_1v_iv_{i+1}u_pv)$ (if $v_i \not\sim u_p$). ⌟

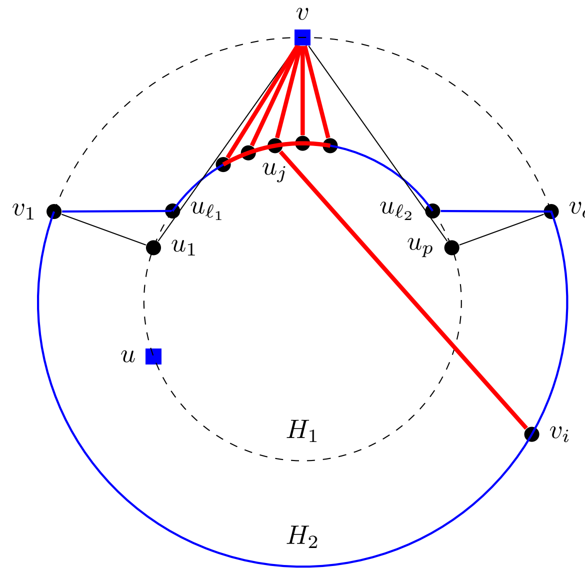

Figure 4: In two congenial holes each covered by a single vertex, there is an uncovered hole.

In what follows we show the existence of a hole in $G - HC$, which contradicts the assumption that HC is a hole cover and thus proves this lemma. Denote by $P_1 = (u_1u_2\ldots u_p)$ and $P_2 = (v_qv_{q+1}\ldots v_0v_1)$. By definition $u \notin P_1$; to show $v \notin P_2$ it suffices to rule out the possibility that $v \in \{v_1, v_q\}$, as by the numbering of H_2, v is in $\{v_1, v_1, \ldots, v_q\}$. According to Lem. 7.1, the two neighbors of v in H_2 are adjacent to either u_1 or u_p; however, by the claim, neither v_0 nor v_{q+1} is adjacent to u_1 or u_p. We now argue that each inner vertex v_i of P_2 is not adjacent to P_1 (see the thick edges in Fig. 4). Suppose to the contrary, v_i is adjacent to P_1. Noting that $v_i \not\sim u_1$, $v_i \not\sim u_p$, and $u_1 \neq u_{p+1}$, Prop. 4.1(3) applies, and we can conclude $v_i \sim v$, which is impossible. (It is immaterial whether $v_i \in H_1$ or not.) Now we construct the hole in $G - HC$ as follows. Claim 1 implies that the length of P_2 is at least 2. If $u_1 \sim v_q$, then $(u_1P_2u_1)$ is such a hole. Otherwise by assumption we have $u_p \sim v_q$. Let $\ell_1 = \max\{i|u_i \sim v_1 \text{ and } 0 \leq i \leq p\}$, and $\ell_2 = \min\{i|u_i \sim v_q \text{ and } \ell_1 \leq i \leq p\}$. Then $(u_{\ell_1}u_{\ell_2}P_2u_{\ell_1})$ will be such a hole (see the solid hole Fig. 4). □

Recall that any hole in a reduced graph contains more than 10 vertices, while a minimal hole cover is a clique and hence local. We can easily derive

LEMMA 7.3. *For any minimal hole cover HC of G_0 and any shortest hole H, there is a vertex $v \in H$ such that $N_{G_0}[v] \not\sim HC$.*

Proof. We show this by construction. By Cors. 6.1 and 7.1, each vertex in G_0 has at most 3 neighbors in H. By Lem. 7.2, HC is a clique and hence $|H \cap HC| \leq 2$. We number the vertices of H in a way that $h_0 \in HC$ and $h_1 \notin HC$, and claim that $v = h_5$ is the asserted vertex. Suppose to the contrary, $N_{G_0}[h_5]$ and HC are adjacent, then there is an h_0-h_5 path P of length at most 3 and all its inner vertices belong to G_0. The case $P = h_0 v h_5$ is impossible, as by Prop. 6.1 and Cor. 6.1, v is adjacent to at most 3 consecutive vertices in H. Now we may assume $P = h_0 v_1 v_2 h_5$, and examine the neighbors of v_1 and v_2 in H. By Cor. 6.1, we have $\texttt{end}_H(v_1) \leq 2$ and $\texttt{begin}_H(v_2) \geq 3$. This means that there is a hole $(v_1 h_i h_{i+1} \dots h_j v_2 v_1)$, where $i = \texttt{end}_H(v_1)$ and $j = \texttt{begin}_H(v_2)$, of length at least 4 and at most 8. □

We now relate minimal hole covers of G_0 to minimal separators in some interval subgraphs. In one direction of the proof, we need the following claim. Observe that in an interval representation of a connected interval graph, the union of all the intervals also forms an interval. Similarly, if there is a point p in the real line such that there are intervals not containing p both to the left and to the right of p, then the set of intervals containing p is a clique separator.

PROPOSITION 7.4. *Let v be a vertex in an interval graph G. If v is not adjacent to any simplicial vertex, then $N[v]$ is a separator of G.*

Proof. We consider an interval representation of G. Without loss of generality, we assume that no two intervals have the same ends. Denote by x the interval with the smallest right end, and y the interval with the largest left end. It is easy to see that x and y are simplicial. If $x \sim y$, then the graph is a complete graph (every interval contains the interval between the left end of y and the right end of x); thus every vertex is adjacent to a simplicial vertex, and the assertion is vacuously true. Therefore, we can assume $x \not\sim y$, and let p be an arbitrary point in interval v. By assumption v is not adjacent to x or y, which means that x is to the left of p and y is to the right of p. As every interval that contains p is in $N[v]$, in the subgraph $G - N[v]$ that contains x and y, no interval contains p; hence x and y are disconnected. In other words, $N[v]$ is an x-y separator. □

According to Lem. 7.3, every minimal hole cover satisfies the condition in the following lemma; hence the lemma applies to all of them. Note that $G_0 - N_{G_0}[v]$ is the same as $G_0 - N[v]$.

LEMMA 7.4. *Let v be a vertex in a shortest hole H of G_0, and X induce a clique nonadjacent to $N_{G_0}[v]$. Set X forms a minimal hole cover of G_0 if and only if X is a minimal separator of $G_0 - N[v]$.*

Proof. It suffices to show that X is a hole cover of G_0 if and only if it is a separator of $G_0 - N[v]$.

⇒ Clearly, each connected component in $G_0 - N[v]$ contains a neighbor of $N[v]$. As X is not adjacent to $N[v]$, the set X cannot fully contain a component of $G_0 - N[v]$, which implies that the number of connected components of $G_0 - N[v] - X$ is no less than that of $G_0 - N[v]$. Therefore, if $G_0 - N[v]$ is not connected, then neither is $G_0 - N[v] - X$, and X makes a trivial separator for $G_0 - N[v]$. In the following argument of this direction we may assume $G_0 - N[v]$ is connected, and it suffices to show that $G_0 - N[v] - X$ is not connected. By Prop. 7.2, $G_0 - X$ is an interval subgraph; as G_0 itself contains no simplicial vertex, any vertex x that is simplicial in $G_0 - X$ must be a neighbor of X: otherwise $N_{G_0 - X}(x) = N_{G_0}(x)$ and cannot be a clique. As $N[v]$ is not adjacent to X by assumption, v is not adjacent to any simplicial vertex of the interval graph $G_0 - X$. Therefore, according to Prop. 7.4, the removal of $N[v]$ disconnects $G_0 - X$. This finishes the proof of the "only if" direction.

⇐ Let us start from a close scrutiny of $G_0 - N[v]$. According to Prop. 6.1, the removal of $N[v]$ transforms each hole into a path of length at least 7; In particular, let P be the path induced by $H \setminus N_H[v]$. In the argument to follow, we show that ends of each such path are connected to the ends of P respectively; the further removal of X separates each path into at most two sub-paths; hence if there is a hole disjoint from X, then the path left by it is able to connect every sub-path and thereby every vertex, which is impossible.

We number the vertices of H such that $v = v_0$; hence $N_H[v] = \{v_{-1}, v_0, v_1\}$ and the ends of P are v_{-2} and v_2. Let H_2 be another hole, and P_2 be the path left. First of all, we show that the two ends of P_2 are adjacent to $\{v_1, v_2\}$ and $\{v_{-1}, v_{-2}\}$, respectively. Number the vertices of H_2 such that $N_{H_2}[v] = \{h_1, \dots, h_p\}$; hence the ends of P_2 are h_0 and h_{p+1}. By Cor. 6.1, $N_H[h_1] \subset \{v_{-2}, v_{-1}, v_0, v_1, v_2\}$, and according to Prop. 7.3, h_0 is adjacent to either $\{v_1, v_2\}$ or $\{v_{-1}, v_{-2}\}$; a symmetric argument works for h_{p+1}. To show they cannot be adjacent to the same end of P, note

- By Lem. 7.1, h_0 and h_{p+1} cannot be both adjacent to v_1.

- Suppose h_0 and h_{p+1} are both adjacent to v_2, then $v_2 \notin H_2$, and we can apply Prop. 4.1 on v_2, v_0, and path $(h_{-1}h_0h_1 \dots h_p h_{p+1} h_{p+2})$ to conclude $v_0 \sim v_2$, which is impossible.

- Suppose h_0 and h_{p+1} are adjacent to v_1 and v_2, respectively. Without loss of generality, $h_0 \sim v_1$ and $h_{p+1} \sim v_2$. Clearly $h_p \neq v_2$ as they have different adjacencies to v_0; likewise, $h_{p+1} \neq v_1$ and $h_{p+1} \neq v_2$. We exclude $h_p = v_1$: then $p = 1$ by $h_0 \sim v_1$, and $v_0 \sim h_0$ by Lem. 7.1, which contradicts the numbering of H_2. Then $h_p \sim v_2$, as otherwise there is a hole $(v_2v_1v_0h_ph_{p+1}v_2)$ or $(v_2v_1h_ph_{p+1}v_2)$; likewise, $h_p \sim v_1$. It follows that, by Cor. 6.1, $h_p \not\sim v_{-1}$, and $p > 1$. On the other hand, $h_0 \not\sim v_{-1}$, as otherwise there is a hole $(h_0v_{-1}v_0v_1h_0)$. We can apply Prop. 4.1 on v_1, v_{-1}, and path $(h_{-1}h_0h_1 \dots h_p v_{p+1})$ to conclude $v_{-1} \sim v_1$, which is impossible.

A symmetric argument works for $\{v_{-1}, v_{-2}\}$, hence we may assume without loss of generality, $h_{p+1} \sim \{v_1, v_2\}$, and then $h_0 \sim \{v_{-1}, v_{-2}\}$. Let ℓ_1 be the smallest index such that $\ell_1 > p$ and $h_{\ell_1} \in N[v_2]$; for its existence, observe that $\ell_1 = p + 1$ if $h_{p+1} \in N[v_2]$, otherwise by Lem. 7.1, $h_{\mathtt{end}_{H_2}(v_1)}$ must be in $N[v_2]$. By construction, every vertex in $\{h_{p+1} \dots, h_{\ell_1}\}$ is adjacent to $N[v]$, thereby no in X. Symmetrically, we can define ℓ_2 to be the largest index such that $\ell_2 \leq 0$ and $h_{\ell_2} \in N[v_{-2}]$.

The argument above applies to every hole in G_0, from which we can also identify such a path. If X intersects every such path, then it makes a hole cover, and we are done; hence we assume otherwise. Let H_1 be a hole disjoint from X, and P_1 be the path induced by $H_1 \setminus N_{H_1}[v]$. As for each path left from a hole, X does not intersect either of its ends, and as X is a clique, after the removal of X, it either remains intact, or is separated into two sub-paths. The two sub-paths are adjacent to either v_{-2} or v_2. On the other hand, both v_{-2} and v_2 are adjacent to P_1. Therefore, all vertices are connected, contradicting the assumption that X is a separator of $G_0 - N[v]$. This finishes the proof. □

We are now ready to prove Thm. 2.3.

Proof. (of Thm. 2.3) Let G_0 be induced by the union of the holes of G. On the one hand, according to Lems. 7.3 and 7.4, each minimal hole cover of G corresponds to a minimal separator of $G_0 - N[v]$ for some vertex v of a shortest hole H. On the other hand, there are at most n minimal separators in $G_0 - N[v]$ for each vertex $v \in H$, which implies a quadratic bound for the total number of minimal hole covers of G. To enumerate them, we try every vertex $v \in H$ and enumerate all minimal separators of $G_0 - N[v]$. □

8 Caterpillar decompositions

This section proves Thm. 2.4 by providing the claimed algorithm for INTERVAL DELETION on nice graphs. Recall that a nice graph is chordal and contains no small AW, and every shallow terminal in a nice graph is simplicial; nice graphs are hereditary. Our algorithm finds an AW satisfying a certain minimality condition, from which we can construct a set of 10 vertices that intersects some minimum interval deletion set. Hence it branches on deleting one of these 10 vertices. The set of all shallow terminals, denoted by $ST(G)$, can be found in polynomial time as follows. For each triple of vertices, we check whether or not they forms the terminals for an AW. If yes, then one of them is necessarily shallow. The following lemma ensures that all shallow terminals can be found as such.

PROPOSITION 8.1. ($\star$) *In a nice graph, all AWs with the same set of terminals have the same shallow terminal.*

It should be noted that this does not rule out the possibility of a vertex being a base terminal of an AW and the shallow terminal of another AW If this happens, these AWs necessarily have at least one different terminal. Recall that by Thm. 2.1, every vertex in $ST(G)$ is simplicial in G. For each †- or ‡-AW, its shallow terminal is in $ST(G)$ by definition, its base terminals might or might not be in $ST(G)$, and all non-terminal vertices cannot be in $ST(G)$ (as they are not simplicial). From Lem. 5.2 we can derive

PROPOSITION 8.2. *Let s be a shallow terminal in a nice graph. There is an AW of which every base vertex is adjacent to all vertices of $N(s) \setminus ST(G)$.*

Proof. Let W be an AW with shallow terminal s and shortest possible base. Applying Lem. 5.2 on any vertex $x \in N(s) \setminus ST(G)$ and W, it cannot be in category "partial" by the minimality of W. Vertex x cannot be in category "none" either, otherwise x is a shallow terminal, contradicting $x \in N(s) \setminus ST(G)$. Thus every vertex in $N(s) \setminus ST(G)$ is in category "full." □

Now that the graph is chordal, it makes sense to discuss its clique tree, which shall be the main structure of this section. No generality will be lost by assuming G is connected. Since no inner vertex of a shortest path can be simplicial, the removal of simplicial vertices will not disconnect a connected graph; hence $G - ST(G)$ is a connected interval graph. This observation suggests a clique tree of G with a very nice structure. A *caterpillar (tree)* is a tree that consists of a central path and all other vertices are leaves connected to it.

Proposition 8.3. ($\star$) *In polynomial time we can build a clique tree $\mathcal{T}$ for a connected nice graph G such that*

- *$\mathcal{T}$ is a caterpillar;*
- *every shallow terminal of G appears only in one leaf node of $\mathcal{T}$; and*
- *every other vertex appears in some node of the central path (possibly leaf nodes as well) of $\mathcal{T}$.*

Within a caterpillar decomposition, we number the nodes in the central path as $K_0, K_1, \ldots$. By Prop. 8.3 and the definition of clique trees, each vertex not in $ST(G)$ is contained in some consecutive nodes of the central path. For each vertex $v \notin ST(G)$, we denote by $\texttt{first}(v)$ and $\texttt{last}(v)$ the smallest and, respectively, largest indices of nodes that contain v. In any †- or ‡-AW, every vertex of the base is non-simplicial, hence belongs to the central path of the caterpillar decomposition. By assumption, $d = |B| \geq 3$ and $b_1 \not\sim b_d$; as a result, the nodes that contain b_1 and b_d are disjoint. When numbering the vertices of the base, we follow the convention that $\texttt{last}(b_1) < \texttt{first}(b_d)$, i.e., base B goes "from left to right." Given a numbering of the base, the base terminals l and r can be distinguished from each other based on their adjacency with b_1 and b_d. Similarly, in the case of a ‡-AW, the centers c_1 and c_2 can be distinguished from each other, as they have different adjacency relations with l and r.

By observing the adjacencies and nonadjacencies between vertices of an AW and their possible positions in an interval representation of $G - ST(G)$, the following is straightforward and hence stated here without proof. In order to avoid pointless repetition, we are again using the same generalized notation for both †- and ‡-AW as stipulated in §4.

Proposition 8.4. *For any chordless path B that is disjoint from $ST(G)$ and has length ≥ 2, it holds that*

$$\text{first}(b_i) \leq \text{last}(b_{i-1}) < \text{first}(b_{i+1}) \leq \text{last}(b_i) < \text{first}(b_{i+2}).$$

Nodes that contain non-terminal vertices of an AW appear consecutively in the central path of $\mathcal{T}$. We would like to identify a minimum set of consecutive nodes whose union contains all non-terminal vertices of the AW.

Definition 8.1. *We define $\amalg[p,q] = \bigcup_{p \leq i \leq q} K_i$ for a pair of indices $p \leq q$, and $\amalg(W) = \amalg[\texttt{last}(b_1), \texttt{first}(b_d)]$ for an AW W. Set $\amalg(W)$ will be referred to as the* container *of W, and we say it is* minimal *if there exists no AW W' such that $\amalg(W') \subset \amalg(W)$.*

Let us observe that every base vertex of W appears in $\amalg(W)$ and no shorter subsequence of nodes contain every base vertex. Moreover, the following proposition shows that the centers also appear in $\amalg(W)$ (recall that $\widehat{N}(B)$ is the set of common neighbors of B and every center is in $\widehat{N}(B)$).

Proposition 8.5. $K_{\texttt{last}(b_1)} \cap K_{\texttt{first}(b_d)} = \widehat{N}(B)$.

Proof. By definition, a vertex of the left side is in K_i for every $\texttt{last}(b_1) \leq i \leq \texttt{first}(b_d)$, and thus belongs to $\widehat{N}(B)$. On the other hand, if a vertex v does not belong to the left side, then either $\texttt{first}(v) > \texttt{last}(b_1)$ or $\texttt{last}(v) < \texttt{last}(b_d)$, which implies $v \not\sim b_1$ or $v \not\sim b_d$ respectively. In either case, we have $v \notin \widehat{N}(B)$. □

In §6, we considered holes of the shortest length and observed that a vertex sees either all or at most 3 vertices in such a hole. Here for an AW whose container is minimal and base consists of the inner vertices of a shortest l-r path specified below, we can observe an analogous statement about the number of base vertices a vertex can see.

Definition 8.2. *Let $W = (s : c_1, c_2 : l, B, r)$ be an AW in a nice graph such that $\amalg(W)$ is minimal. We say B is a* short *base if (lBr) is a shortest l-r path in the subgraph induced by $\left(\amalg(W) \setminus \widehat{N}(B)\right) \cup \{l, r\}$.*

The following lemma shows that if the base is not short, then we can get an AW with a shorter base. In particular, this implies that a vertex of $\amalg(W) \setminus \widehat{N}(B)$ can see at most 3 consecutive vertices of the base.

Lemma 8.1. *Let $W = (s; c_1, c_2; l, B, r)$ be an AW such that $\amalg(W)$ is minimal. Then there is an W' such that $\amalg(W') = \amalg(W)$ and W has a short base.*

Proof. We show that if (lPr) is a chordless l-r path in the subgraph induced by $\left(\amalg(W) \setminus \widehat{N}(B)\right) \cup \{l, r\}$, then we can replace the base B of W by P to obtain another AW $W_P = (s : c_1, c_2 : l, P, r)$. Clearly the center(s) of W belong to $\widehat{N}(B)$, thereby adjacent to every other vertex in $\amalg(W)$, and hence to P. It is also easy to verify that no vertex in $\amalg(W) \setminus \widehat{N}(B)$ is adjacent to s: if such a vertex exists, then Lem. 5.2 classifies it as "partial" with respect to W, hence there is another AW W' such that $B' \subset B$ and $\amalg(W') \subset \amalg(W)$, which contradicts the minimality of $\amalg(W)$. Therefore, W_p is indeed an AW. Letting b'_1 and $b'_{d'}$ be the first and, respectively, last vertices of P, the selection of P implies $\texttt{last}(b'_1) \geq \texttt{last}(b_1)$ and $\texttt{first}(b'_{d'}) \leq \texttt{first}(b_d)$, hence $\amalg(W_P) \subseteq \amalg(W)$; as the latter is already minimal, they must be equal. Therefore, if the base of W is not short, then we can find another AW with the same container and shorter base. Applying this argument repeatedly will eventually procure an AW with the same container and having a short base. □

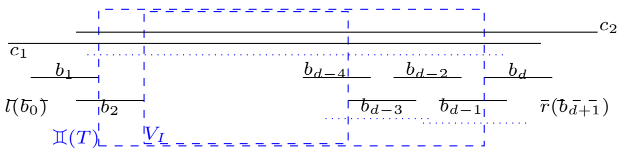

Figure 5: Non-terminal vertices of a leftmost minimal AW, represented by intervals.
base terminals are dashed as they might be in $ST(G)$; possible vertices in another AW are dotted.

With all pertinent definitions and observations, we are now ready to present the main lemma of this section which justifies our branching rule. Without an upper bound on the number of vertices in an AW—in particular, the length of its base can be arbitrarily long—trying each vertex in it cannot be done in FPT time. Thus we have to avoid most but a (small) constant number of base vertices to procure the claimed algorithm. To further decrease the number of vertices we need to consider, observing that the central path of the caterpillar decomposition has a linear structure, we start from the *leftmost* minimal container. By definition, minimal containers cannot properly contain each other, and thus the one with smallest begin-index also has the smallest end-index. In particular, the leftmost minimal container is unique, though it might be observed by more than one AWs, and can be identified in polynomial time. With this additional condition, if another AW intersects $\mathbb{I}(W)$, it has to come "from the right."

Let W be an AW of leftmost minimal container and having a short base. We claim that there is a minimum interval deletion set that breaks W in a canonical way: it contains either one of a constant number of specific vertices of W, or a specific minimum separator (details are given below) breaking the base of W. Therefore, by branching into a constant number of directions, we can guess one vertex of this interval deletion set. Lem. 8.2 below presents this result in a way that allows us to branch into at most 14 directions; using this lemma would result in a $14^k \cdot n^{O(1)}$ time algorithm. With further technical work, the number of branches can be reduced to 10, which is needed to achieve the running time $10^k \cdot n^{O(1)}$ claimed in Thm. 1.1. The proof of this improvement is deferred to the full version: we believe that the underlying ideas are easier to understand in the simplified proof below, and it is already sufficient to establish fixed-parameter tractability.

For each $\texttt{last}(b_1) \le i < \texttt{first}(b_{d-1})$, let us define $S_i = K_i \cap K_{i+1}$ to be the *ith separator.* Note that S_i contains $\widehat{N}(B)$ as a proper subset.

Lemma 8.2. *Let $\mathcal{T}$ be a caterpillar decomposition of a nice graph G, and $W = (s : c_1, c_2 : l, B, r)$ be an AW in G such that (1) $\mathbb{I}(W)$ is minimal; (2) B is a short base; and (3) $\texttt{first}(b_d)$ is the smallest among all AWs.*

Let ℓ be the minimum index such that $\texttt{last}(b_2) \le \ell < \texttt{first}(b_{d-5})$ and the cardinality of S_ℓ is minimum among $\{S_i | \texttt{last}(b_2) \le i < \texttt{first}(b_{d-5})\}$. There is a minimum interval deletion set to G that either contains one of the 13 vertices

$$V_B = \{s, c_1, c_2, l, b_1, b_2, b_{d-5}, b_{d-4}, b_{d-3}, b_{d-2}, b_{d-1}, b_d, r\},$$

or the whole set $X = S_\ell \setminus N$, where $N = \widehat{N}(B)$.

Proof. We prove by construction. Let Q be any minimum interval deletion set; we may assume $Q \cap V_B = \emptyset$, and $X \not\subseteq Q$, as otherwise Q satisfies the asserted condition and we are finished. We claim $Q' = (Q \setminus V_I) \cup X$, where $V_I = \mathbb{I}[\texttt{last}(b_3), \texttt{first}(b_{d-6})] \setminus N$, is the desired interval deletion set, which fully contains X in particular. By definition of V_I, any vertex $z \in V_I$ is adjacent to some vertex b_i for $4 \le i \le d-7$, then as B is short and $z \notin N$, we have

$$\begin{aligned} &\text{first}(b_2) \le \text{last}(b_1) < \text{first}(z) \\ &\qquad \le \text{last}(z) < \text{first}(b_{d-4}) \le \text{last}(b_{d-5}). \end{aligned} \tag{8.1}$$

As G is chordal, all minimal forbidden induced subgraphs in G are AWs. To show that Q' makes an interval deletion set to G, it suffices to argue that if there exists an AW W' avoiding Q' then we can also find an AW, not necessarily the same as W', avoiding Q. Suppose $W' = (s' : c'_1, c'_2 : l', B', r')$ is an AW in $G - Q'$. By the construction of Q', this AW must intersect $V_I \setminus X$; let $u \in W' \cap (V_I \setminus X)$. Clearly, u can neither be s', as $u \notin ST(G)$, nor r', as otherwise according to Prop. 8.4, $\texttt{first}(b'_{d'}) < \texttt{first}(u) < \texttt{first}(b_d)$, contradicting the selection of W. The following claim further rules out the possibility that $u \in \{c'_1, c'_2\}$.

Claim 4. *For each vertex $v \in \mathbb{I}[0, \texttt{first}(b_{d-2})] \setminus N$, we have $\texttt{last}(v) < \texttt{first}(b_d)$, and $v \not\sim ST(G)$.*

Proof. By definition, if v is adjacent to B, then $v \sim b_i$ for some $i \le d-3$. If $v \sim b_d$, then B is not a short base, as there would be a a shorter (not necessarily chordless) l-r path $(l, \ldots, b_i, v, b_d, r)$. Therefore, $v \not\sim b_d$ and it follows that $\texttt{last}(v) < \texttt{first}(b_d)$. Suppose to the contrary of the second assertion, v is adjacent to the shallow terminal x of some AW W_1. We apply Lem. 5.2(2) on v and W_1. As $v \notin ST(G)$, it has to be in categories "full" or "partial." In either case, there exists an AW whose base is fully contained in $\mathbb{I}[\texttt{first}(v), \texttt{last}(v)]$, contradicting the selection of W. ⌟

Therefore, either $u = l'$ or $u \in B'$. Now we focus on the chordless path $l'B'r'$, which we shall refer to by P', and how it reaches u when going from r' to l'. Recall that every vertex of B' appears in the central path of the caterpillar decomposition.

Claim 5. $B' \cap N = \emptyset$.

Proof. Suppose the contrary and let x be a vertex in $B' \cap N$. By definition of N and (8.1), we have $\texttt{first}(x) < \texttt{first}(u) \leq \texttt{last}(u) < \texttt{last}(x)$. Then every neighbor of u, which is not in $ST(G)$ according to Claim 4, is thus adjacent to x. As x and u are both in the chordless path P', vertex u has to be one end of it. More specifically, $u = l'$ and $x = b'_1$. A further consequence is that u is the only vertex in $W' \cap V_I$: the argument above applies to any vertex $u' \in W' \cap V_I$, and thus $u' = l' = u$.

Now we show, for any vertex w in $X \setminus Q$, which is nonempty by assumption, it has the same neighbors as u in W', and hence $(s' : c'_1, c'_2 : w, B', r')$ is an AW in $G - Q$, contradicting the assumption that Q is an interval deletion set to G. Observe that any vertex in N is adjacent to both u and w.

- The assumption $w \notin N$ implies $w \not\sim s$.
- By the selection of W, we have $\texttt{last}(c'_i) \geq \texttt{first}(b'_d) \geq \texttt{first}(b_d)$ for both $i = 1, 2$. If c'_i, where $i = 1$ or 2, is adjacent to one of u and w, then (8.1) implies $\texttt{first}(c'_i) < \texttt{last}(b_{d-5})$; as B is short, c'_i must be in N, and then adjacent to both u and w.
- Vertex b'_1 $(= x)$ is in N, hence adjacent to w.
- By definition, $b'_3 \sim b'_2$ and $b'_3 \not\sim b'_1 (\in N)$ imply $\texttt{last}(b'_2) \geq \texttt{first}(b'_3) > \texttt{last}(b'_1) \geq \texttt{first}(b_d)$. On the other hand, $b'_2 \not\sim u$ implies $b'_2 \notin N$. Then as B is short, $\texttt{first}(b'_2) > \texttt{last}(b_{d-5})$. Therefore, from (8.1) we can conclude that $\texttt{first}(b'_i) > \texttt{last}(w)$ for $2 \leq i \leq d' + 1$, and thus $w \not\sim b'_i$. ⌟

Claim 6. $c'_2 \in N$.

Proof. As $u = b'_i$ for some $0 \leq i \leq d'$, Prop. 8.4 and (8.1) imply $\texttt{first}(b'_1)\,(\leq \texttt{last}(l')) \leq \texttt{last}(u) < \texttt{last}(b_{d-5})$. By Claim 5, b'_1 is not in N and adjacent to at most 3 vertices of B; thus $\texttt{last}(b'_1) < \texttt{first}(b_{d-2}) \leq \texttt{last}(b_{d-3})$. On the other hand, by the selection of W, we have $\texttt{last}(c'_2) \geq \texttt{first}(b'_d) \geq \texttt{first}(b_d)$. Therefore, c'_2 is adjacent to at least 4 vertices of B and is in N. ⌟

From Claim 6 we can conclude $c'_2 \sim u$ and then $u \in B'$. By Prop. 8.4, $\texttt{first}(b'_1) \leq \texttt{first}(u) < \texttt{first}(b_{d-4})$. Then from Claim 4 and the fact $l' \sim b'_1$, it can be inferred that $l' \notin ST(G)$. Now $\texttt{last}(l')$ is defined, and $\texttt{last}(l') < \texttt{first}(c'_2) \leq \texttt{last}(b_1)$; the selection of W implies $\texttt{first}(b'_{d'}) \geq \texttt{first}(b_d)$. Therefore, the l'-$b'_{d'}$ path $l'B'$ has to go through X, and we end with a contradiction. This verifies that Q' is an interval deletion set to G, and it remains to show that Q' is minimum, from which the lemma follows.

Claim 7. $|Q'| \leq |Q|$.

Proof. It will suffice to show that $Q \cap V_I$ makes a b_2-b_{d-5} separator in $G - N$, and then the claim ensues as

$$|Q'| = |Q \setminus V_I| + |X| \leq |Q \setminus V_I| + |Q \cap V_I| = |Q|.$$

Suppose to the contrary, there is a chordless b_2-b_{d-5} path P. We can extend P into an l-r path $P^+ = (lb_1Pb_{d-4}b_{d-3}b_{d-2}b_{d-1}b_dr)$, which is disjoint from Q and N. Within P^+ there is a chordless l-r path (lB_1r). By assumption, $\{s, c_1, c_2\} \cap Q = \emptyset$; every vertex in B_1 satisfies the condition Claim 4, and hence nonadjacent to s. Thus, $(s : c_1, c_2 : l, B_1, r)$ is an AW in $G - Q$, which is impossible. □

As shown in Lem. 8.3, set V_B can be further improved to contain 9 vertices, and hence we only need to fork into 10 branches.

Lemma 8.3. ($\star$) *Let $\mathcal{T}$ be a caterpillar decomposition of a nice graph G, and $W = (s : c_1, c_2 : l, B, r)$ be an AW in G such that (1) $\amalg(W)$ is minimal; (2) B is a short base; and (3) $\texttt{first}(b_d)$ is the smallest among all AWs.*

Let ℓ be the minimum index such that $\texttt{last}(b_1) \leq \ell < \texttt{first}(b_{d-2})$ and the cardinality of S_ℓ is minimum among $\{S_i | \texttt{last}(b_1) \leq i < \texttt{first}(b_{d-2})\}$. There is a minimum interval deletion set to G that either contains one of the 9 vertices

$$V_B = \{s, c_1, c_2, l, b_1, b_{d-2}, b_{d-1}, b_d, r\},$$

or the whole set $X = S_\ell \setminus N$, where $N = \widehat{N}(B)$.

To complete the proof of Thm. 2.4, we need one last piece of the jigsaw, i.e., to find the AW required by Lem. 8.3.

Proof. (of Thm. 2.4) Based on Lem. 8.3, it suffices to show how to find such an AW, and then the standard branching will deliver the claimed algorithm. For any triple of vertices $\{x, y, z\}$ and pair of indices $\{p, q\}$ for the nodes in the central path of the caterpillar decomposition, we can check whether or not there is an AW W whose terminals are $\{x, y, z\}$ and non-terminal vertices are fully contained in $\amalg[p, q]$. Therefore, in $O(n^6)$ time we are able to find the correct terminals and indices, from which the short base B can also be easily constructed. This finishes the construction of the AW required by Lem. 8.3. □

9 Concluding remarks

We have classified INTERVAL DELETION to be FPT by presenting a $c^k \cdot n^{O(1)}$ algorithm with $c = 10$. The constant c might be improvable, and let us have a brief discussion on how to achieve this. The current constant 10 comes from Reduction 1 and Thm. 2.4. The constant in Reduction 1 is not tight, and it can be replaced by 8. We choose the current number for the convenience for later argument; for example, if we do not break AWs of size 9 in preprocessing, then we have to use a far more complicated proof for Prop. 8.1. In other words, the real dominating step is to break ATs in nice graphs, where we need to branch into 10 cases. As a nice graph exhibits a linear structure, it might help to apply dynamic programming here. To further lower the constant c, we need to break small forbidden induced subgraphs in a better way then the brute-force in our algorithm. So a natural question is: Can it be $c = 2$?

It is known that CHORDAL COMPLETION can be solved in polynomial time if the input graph is a circular-arc graph [21] while INTERVAL COMPLETION remains NP-hard on chordal graphs [28]. It would be interesting to inquire the complexity of INTERVAL DELETION on chordal graphs and other graph classes. At least, can it be solved in polynomial time if the input graph is nice, which, if positively answered, would suggest that all the troubles are small forbidden subgraphs.

As having been explored in [27], we would also like to ask which other problems can be formulated as or reduced to INTERVAL DELETION and then solved with our algorithm. Both practical and theoretical consequences are worth further investigation.

Acknowledgment. We would like to thank one of the referees for a particularly careful reading of the manuscript, which leads to a number of improvements in the paper.

References

[1] Isolde Adler, Martin Grohe, and Stephan Kreutzer. Computing excluded minors. In *SODA*, pages 641–650. SIAM, 2008.

[2] Amotz Bar-Noy, Reuven Bar-Yehuda, Ari Freund, Joseph Naor, and Baruch Schieber. A unified approach to approximating resource allocation and scheduling. *Journal of the ACM*, 48(5):1069–1090, 2001.

[3] Seymour Benzer. On the topology of the genetic fine structure. *Proceedings of the National Academy of Sciences*, 45(11):1607–1620, 1959.

[4] Kellogg S. Booth and George S. Lueker. Testing for the consecutive ones property, interval graphs, and graph planarity using *PQ*-tree algorithms. *Journal of Computer and System Sciences*, 13(3):335 – 379, 1976. A preliminary version appeared in STOC 1975.

[5] Peter Buneman. A characterization of rigid circuit graphs. *Discrete Mathematics*, 9(4):205–212, 1974.

[6] Leizhen Cai. Fixed-parameter tractability of graph modification problems for hereditary properties. *Information Processing Letters*, 58(4):171–176, 1996.

[7] Yixin Cao. An efficient branching algorithm for interval completion. arXiv:1306.3181, 2013.

[8] Yixin Cao, Jianer Chen, and Yang Liu. On feedback vertex set: New measure and new structures. In *SWAT*, volume 6139 of *LNCS*, pages 93–104. Springer, 2010.

[9] Yixin Cao and Dániel Marx. Chordal editing is fixed-parameter tractable. Manuscript, 2013.

[10] Jianer Chen, Yang Liu, Songjian Lu, Barry O'Sullivan, and Igor Razgon. A fixed-parameter algorithm for the directed feedback vertex set problem. *Journal of the ACM*, 55(5):21:1–19, 2008. A preliminary version appeared in STOC 2008.

[11] Gabriel A. Dirac. On rigid circuit graphs. *Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg*, 25(1):71–76, 1961.

[12] Rodney G. Downey and Michael R. Fellows. *Parameterized Complexity*. Springer, 1999.

[13] Fedor V. Fomin, Daniel Lokshtanov, Neeldhara Misra, and Saket Saurabh. Planar F-deletion: Approximation and optimal FPT algorithms. In *FOCS*, pages 470–479. IEEE Computer Society, 2012.

[14] Delbert R. Fulkerson and Oliver A. Gross. Incidence matrices and interval graphs. *Pacific Journal of Mathematics*, 15(3):835–855, 1965.

[15] Tibor Gallai. Transitiv orientierbare graphen. *Acta Mathematica Academiae Scientiarum Hungaricae*, 18:25–66, 1967.

[16] György Hajós. (problem 65) Über eine art von graphen. *Internationale Mathematische Nachrichten*, 11, 1957.

[17] Haim Kaplan, Ron Shamir, and Robert E. Tarjan. Tractability of parameterized completion problems on chordal, strongly chordal, and proper interval graphs. *SIAM Journal on Computing*, 28(5):1906–1922, 1999. A preliminary version appeared in FOCS 1994.

[18] Ken-ichi Kawarabayashi. Planarity allowing few error vertices in linear time. In *FOCS*, pages 639–648. IEEE Computer Society, 2009.

[19] Ken-ichi Kawarabayashi and Bruce A. Reed. An (almost) linear time algorithm for odd cyles transversal. In *SODA*, pages 365–378. SIAM, 2010.

[20] David George Kendall. Incidence matrices, interval graphs and seriation in archaeology. *Pacific Journal of Mathematics*, 28:565–570, 1969.

[21] Ton Kloks, Dieter Kratsch, and C. K. Wong. Minimum fill-in on circle and circular-arc graphs. *Journal of Algorithms*, 28(2):272–289, 1998. A preliminary version appeared in ICALP 1996.

[22] C. G. Lekkerkerker and J. Ch. Boland. Representation of a finite graph by a set of intervals on the real line. *Fundamenta Mathematicae*, 51:45–64, 1962.

[23] John M. Lewis and Mihalis Yannakakis. The node-

deletion problem for hereditary properties is NP-complete. *Journal of Computer and System Sciences*, 20(2):219–230, 1980.

[24] Carsten Lund and Mihalis Yannakakis. The approximation of maximum subgraph problems. In *ICALP*, volume 700 of *LNCS*, pages 40–51. Springer, 1993.

[25] Dániel Marx. Chordal deletion is fixed-parameter tractable. *Algorithmica*, 57(4):747–768, 2010.

[26] Dániel Marx and Ildikó Schlotter. Obtaining a planar graph by vertex deletion. *Algorithmica*, 62(3-4):807–822, 2012.

[27] N. S. Narayanaswamy and R. Subashini. d-COS-R is FPT via interval deletion. To appear in IPEC 2013. Available online at arXiv:1303.1643., 2013.

[28] Sheng-Lung Peng and Chi-Kang Chen. On the interval completion of chordal graphs. *Discrete Applied Mathematics*, 154(6):1003–1010, 2006.

[29] Bruce A. Reed, Kaleigh Smith, and Adrian Vetta. Finding odd cycle transversals. *Operations Research Letters*, 32(4):299–301, 2004.

[30] Pim van 't Hof and Yngve Villanger. Proper interval vertex deletion. *Algorithmica*, 65(4):845–867, 2013.

[31] Yngve Villanger, Pinar Heggernes, Christophe Paul, and Jan Arne Telle. Interval completion is fixed parameter tractable. *SIAM Journal on Computing*, 38(5):2007–2020, 2009.

A Examples used in proofs

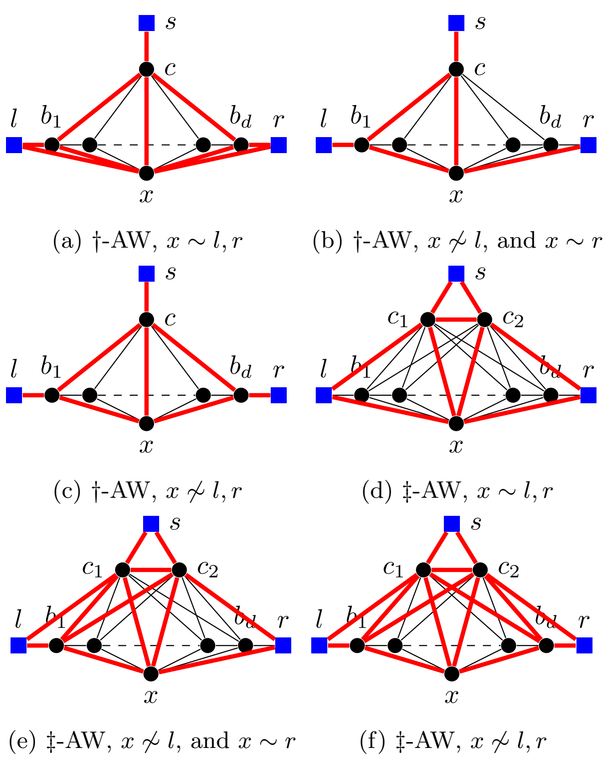

(a) †-AW, $x \sim l, r$ (b) †-AW, $x \not\sim l$, and $x \sim r$

(c) †-AW, $x \not\sim l, r$ (d) ‡-AW, $x \sim l, r$

(e) ‡-AW, $x \not\sim l$, and $x \sim r$ (f) ‡-AW, $x \not\sim l, r$

Figure 6: $x \in \widehat{N}(B)$ and s [Lem. 5.1].

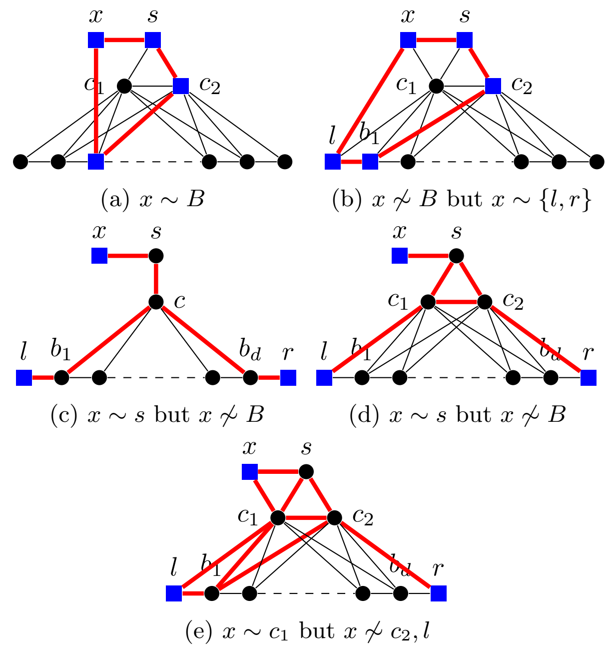

(a) $x \sim B$ (b) $x \not\sim B$ but $x \sim \{l, r\}$

(c) $x \sim s$ but $x \not\sim B$ (d) $x \sim s$ but $x \not\sim B$

(e) $x \sim c_1$ but $x \not\sim c_2, l$

Figure 7: $x \in N(s)$ and centers [Lem. 5.2].

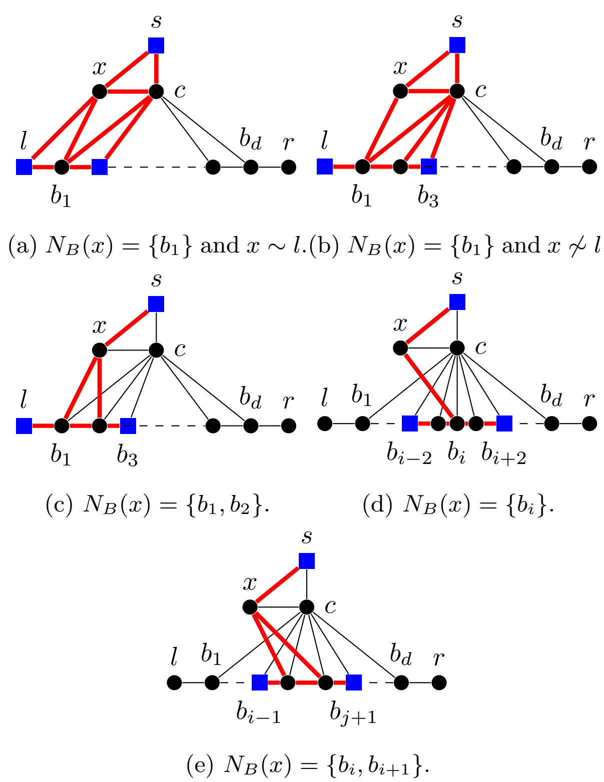

(a) $N_B(x) = \{b_1\}$ and $x \sim l$. (b) $N_B(x) = \{b_1\}$ and $x \not\sim l$.

(c) $N_B(x) = \{b_1, b_2\}$. (d) $N_B(x) = \{b_i\}$.

(e) $N_B(x) = \{b_i, b_{i+1}\}$.

Figure 8: Vertex x in category "partial" [Lem. 5.2].