Direct and Certifying Recognition of Normal Helly Circular-Arc Graphs in Linear Time*

Yixin Cao[1,2]

[1] Institute for Computer Science and Control, Hungarian Academy of Sciences
[2] Department of Computing, The Hong Kong Polytechnic University

Abstract. A normal Helly circular-arc graph is the intersection graph of arcs on a circle of which no three or less arcs cover the whole circle. Lin et al. [Discrete Appl. Math. 2013] presented the first recognition algorithm for this graph class by characterizing circular-arc graphs that are not in it. They posed as an open problem to design a direct recognition algorithm, which is resolved by the current paper. When the input is not a normal Helly circular-arc graph, our algorithm finds in linear time a minimal forbidden induced subgraph. Grippo and Safe [arXiv:1402.2641] recently reported the forbidden induced subgraphs characterization of normal Helly circular-arc graphs. The correctness proof of our algorithm provides, as a byproduct, an alternative proof to this characterization.

1 Introduction

This paper will be only concerned with simple undirected graphs. A graph is a *circular-arc graph* if its vertices can be assigned to arcs on a circle such that two vertices are adjacent iff their corresponding arcs intersect. Such a set of arcs is called a *circular-arc model* of this graph. If some point on the circle is not in any arc in the model, then the graph is an *interval graph*, and it can be represented by a set of intervals on the real line, which is called an *interval model*. Circular-arc graphs and interval graphs are two of the most famous intersection graph classes, and both have been studied intensively for decades. However, in contrast to the nice result of Lekkerkerker and Boland [5], characterizing circular-arc graphs by forbidden induced subgraphs remains a notorious open problem in this area.

The complication of circular-arc graphs should be attributed to two special intersection patterns of circular-arc models that are not possible in interval models. The first is two arcs intersecting in both ends, and a circular-arc model is called *normal* if no such pair exists. The second is a set of arcs intersecting pairwise but containing no common point, and a circular-arc model is called *Helly* if no such set exists. Normal and Helly circular-arc models are precisely those with no set of three or less arcs covering the whole circle [10,6]. A graph that admits such a model is called a *normal Helly circular-arc graph*.

One fundamental problem on a graph class is its recognition, i.e., to efficiently decide whether a given graph belongs to this class or not. For intersection graph

* Supported by ERC under the grant 280152 and OTKA under the grant NK105645.

J. Chen, J.E. Hopcroft, and J. Wang (Eds.): FAW 2014, LNCS 8497, pp. 13–24, 2014.

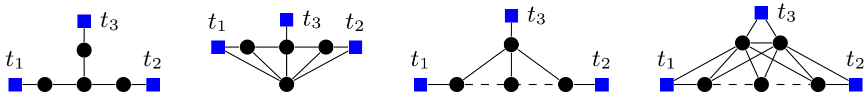

(a) long claw (b) whipping top (c) † (d) ‡

Fig. 1. Chordal minimal forbidden induced graphs

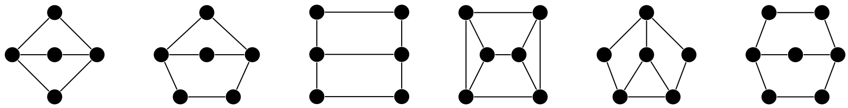

(a) $K_{2,3}$ (b) twin-C_5 (c) domino (d) $\overline{C_6}$ (e) FIS-1 (f) FIS-2

Fig. 2. Non-chordal and finite minimal forbidden induced graphs

classes, all recognition algorithms known to the author provide an intersection model when the membership is asserted. Most of them, on the other hand, simply return "NO" otherwise, while one might also want some verifiable certificate for some reason [9]. A recognition algorithm is *certifying* if it provides both positive and negative certificates. A minimal forbidden (induced) subgraph is arguably the simplest and most preferable among all forms of negative certificates [3].

For example, a graph is an interval graph iff it contains neither hole nor any graph in Fig. 1 [5]. Recall that a graph is *chordal* if it contains no holes. Kratsch et al. [4] reported a certifying recognition algorithm for interval graphs, which in linear time returns either an interval model of an interval graph or a forbidden induced subgraph for a non-interval graph. Although the forbidden induced subgraph returned by [4] is unnecessarily minimal, a minimal one can be easily retrieved from it (see [7] for another approach). Likewise, a minimal forbidden induced subgraph of chordal graphs, i.e., a hole, can be detected from a non-chordal graph in linear time [11]. However, although a circular-arc model of a circular-arc graph can be produced in linear time [8], it remains a challenging open problem to find a negative certificate for a non-circular-arc graph.

Indeed, all efforts attempting to characterize circular-arc graphs by forbidden induced subgraphs have been of no avail. For normal Helly circular-arc graphs, partial results were reported by [6], who listed all Helly circular-arc graphs that are not normal Helly circular-arc graphs. Very recently, Grippo and Safe completed this task by proving the following result. A wheel (resp., C^*) comprises a hole and another vertex completely adjacent (resp., nonadjacent) to it.

Theorem 1 ([2]). *A graph is a normal Helly circular-arc graph iff it contains no C^*, wheel, or any graph depicted in Figs. 1 and 2.*

It is easy to use definition to verify that a normal Helly circular-arc graph is chordal iff it is an interval graph. An interval model is always a normal and Helly circular-arc model, but an interval graph might have circular-arc model that is neither normal nor Helly, e.g., K_4. On the other hand,

Theorem 2 ([10,6]). *If a normal Helly circular-arc graph G is not chordal, then every circular-arc model of G is normal and Helly.*

These observations inspire us to recognize normal Helly circular-arc graphs as follows. If the input graph is chordal, it suffices to check whether it is an interval graph. Otherwise, we try to build a circular-arc model of it, and if success, verify whether the model is normal and Helly. Lin et al. [6] showed that this approach can be implemented in linear time. Moreover, if there exists a set of at most three arcs covering the circle, then their algorithm returns it as a certificate.

This algorithm, albeit conceptually simple, suffers from twofold weakness. First, it needs to call some recognition algorithm for circular-arc graphs, while all known algorithms are extremely complicated. Second, it is very unlikely to deliver a negative certificate in general. Therefore, Lin et al. [6] posed as an open problem to design a direct recognition algorithm for normal Helly circular-arc graphs, which would be desirable for both efficiency and the detection of negative certificates. The main result of this paper is the following algorithm—$n := |V(G)|$ and $m := |E(G)|$ are used throughout:

Theorem 3. *There is an $O(n+m)$-time algorithm that given a graph G, either constructs a normal and Helly circular-arc model of G, or finds a minimal forbidden induced subgraph of G.*

We remark that the proof of Thm. 3 will not rely on Thm. 1. Indeed, since our algorithm always finds a subgraph specified in Thm. 1 when the graph is not a normal Helly circular-arc graph, the correctness proof of our algorithm provides another proof of Thm. 1.

Let us briefly discuss the basic idea behind our disposal of a non-chordal graph G. If G is a normal Helly circular-arc graph, then for any vertex v of G, both $N[v]$ and its complement induce nonempty interval subgraphs. The main technical difficulty is how to combine interval models for them to make a circular-arc model of G. For this purpose we build an auxiliary graph $\mho(G)$ by taking two identical copies of $N[v]$ and appending them to the two ends of $G - N[v]$ respectively. The shape of symbol $\mho$ is a good hint for understanding the structure of the auxiliary graph. We show that $\mho(G)$ is an interval graph and more importantly, a circular-arc model of G can be produced from an interval model of $\mho(G)$. On the other hand, if G is not a normal Helly circular-arc graph, then $\mho(G)$ cannot be an interval graph. In this case we use the following procedure to obtain a minimal forbidden induced subgraph of G.

Theorem 4. *Given a minimal non-interval induced subgraph of $\mho(G)$, we can in $O(n+m)$ time find a minimal forbidden induced subgraph of G.*

The crucial idea behind our certifying algorithm is a novel correlation between normal Helly circular-arc graphs and interval graphs, which can be efficiently used for algorithmic purpose. This was originally proposed in the detection of small forbidden induced subgraph of interval graphs [1], i.e., the opposite direction of the current paper. In particular, in [1] we have used a similar definition of the auxiliary graph and pertinent observations. However, the main structural analyses, i.e., the detection of forbidden induced subgraphs, divert completely. For example, the most common forbidden induced subgraphs in [1] are 4- and 5-holes, which, however, are allowed in normal Helly circular-arc graphs. Their existence makes the interaction between $N[v]$ and $G - N[v]$ far more subtle, and thus the detection of minimal forbidden induced subgraphs in the current paper is significantly more complicated than that of [1].

2 The Recognition Algorithm

All graphs are stored as adjacency lists. We use the customary notation $v \in G$ to mean $v \in V(G)$, and $u \sim v$ to mean $uv \in E(G)$. Exclusively concerned with induced subgraphs, we use F to denote both a subgraph and its vertex set.

Consider a circular-arc model $\mathcal{A}$. If every point of the circle is contained in some arc in $\mathcal{A}$, then we can find an inclusive-wise minimal set X of arcs that cover the entire circle. If $\mathcal{A}$ is normal and Helly, then X consists of at least four arcs and thus corresponds to a hole. Therefore, a normal Helly circular-arc graph G is chordal iff it is an interval graph, for which it suffices to call the algorithms of [4,7]. We are hence focused on graphs that are not chordal. We call the algorithm of Tarjan and Yannakakis [11] to detect a hole H.

Proposition 1. *Let H be a hole of a circular-arc graph G. In any circular-arc model of G, the union of arcs for H covers the whole circle, i.e., $N[H] = V(G)$.*

Indices of vertices in H should be understood as modulo $|H|$, e.g., $h_0 = h_{|H|}$. By Prop. 1, every vertex should have neighbors in H. We use $N_H[v]$ as a shorthand for $N[v] \cap H$, regardless of whether $v \in H$ or not. We start from characterizing $N_H[v]$ for every vertex v: we specify some forbidden structures not allowed to appear in a normal Helly circular-arc graph, and more importantly, we show how to find a minimal forbidden induced subgraph if one of these structures exists. The fact that they are forbidden can be easily derived from the definition and Prop. 1. Due to the lack of space, their proofs, mainly on the detection of minimal forbidden induced subgraphs, are deferred to the full version.

Lemma 1. *For every vertex v, we can in $O(d(v))$ time find either a proper sub-path of H induced by $N_H[v]$, or a minimal forbidden induced subgraph.*

We designate the ordering $h_0, h_1, h_2, \cdots$ of traversing H as *clockwise*, and the other *counterclockwise*. In other words, edges h_0h_1 and h_0h_{-1} are clockwise and counterclockwise, respectively, from h_0. Now let P be the path induced by $N_H[v]$. We can assign a direction to P in accordance to the direction of H, and then we have clockwise and counterclockwise ends of P. For technical reasons, we assign canonical indices to the ends of the path P as follows.

Definition 1. *For each vertex $v \in G$, we denote by $\texttt{first}(v)$ and $\texttt{last}(v)$ the indices of the counterclockwise and clockwise, respectively, ends of the path induced by $N_H[v]$ in H satisfying*

- $-|H| < \texttt{first}(v) \leq 0 \leq \texttt{last}(v) < |H|$ *if* $h_0 \in N_H[v]$; *or*
- $0 < \texttt{first}(v) \leq \texttt{last}(v) < |H|$, *otherwise.*

It is possible that $\texttt{last}(v) = \texttt{first}(v)$, when $|N_H[v]| = 1$. In general, $\texttt{last}(v) - \texttt{first}(v) = |N_H[v]| - 1$, and $v = h_i$ or $v \sim h_i$ for each i with $\texttt{first}(v) \leq i \leq \texttt{last}(v)$. The indices $\texttt{first}(v)$ and $\texttt{last}(v)$ can be easily retrieved from Lem. 1, with which we can check the adjacency between v and any vertex $h_i \in H$ in constant time. Now consider the neighbors of more than one vertices in H.

Lemma 2. *Given a pair of adjacent vertices u, v s.t. $N_H[u]$ and $N_H[v]$ are disjoint, then in $O(n+m)$ time we can find a minimal forbidden induced subgraph.*

Lemma 3. *Given a set U of two or three pairwise adjacent vertices such that 1) $\bigcup_{u \in U} N_H[u] = H$; and 2) for every $u \in U$, each end of $N_H[u]$ is adjacent to at least two vertices in U, then we can in $O(n+m)$ time find a minimal forbidden induced subgraph.*

Let $T := N[h_0]$ and $\overline{T} := V(G) \setminus T$. As we have alluded to earlier, we want to duplicate T and append them to different sides of $\overline{T}$. Each edge between $v \in T$ and $u \in \overline{T}$ will be carried by only one copy of T, and this is determined by its direction specified as follows. We may assume that none of the Lems. 1, 2, and 3 applies to v or/and u, as otherwise we can terminate the algorithm by returning the forbidden induced subgraph found by them. As a result, u is adjacent to either $\{h_{\texttt{first}(v)}, \cdots, h_{-1}\}$ or $\{h_1, \cdots, h_{\texttt{last}(v)}\}$ but not both. The edge uv is said to be clockwise from T if $u \sim h_i$ for $1 \leq i \leq \texttt{last}(v)$, and counterclockwise otherwise. Let E_{c}(resp., E_{cc}) denote the set of edges clockwise (resp., counterclockwise) from T, and let T_{c} (resp., T_{cc}) denote the subsets of vertices of T that are incident to edges in E_{c} (resp., E_{cc}). Note that $\{E_{\text{cc}}, E_{\text{c}}\}$ partitions edges between T and $\overline{T}$, but a vertex in T might belong to both T_{cc} and T_{c}, or neither of them. We have now all the details for the definition and construction of the auxiliary graph $\mho(G)$, which can be done in linear time.

Definition 2. *The vertex set of $\mho(G)$ consists of $\overline{T} \cup L \cup R \cup \{w\}$, where L and R are distinct copies of T, i.e., for each $v \in T$, there are a vertex v^l in L and another vertex v^r in R, and w is a new vertex distinct from $V(G)$. For each edge $uv \in E(G)$, we add to the edge set of $\mho(G)$*

- *an edge uv if neither u nor v is in T;*
- *two edges $u^l v^l$ and $u^r v^r$ if both u and v are in T; or*
- *an edge uv^l or uv^r if $uv \in E_{\text{c}}$ or $uv \in E_{\text{cc}}$ respectively ($v \in T$ and $u \in \overline{T}$).*

Finally, we add an edge wv^l for every $v \in T_{\text{cc}}$.

Lemma 4. *The numbers of vertices and edges of $\mho(G)$ are upper bounded by $2n$ and $2m$ respectively. Moreover, an adjacency list representation of $\mho(G)$ can be constructed in $O(n+m)$ time.*

In an interval model, each vertex v corresponds to a closed interval $I_v = [\texttt{lp}(v), \texttt{rp}(v)]$. Here $\texttt{lp}(v)$ and $\texttt{rp}(v)$ are the left and right *endpoints* of I_v respectively, and $\texttt{lp}(v) < \texttt{rp}(v)$. We use unit-length circles for circular-arc models, where every point has a positive value in $(0, 1]$. Each vertex v corresponds to a closed arc $A_v = [\texttt{ccp}(v), \texttt{cp}(v)]$. Here $\texttt{ccp}(v)$ and $\texttt{cp}(v)$ are counterclockwise and clockwise endpoints of A_v respectively; $0 < \texttt{ccp}(v), \texttt{cp}(v) \leq 1$ and they are assumed to be distinct. It is worth noting that possibly $\texttt{cp}(v) < \texttt{ccp}(v)$; such an arc necessarily contains the point 1.

Lemma 5. *If G is a normal Helly circular-arc graph, then $\mho(G)$ is an interval graph.*

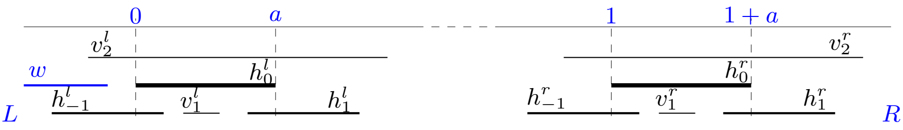

Fig. 3. Illustration for Lem. 5

As shown in Fig. 3, it is intuitive to transform a normal Helly circular-arc model of G to an interval model of $\mho(G)$. Note that for any vertex $v \in T$, an induced (v^l, v^r)-path corresponds to a cycle whose arcs cover the entire circle. The main thrust of our algorithm will be a process that does the reversed direction, which is nevertheless far more involved.

Theorem 5. *If $\mho(G)$ is an interval graph, then we can in $O(n+m)$ time build a circular-arc model of G.*

Proof. We can in $O(n+m)$ time build an interval model $\mathcal{I}$ for $\mho(G)$. By construction, $(wh^l_{-1}h^l_0h^l_1h_2 \cdots h_{-2}h^r_{-1}h^r_0h^r_1)$ is an induced path of $\mho(G)$; without loss of generality, assume it goes "from left to right" in $\mathcal{I}$. We may assume $\texttt{rp}(w) = 0$ and $\max_{u \in \overline{T}} \texttt{rp}(u) = 1$, while no other interval in $\mathcal{I}$ has 0 or 1 as an endpoint. Let $a = \texttt{rp}(h^l_0)$. We use $\mathcal{I}$ to construct a set of arcs for $V(G)$ as follows. For each $u \in \overline{T}$, let $A_u := [\texttt{lp}(u), \texttt{rp}(u)]$, which is a subset of $(a, 1]$. For each $v \in T$, let

$$A_v := \begin{cases} [\texttt{lp}(v^r), \texttt{rp}(v^l)] & \text{if } v \in T_{\text{cc}}, \\ [\texttt{lp}(v^l), \texttt{rp}(v^l)] & \text{otherwise.} \end{cases}$$

It remains to verify that the arcs obtained as such represent G, i.e., a pair of vertices u, v of G is adjacent iff A_u and A_v intersect. This holds trivially when $u, v \notin T$; hence we may assume without loss of generality that $v \in T$. By construction, $a < \texttt{lp}(u) < \texttt{rp}(u) \leq 1$ for every $u \in \overline{T}$. Note that $v^l \sim w$ and $v^r \sim \overline{T}$ for every $v \in T_{\text{cc}}$, which implies that $\texttt{lp}(v^l) < 0$ iff $\texttt{lp}(v^r) < 1$ iff $v \in T_{\text{cc}}$.

Assume first that u is also in T, then $u \sim v$ in G if and only if $u^l \sim v^l$ in $\mho(G)$. They are adjacent when both $u, v \in T_{\text{cc}}$, and since $\texttt{lp}(v^l), \texttt{lp}(u^l) < 0$, both A_u and A_v contains the point 1 and thus intersect. If neither u nor v is in T_{cc}, then $\texttt{lp}(v^l), \texttt{lp}(u^l) > 0$, and $u \sim v$ if and only if $A_u = [\texttt{lp}(u^l), \texttt{rp}(u^l)]$ and $A_v = [\texttt{lp}(v^l), \texttt{rp}(v^l)]$ intersect. Otherwise, assume, without loss of generality, that $\texttt{lp}(v^l) < 0 < \texttt{lp}(u^l)$, then $u \sim v$ in G if and only if $0 < \texttt{lp}(u^l) < \texttt{rp}(v^l)$, which implies A_u and A_v intersect (as both contain $[\texttt{lp}(u^l), \texttt{rp}(v^l)]$).

Assume now that u is not in T, and then $u \sim v$ in G if and only if either $u \sim v^l$ or $u \sim v^r$ in $\mho(G)$. In the case $u \sim v^l$, we have $\texttt{lp}(v^l) \leq a < \texttt{lp}(u) \leq \texttt{rp}(v^l)$; since both A_u and A_v contain $[\texttt{lp}(u), \texttt{rp}(v^l)]$, which is nonempty, they intersect. In the case $u \sim v^r$, we have $\texttt{lp}(v^r) < \texttt{rp}(u) \leq 1$; since both A_u and A_v contain $[\texttt{lp}(v^r), \texttt{rp}(u)]$, they intersect. Otherwise, $u \not\sim v$ in G and $\texttt{lp}(v^l) < \texttt{rp}(v^l) < \texttt{lp}(u) < \texttt{rp}(u) < \texttt{lp}(v^r) < \texttt{rp}(v^r)$, then A_u and A_v are disjoint. □

We are now ready to present the recognition algorithm in Fig. 4, and prove Thm. 3. Recall that Lin et al. [6] have given a linear-time algorithm for verifying whether a circular-arc model is normal and Helly.

Algorithm **nhcag**(G)
Input: a graph G.
Output: a normal Helly circular-arc model, or a forbidden induced subgraph.

1 test the chordality of G and find a hole H if not;
  if G is chordal **then** verify whether G is an interval graph or not;
2 construct the auxiliary grpah $\mho(G)$;
3 **if** $\mho(G)$ is not an interval graph **then**
    call Thm. 4 to find a forbidden induced subgraph;
4 **call** 5 to build a circular-arc $\mathcal{A}$ model of G;
5 verify whether $\mathcal{A}$ is normal and Helly.

Fig. 4. The recognition algorithm for normal Helly circular-arc graphs

Proof (Thm. 3). Step 1 is clear. Steps 2-4 follow from Lem. 4, Thm. 4, and Lem. 5, respectively. If model $\mathcal{A}$ built in step 4 is not normal and Helly, then we can in linear time find a set of two or three arcs whose union covers the circle. Their corresponding vertices satisfy Lem. 3, and this concludes the proof. □

It is worth noting that if we are after a recognition algorithm (with positive certificate only), then we can simply return "NO" if the hypothesis of step 3 is true (justified by Lem. 5) and the algorithm is already complete.

3 Proof of Theorem 4

Recall that Thm. 4 is only called in step 3 of algorithm nhcag; the graph is then not choral and we have a hole H. In principle, we can pick any vertex as h_0. But for the convenience of presentation, we require it satisfies some additional conditions. If some vertex v is adjacent to four or more vertices in H, i.e., $\texttt{last}(v) - \texttt{first}(v) > 2$, then $v \notin H$. We can thus use $(h_{\texttt{first}(v)} v h_{\texttt{last}(v)})$ as a short cut for the sub-path induced by $N_H[v]$, thereby yielding a strictly shorter hole. This condition, that h_0 cannot be bypassed as such, is formally stated as:

Lemma 6. *We can in $O(n+m)$ time find either a minimal forbidden induced subgraph, or a hole H such that $\{h_{-1}, h_0, h_1\} \subseteq N_H[v]$ for some v if and only if $N_H[v] = \{h_{-1}, h_0, h_1\}$.*

This linear-time procedure can be called before step 2 of algorithm nhcag, and it does not impact the asymptotic time complexity of the algorithm, which remains linear. Henceforth we may assume that H satisfies the condition of Lem. 6. During the construction of $\mho(G)$, we have checked $N_H[v]$ for every vertex v, and Lem. 1 was called if it applies. Thus, for the proof of Thm. 4 in this section, we may assume that $N_H[v]$ always induces a proper sub-path of H.

Each vertex x of $\mho(G)$ different from w is uniquely defined by a vertex of G, which is denoted by $\phi(x)$. We say that x is *derived from* $\phi(x)$. For example, $\phi(v^l) = \phi(v^r) = v$ for $v \in T$. By abuse of notation, we will use the same letter for a vertex $u \in \overline{T}$ of G and the unique vertex of $\mho(G)$ derived from u, i.e., $\phi(u) = u$ for $u \in \overline{T}$; its meaning is always clear from the context. We can mark $\phi(x)$ for each vertex of $\mho(G)$ during its construction. For a set U of vertices not containing w, we define $\phi(U) := \{\phi(v) : v \in U\}$; possibly $|\phi(U)| \neq |U|$.

By construction, if a pair of vertices x and y (different from w) is adjacent in $\mho(G)$, then $\phi(x)$ and $\phi(y)$ must be adjacent in G as well. The converse is unnecessarily true, e.g., $u \not\sim v^r$ for any vertex $v \in T_{\mathrm{c}}$ and edge $uv \in E_{\mathrm{c}}$, and $u^l \not\sim v^r$ and $u^r \not\sim v^l$ for any pair of adjacent vertices $u, v \in T$. We say that a pair of vertices x, y of $\mho(G)$ is a *bad pair* if $\phi(x) \sim \phi(y)$ in G but $x \not\sim y$ in $\mho(G)$. By definition, w does not participate in any bad pair, and at least one vertex of a bad pair is in $L \cup R$. Note that any induced path of length d between a bad pair x, y with $x = v^l$ or v^r can be extended to a (v^l, v^r)-path with length $d + 1$.

Figure 3 shows that if G is a normal Helly circular-arc graph, then for any $v \in T$, the distance between v^l and v^r is at least 4. We now see what happens when this necessary condition is not satisfied by $\mho(G)$. By definition of $\mho(G)$, there is no edge between L and R; for any $v \in T$, there is no vertex adjacent to both v^l and v^r. In other words, for every $v \in T$, the distance between v^l and v^r is at least 3. The following observation can be derived from Lems. 1 and 2.

Lemma 7. *Given a (v^l, v^r)-path P of length 3 for some $v \in T$, we can in $O(n + m)$ time find a minimal forbidden induced subgraph of G.*

Proof. Let $P = (v^l x y v^r)$. Note that P must be a shortest (v^l, v^r)-path, and $w \notin P$. The inner vertices x and y cannot be both in $L \cup R$; without loss of generality, let $x \in \overline{T}$. Assume first that $y \in \overline{T}$ as well, i.e., $vx \in E_{\mathrm{c}}$ and $vy \in E_{\mathrm{cc}}$. By definition, $v \in T_{\mathrm{c}} \cap T_{\mathrm{cc}}$, and then v is adjacent to both h_{-1} and h_1. If follows from Lem. 6 that $N_H[v] = \{h_{-1}, h_0, h_1\}$, and then $x \sim h_1$ and $y \sim h_{-1}$. If $x \sim h_{-1}$, i.e., $\texttt{last}(x) = |H| - 1$, then we call Lem. 2 with v and x. If $\texttt{last}(x) < \texttt{first}(y)$, then we call Lem. 1 with x and y. In the remaining case, $\texttt{first}(y) \leq \texttt{last}(x) < |H| - 1$, and $(v x h_{\texttt{last}(x)} \cdots h_{-1} v)$ is a hole of G; this hole is completely adjacent to y, and thus we find a wheel.

Now assume that, without loss of generality, $y = u^r \in R$. If $\texttt{last}(v) \geq \texttt{first}(y)$, then we call Lem. 2 with v and y. Otherwise, $(v h_{\texttt{last}(v)} \cdots h_{\texttt{first}(y)} u v)$ is a hole of G; this hole is completely adjacent to x, and thus we find a wheel. □

If G is a normal Helly circular-arc graph, then in a circular-arc model of G, all arcs for T_{cc} and T_{c} contain $\texttt{ccp}(h_0)$ and $\texttt{cp}(h_0)$ respectively. Thus, both T_{cc} and T_{c} induce cliques. This observation is complemented by

Lemma 8. *Given a pair of nonadjacent vertices $u, x \in T_{cc}$ (or T_c), we can in $O(n+m)$ time find a minimal forbidden induced subgraph of G.*

Proof. By definition, we can find $uv, xy \in E_{cc}$. We have three (possibly intersecting) chordless paths $h_0h_1h_2$, h_0uv, and h_0xy. If both u and x are adjacent to h_1, then we return $(uh_{-1}xh_1u)+h_0$ as a wheel. Hence we may assume $x \not\sim h_1$.

If $u \sim h_1$, then by Lem. 6, $N_H[u] = \{h_{-1}, h_0, h_1\}$. We consider the subgraph induced by the set of distinct vertices $\{h_0, h_1, h_2, u, v, x\}$. If v is adjacent to h_0 or h_1, then we can call Lem. 3 with u, v. By assumption, h_0, h_1, and u make a triangle; x is adjacent to neither u nor h_1; and h_2 is adjacent to neither h_0 nor u. Thus, only uncertain adjacencies in this subgraph are between v, x, and h_2. The subgraph is hence isomorphic to (1) FIS-1 if there are two edges among v, x, and h_2; (2) $\overline{C_6}$ if v, x, and h_2 are pairwise adjacent; or (3) net if v, x, and h_2 are pairwise nonadjacent. In the remaining cases there is precisely one edge among v, x, and h_2. We can return a C^*, e.g., $(vxh_0uv)+h_2$ when the edge is vx.

Assume now that u, x, and h_1 are pairwise nonadjacent. We consider the subgraph induced by $\{h_0, h_1, h_2, u, v, x, y\}$, where the only uncertain relations are between v, y, and h_2. The subgraph is thus isomorphic to (1) $K_{2,3}$ if all of them are identical; or (2) twin-C_5 if two of them are identical, and adjacent to the other. If two of them are identical, and nonadjacent to the other, then the subgraph contains a C^*, e.g., $(vuh_0xv)+h_2$ when $v = y$. In the remaining cases, all of v, y, and h_2 are distinct, and then the subgraph (1) is isomorphic to long claw if they are pairwise nonadjacent; (2) contains net $\{h_1, h_2, u, v, x, y\}$ if they are pairwise adjacent; or (3) is isomorphic to FIS-2 if there are two edges among them. If there is one edge among them, then the subgraph contains a C^*, e.g., $(vuh_0xyv)+h_2$ when the edge is vy.

A symmetrical argument applies to T_c. the runtime is clearly $O(n+m)$. □

It can be checked in linear time whether T_{cc} and T_c induce cliques. When it is not, a pair of nonadjacent vertices can be found in the same time. By Lem. 8, we may assume hereafter that T_{cc} and T_c induce cliques. Recall that $N(w) \subseteq T_{cc}$; as a result, w is simplicial and participates in no holes.

Proposition 2. *Given a (h_0^l, h_0^r)-path nonadjacent to h_i for some $1 < i < |H| - 1$, we can in $O(n+m)$ time find a minimal forbidden induced subgraph.*

We are now ready to prove Thm. 4, which is separated into three statements, the first of which considers the case when $\mho(G)$ is not chordal.

Lemma 9. *Given a hole C of $\mho(G)$, we can in $O(n+m)$ time find a minimal forbidden induced subgraph of G.*

Proof. Let us first take care of some trivial cases. If C is contained in L or R or $\overline{T}$, then by construction, $\phi(C)$ is a hole of G. This hole is either nonadjacent or completely adjacent to h_0 in G, whereupon we can return $\phi(C) + h_0$ as a C^* or wheel respectively. Since L and R are nonadjacent, it must be one of the cases above if C is disjoint from $\overline{T}$. Henceforth we may assume that C intersects $\overline{T}$ and, without loss of generality, L; it might intersect R as well, but this fact is irrelevant in the following argument. Then we can find an edge x_1x_2 of C such that $x_1 \in L$ and $x_2 \in \overline{T}$, i.e., $x_1x_2 \in E_c$.

Let $a := \texttt{last}(\phi(x_1))$. Assume first that $x_2 = h_a$; then we must have $a > 1$. Let x_3 and x_4 be the next two vertices of C. Note that $x_3 \notin L$, i.e., $x_3 \not\sim h_0^l$; otherwise $x_1 \sim x_3$, which is impossible. If $x_3 \sim h_{a-2}$ (or h_{a-2}^l when $a = 3$), then $\phi(\{x_1, x_2, x_3\}) \cup \{h_{a-2}\}$ induces a hole of G, and we can return it and h_{a-1} as a wheel. Note that $x_4 \not\sim h_a$ as they are non-consecutive vertices of the hole C. We now argue that $\texttt{last}(\phi(x_4)) < a$. Suppose for contradiction, $\texttt{first}(\phi(x_4)) > a$. We can extend the (x_3, x_1)-path P in C that avoids x_2 to a (h_0^l, h_0^r)-path avoiding the neighborhood of h_a, which allows us to call Prop. 2. We can call Lem. 2 with x_3 and x_4 if $\texttt{first}(\phi(x_3)) = a$. In the remaining case, $\texttt{first}(\phi(x_3)) = a - 1$. Let x be the first vertex in P that is adjacent to h_{a-2} (or h_{a-2}^l if $a \leq 3$); its existence is clear as x_1 satisfies this condition. Then $\phi(\{x_3, \ldots, x, h_{a-2}, x_1, x_2\})$ induces a hole of G, and we can return it and h_{a-1} as a wheel.

Assume now that h_a is not in C. Denote by P the (x_2, x_1)-path obtained from C by deleting the edge x_1x_2. Let x be the first neighbor of h_{a+1} in P, and let y be either the first neighbor of h_{a-1} in the (x, x_1)-path or the other neighbor of x_1 in C. It is easy to verify that $\phi(\{x_1, \cdots, x, \cdots y, x_2\})$ induces a hole of G, which is completely adjacent to h_a, i.e., we have a wheel. □

In the rest $\mho(G)$ will be chordal, and thus we have a chordal non-interval subgraph F of $\mho(G)$. This subgraph is isomorphic to some graph in Fig. 1, on which we use the following notation. It is immediate from Fig. 1 that each of them contains precisely three simplicial vertices (squared vertices), which are called *terminals*, and others (round vertices) are *non-terminal vertices*. In a long claw or †, for each $i = 1, 2, 3$, terminal t_i has a unique neighbor, denoted by u_i.

Proposition 3. *Given a subgraph F of $\mho(G)$ in Fig. 1, we can in $O(n+m)$ time find either all bad pairs in F or a forbidden induced subgraph of G.*

Lemma 10. *Given a subgraph F of $\mho(G)$ in Fig. 1 that does not contain w, we can in $O(n+m)$ time find a minimal forbidden induced subgraph of G.*

Proof. We first call Prop. 3 to find all bad pairs in F. If F has no bad pair, then we return the subgraph of G induced by $\phi(F)$, which is isomorphic to F. Let x, y be a bad pair with the minimum distance in F; we may assume that it is 3 or 4, as otherwise we can call Lem. 7. Noting that the distance between a pair of non-terminal vertices is at most 2, we may assume that without loss of generality, x is a terminal of F. We break the argument based on the type of F.

Long claw. We may assume that $x = t_1$ and $y \in \{u_2, t_2\}$; other situations are symmetrical. Let P be the unique (x, y)-path in F. If $\phi(t_3)$ is nonadjacent to $\phi(P)$, then we return $\phi(P) + \phi(t_3)$ as a C^*; we are thus focused on the adjacency between $\phi(t_3)$ and $\phi(P)$. If $y = t_2$, then by the selection of x, y (they have the minimum distance among all bad pairs), $\phi(t_3)$ can be only adjacent to $\phi(t_1)$ and/or $\phi(t_2)$. We return either $\phi(F)$ as an FIS-2 , or $\phi(\{t_1, t_2, t_3, u_1, u_2, u_3\})$ as a net. In the remaining cases, $y = u_2$, and $\phi(t_3)$ can only be adjacent to $\phi(u_1)$, $\phi(u_2)$, and/or $\phi(t_1)$. We point out that possibly $\phi(t_2) = \phi(t_1)$, which is irrelevant as $\phi(t_2)$ will not be used below. If $\phi(t_3)$ is adjacent to both $\phi(u_1)$ and $\phi(u_2)$ in G, then we get a $K_{2,3}$. Note that this is the only case when $\phi(t_1) = \phi(t_3)$. If $\phi(t_3)$ is adjacent to both $\phi(t_1)$ and $\phi(u_2)$ in G, then we get an FIS-1. If $\phi(t_3)$ is adjacent to only $\phi(u_2)$ or only $\phi(t_1)$ in G, then we get a domino or twin-C_5 , respectively. The situation that $\phi(t_3)$ is adjacent to $\phi(u_1)$ but not $\phi(u_2)$ is similar as above.

†. Consider first that $x = t_1$ and $y = t_3$, and let $P = (t_1 u_1 u_3 t_3)$. If $\phi(t_2)$ is nonadjacent to the hole induced by $\phi(P)$, then we return $\phi(P)$ and $\phi(t_2)$ as a C^*. If $\phi(t_2)$ is adjacent to $\phi(t_3)$ or $\phi(u_1)$, then we get a domino. If $\phi(t_2)$ is adjacent to $\phi(t_1)$, then we get a twin-C_5. If $\phi(t_2)$ is adjacent to $\phi(t_1)$ and precisely one of $\{\phi(t_3), \phi(u_1)\}$, then we get an FIS-1. If $\phi(t_2)$ is adjacent to both $\phi(t_3)$ and $\phi(u_1)$, then we get a $K_{2,3}$; here the adjacency between $\phi(t_2)$ and $\phi(t_1)$ is immaterial. A symmetric argument applies when $\{t_2, t_3\}$ is a bad pair. In the remaining case, neither $\phi(t_1)$ nor $\phi(t_2)$ is adjacent to $\phi(t_3)$. Therefore, a bad pair must be in the path $F - N[t_3]$, which is nonadjacent to $\phi(t_3)$, then we get a C^*.

The whipping top and ‡ are straightforward and omitted. □

Lemma 11. *Given a subgraph F of $\mho(G)$ in Fig. 1 that contains w, we can in $O(n+m)$ time find a minimal forbidden induced subgraph of G.*

\\ Note that $0 \leq \texttt{last}(\phi(x_1)), \texttt{last}(\phi(x_2)) \leq 1$.
1 **if** $\texttt{last}(\phi(x_1)) = 1$ and $y_1 \sim h_2$ **then**
  call Lem. 2 with $(y_1 \phi(x_1) h_1 h_2 y_1)$ and $\{\phi(x_2), y_2\}$;
  if $\texttt{last}(\phi(x_1)) = 0$ and $y_1 \sim h_1$ **then**
  call Lem. 2 with $(y_1 \phi(x_1) h_0 h_1 y_1)$ and $\{\phi(x_2), y_2\}$;
  if $y_2 \sim h_{\texttt{last}(\phi(x_2))+1}$ **then** symmetric as above;
2 **if** $\texttt{last}(\phi(x_1)) = \texttt{last}(\phi(x_2))$ **then**
  return $\{y_1, \phi(x_1), y_2, \phi(x_2), h_{\texttt{last}(\phi(x_2))}, h_{\texttt{last}(\phi(x_2))+1}\}$ as a †;
\\ assume from now that $\texttt{last}(\phi(x_1)) = 1$ and $\texttt{last}(\phi(x_2)) = 0$.
3 **if** $\phi(x_2) \sim h_2$ **then return** $(\phi(x_2) h_0 h_1 h_2 \phi(x_2)) + y_1$ as a C^*;
4 **if** $y_2 \not\sim h_{-1}$ **then return** $\{y_1, h_{-1}, \phi(x_1), y_2, \phi(x_2), h_0, h_1\}$ as a ‡;
  if $y_2 \sim h_{-1}$ **then return** $\{y_1, h_{-1}, y_2, \phi(x_2), h_0, h_1\}$ as a †.

Fig. 5. Procedure for Lem. 11

Proof. Since w is simplicial, it has at most 2 neighbors in F. If w has a unique neighbor in F, then we can use a similar argument as Lem. 10. Now let x_1, x_2 be the two neighbors of w in F. If there exists some vertex $u \in \overline{T}$ adjacent to both $\phi(x_1)$ and $\phi(x_2)$ in G, which can be found in linear time, then we can use it replace w. Hence we assume there exists no such vertex. By assumption, we can find two distinct vertices $y_1, y_2 \in \overline{T}$ such that $\phi(x_1)y_1, \phi(x_2)y_2 \in E_{\text{cc}}$; note that $\phi(x_1) \not\sim y_2$ and $\phi(x_2) \not\sim y_1$ in G. As a result, y_1 and y_2 are nonadjacent; otherwise, $\{y_1, y_2\}$ and the counterparts of $\{x_1, x_2\}$ in R induce a hole of $\mho(G)$, which is impossible. We then apply the procedure described in Fig. 5.

We now verify the correctness of the procedure. Since each step—either directly or by calling a previously verified lemma—returns a minimal forbidden induced subgraph of G, all conditions of previous steps are assumed to not hold in a later step. By Lem. 6, $\texttt{last}(\phi(x_1))$ and $\texttt{last}(\phi(x_2))$ are either 0 or 1. Step 1 considers the case where $y_1 \sim h_{\texttt{last}(\phi(x_1))+1}$. By Lem. 3, $y_1 \not\sim h_{\texttt{last}(\phi(x_1))}$. Thus, $(y_1\phi(x_1)h_1h_2y_1)$ or $(y_1\phi(x_1)h_0h_1y_1)$ is a hole of G, depending on $\texttt{last}(\phi(x_1))$ is 0 or 1. In the case $(y_1\phi(x_1)h_1h_2y_1)$, only $\phi(x_1)$ and h_1 can be adjacent to $\phi(x_2)$; they are nonadjacent to y_2. Likewise, in the case $(y_1\phi(x_1)h_0h_1y_1)$, vertices $\phi(x_1)$ and h_0 are adjacent to $\phi(x_2)$ but not y_2, while h_1 can be adjacent to only one of $\phi(x_2)$ and y_2. Thus, we can call Lem. 2. A symmetric argument applies when $y_2 \sim h_{\texttt{last}(\phi(x_2))+1}$. Now that the conditions of step 1 do not hold true, step 2 is clear from assumption. Henceforth we may assume without loss of generality that $\texttt{last}(\phi(x_1)) = 1$ and $\texttt{last}(\phi(x_2)) = 0$. Consequently, $\texttt{last}(y_1) = |H| - 1$ (Lem. 2). Because we assume that the condition of step 1 does not hold, $y_1 \not\sim h_2$; this justifies step 3. Step 4 is clear as y_1 is always adjacent to h_{-1}. □

References

1. Cao, Y.: Linear recognition of almost (unit) interval graphs (2014) (manuscript)
2. Grippo, L., Safe, M.: On circular-arc graphs having a model with no three arcs covering the circle. arXiv:1402.2641 (2014)
3. Heggernes, P., Kratsch, D.: Linear-time certifying recognition algorithms and forbidden induced subgraphs. Nord. J. Comput. 14, 87–108 (2007)
4. Kratsch, D., McConnell, R., Mehlhorn, K., Spinrad, J.: Certifying algorithms for recognizing interval graphs and permutation graphs. SIAM J. Comput. 36, 326–353 (2006)
5. Lekkerkerker, C., Boland, J.: Representation of a finite graph by a set of intervals on the real line. Fund. Math. 51, 45–64 (1962)
6. Lin, M., Soulignac, F., Szwarcfiter, J.: Normal Helly circular-arc graphs and its subclasses. Discrete Appl. Math. 161, 1037–1059 (2013)
7. Lindzey, N., McConnell, R.M.: On finding Tucker submatrices and Lekkerkerker-Boland subgraphs. In: Brandstädt, A., Jansen, K., Reischuk, R. (eds.) WG 2013. LNCS, vol. 8165, pp. 345–357. Springer, Heidelberg (2013)
8. McConnell, R.: Linear-time recognition of circular-arc graphs. Algorithmica 37, 93–147 (2003)
9. McConnell, R., Mehlhorn, K., Näher, S., Schweitzer, P.: Certifying algorithms. Computer Science Review 5, 119–161 (2011)
10. McKee, T.: Restricted circular-arc graphs and clique cycles. Discrete Math. 263, 221–231 (2003)
11. Tarjan, R., Yannakakis, M.: Simple linear-time algorithms to test chordality of graphs, test acyclicity of hypergraphs, and selectively reduce acyclic hypergraphs. SIAM J. Comput. 13, 566–579. Addendum in the same journal 14, 254–255 (1985)