THE LBFS STRUCTURE AND RECOGNITION OF INTERVAL GRAPHS

DEREK G. CORNEIL*, STEPHAN OLARIU†, AND LORNA STEWART‡

Abstract. A graph is an interval graph if it is the intersection graph of intervals on a line. Interval graphs are known to be the intersection of chordal graphs and asteroidal triple-free graphs, two families where the well-known Lexicographic Breadth First Search (LBFS) plays an important algorithmic and structural role. In this paper we show that interval graphs have a very rich LBFS structure and that, by exploiting this structure, one can design a linear time, easily implementable, interval graph recognition algorithm.

Key words. lexicographic breadth first search, interval graphs, chordal graphs, AT-free graphs, structure, recognition, algorithm

AMS subject classifications. 05C62, 05C75, 05C85, 68R10

1. Introduction. *Interval graphs* (namely, the intersection graphs of intervals on a line) arise naturally in the process of modeling real-life situations, especially those involving time dependencies or other restrictions that are linear in nature. Fifty years ago, Benzer [1] used interval graphs to model genetic structure and, since then, dozens of papers have described applications of interval graphs to such diverse areas as archaeology, biology, psychology, sociology, management, genetics, engineering, scheduling, transportation and others. For example, Choi and Farach-Colton [4] make use of the algorithm presented in this paper to solve the sequence assembly problem significantly faster than previous methods. For a wealth of information about interval graphs the interested reader is referred to [15], where many of the above applications are summarized.

Lexicographic Breadth First Search (LBFS) is a search paradigm developed in 1976 by Rose, Tarjan and Lueker [30] for the efficient recognition of *chordal graphs*, namely graphs that contain no induced cycle of size greater than three. Their fundamental paper showed that a graph is chordal if and only if any LBFS will result in an ordering of the vertex set that satisfies the condition that no vertex in the ordering is adjacent to two nonadjacent vertices that occur before it in the ordering.

Little work was done on LBFS for almost two decades after the Rose, Tarjan and Lueker paper. Jamison and Olariu [20] were probably the first to use LBFS in a context other than chordal graph recognition. Recently LBFS has received a great deal of attention (see [6] for a survey), including the work by the authors on *asteroidal triple-free (AT-free) graphs*. An *asteroidal triple* is an independent triple of vertices such that between each pair of vertices in the triple there is a path that avoids the neighbourhood of the third vertex. AT-free graphs were introduced by Lekkerkerker and Boland [25] who produced the following characterization of interval graphs.

THEOREM 1.1. *[25] A graph is an interval graph if and only if it is chordal and asteroidal triple-free.*

In light of this theorem and the LBFS results, both algorithmic and structural, for

*Department of Computer Science, University of Toronto, Toronto, Ontario, Canada (dgc@cs.utoronto.ca). Financial support received from NSERC and CITO

†Department of Computer Science, Old Dominion University, Norfolk, Virginia, U.S.A (olariu@cs.odu.edu). Supported, in part, by NSF grant CCR-9522093

‡Department of Computing Science, University of Alberta, Edmonton, Alberta, Canada (stewart@cs.ualberta.ca). Financial support received from NSERC

chordal and AT-free graphs, one might expect that interval graphs also are amenable to LBFS algorithms and have a rich LBFS structure. In this paper we show that this is in fact the case and that the LBFS structure leads to an easily implementable linear time interval graph recognition algorithm.[1] This algorithm involves a preprocessing arbitrary LBFS sweep followed by five LBFS sweeps with specific tie breaking rules.

The following characterization presented in [14] turned out to be the workhorse of the vast majority of recognition algorithms for interval graphs.

THEOREM 1.2. *[14] A graph is an interval graph if and only if its maximal cliques can be linearly ordered in such a way that for every vertex in the graph the maximal cliques to which it belongs occur consecutively in the linear order.*

A crude implementation of the characterization in Theorem 1.2 yields a recognition algorithm that runs in time proportional to the cube of the number of vertices in the graph. Later, Booth and Lueker [2] used PQ-trees to compute a required linear order on the set of maximal cliques, if such an order exists. In the process, they reduced the complexity to linear in the size of the graph, which is best possible. In spite of this, the algorithm of [2] was less than perfect. For one thing, a PQ-tree is a complicated data structure and, not surprisingly, the resulting algorithm was rather involved. Roughly ten years later, Korte and Möhring [21] streamlined the recognition algorithm of Booth and Lueker, using a variant of PQ-trees that they called MPQ-trees, while still using the characterization by maximal cliques as the focal point of their algorithm. Hsu and McConnell [19] further streamlined PQ-tree manipulations with the introduction of PC-trees. Recently, Kratsch, McConnell, Mehlhorn, and Spinrad [22], modified the algorithm of [21] to produce a certifying interval graph recognition algorithm. Hsu [17] demonstrated that PQ-trees are not essential for recognizing interval graphs by using modular decomposition techniques and not relying on Theorem 1.2. Although still linear time, this algorithm was also difficult to implement since there was no known easily implementable linear time algorithm to perform modular decomposition. More recently, Hsu and Ma [18] gave a simple modular decomposition algorithm for chordal graphs, based on a variation of LBFS, as well as an algorithm for linearly ordering the maximal cliques of a prime (with respect to modular decomposition) interval graph. The resulting linear time interval graph recognition algorithm is easier to implement, does not use PQ-trees, and is again based on the characterization of Theorem 1.2. Habib, McConnell, Paul and Viennot [16] also developed a linear time algorithm that is relatively easy to implement. Their algorithm, which is completely different from ours, uses LBFS to determine that the given graph is chordal and to produce a clique tree. It then manipulates the clique tree into a clique path if the graph is an interval graph, thus satisfying Theorem 1.2.

Instead of using Theorem 1.2, our algorithm is based on the following characterization of interval graphs independently observed by many researchers including [27, 28, 29].

THEOREM 1.3. *[27, 28, 29] A graph is an interval graph if and only if there exists a linear order $\prec$ on the set of its vertices such that for every choice of vertices u, v, w, with $u \prec v$ and $v \prec w$,*

$$uw \in E \implies uv \in E. \tag{1.1}$$

[1] A previous version of this algorithm without the first two sweeps was incorrectly claimed to work [10]. See §5 for details.

We call any ordering of vertices satisfying condition (1.1) of Theorem 1.3 an *I-ordering*. If an edge uw does not satisfy this condition (i.e. $uv \notin E$), then the edge uw is called an *umbrella*. The term *umbrella-free ordering* will be used synonymously with I-ordering.

Our paper is organized as follows. Section 2 reviews graph-theoretic concepts, establishes notation and terminology and introduces both the standard LBFS as well as a variant, LBFS+. Since the paper presents two major contributions, namely the LBFS structure of interval graphs and the new interval graph recognition algorithm, the paper is divided into two parts. Sections 3 and 4 constitute the first part and provide the LBFS structural results on graph families containing interval graphs and on interval graphs respectively. Part 2 starts with §5 where the interval graph recognition algorithm together with an example are presented. The proof of correctness of the algorithm and its easy linear time implementation are presented in §6 and §7 respectively. Section 8 contains concluding remarks and open problems.

We have tried to adopt a consistent approach to the naming of our results. For the most part, theorems are results that may be of independent interest, lemmas are important steps in the development of the paper, and claims are intermediate steps leading up to the various lemmas and theorems. Since the paper introduces many new terms and concepts, an index is provided at the end of the paper. In the proof of correctness of the algorithm, various variables are assigned specific meanings; in the index such variables, as well as definitions associated with the proof, are annotated with (PC).

2. Background. We now provide the background information needed for the paper. First we establish various notation and definitions; we then describe both generic LBFS and a variant of it, LBFS+.

2.1. Notation and definitions. All the graphs in this work are finite with no loops or multiple edges. Let $G = (V, E)$. We use $d(u, v)$ to denote the distance between vertices u and v in G. The *eccentricity* of vertex v, denoted $ecc(v)$, is $max_{u \in V} d(u, v)$, and the *diameter* of graph G, denoted $diam(G)$, is $max_{u,v \in V} d(u, v)$. The *neighbourhood* of vertex $v \in V$, denoted $N(v)$, is the set of vertices adjacent to v. The *closed neighbourhood* of v, denoted $N[v]$ is $N(v) \cup \{v\}$. For $W \subseteq V$, we let $N(W)$, the *neighbourhood of* W, denote $\{v \in V \setminus W$ such that v is adjacent to at least one vertex in $W\}$. Given a set W of vertices of G, we say that vertex v is *universal* with respect to W whenever v is adjacent to all the vertices of $W \setminus v$. We often abuse the language and say that v is universal with respect to the subgraph induced by W. (Note we allow $v \in W$ or $v \in V \setminus W$.) We let K_G denote the vertices in G that are universal to G; clearly, this set is a clique, justifying the use of "K_G". By convention, $K_G = \emptyset$ if G has no universal vertices. A vertex v is *homogeneous* to a set X if v is adjacent to all vertices of X or to none of them. Set of vertices $M \subseteq V$ is a *module* of G if, for all $v \in V \setminus M$, v is homogeneous to M. A nontrivial module is one of size greater than one and less than $|V|$. A vertex v is said to be *simplicial* if its neighbours are pairwise adjacent. A subset of vertices S is a *separator* if $G \setminus S$ is disconnected; S is a *minimal separator* if no proper subset of S is also a separator. We refer to a path[2] joining vertices x and y as an x, y-path. A vertex u *intercepts a path* π if u is on π or adjacent to at least one vertex on π; otherwise, u is said to *miss* π. We also say that π misses u. For vertices u, v in G, we let $D(u, v)$ denote the set of vertices that intercept all u, v-paths. Vertex pair (u, v) is said to be a *dominating pair*

[2] Unless stated otherwise, we assume that the paths are induced.

whenever $D(u, v) = V$. We say that vertices u and v are *unrelated with respect to* y if $u \notin D(v, y)$ and $v \notin D(u, y)$. A vertex y of G is said to be *admissible* if there are no vertices in G unrelated with respect to y.

If G is an AT-free graph then vertex y of G is termed *pokable* if the graph obtained from G by adding a pendant vertex adjacent to y is AT-free. It was shown in [9] that in an AT-free graph every admissible vertex is also pokable. A *pokable dominating pair* is a dominating pair such that both vertices are pokable. A vertex x is a *pokable dominating pair vertex* if x is pokable and there exists y such that (x, y) is a dominating pair.

To simplify the notation, throughout the remainder of the paper we deliberately blur the distinction between a set S of vertices and the subgraph $G[S]$ (or G_S) it induces.

2.2. LBFS and LBFS+. Let $G = (V, E)$ be a graph and let u be a vertex of G. We now reproduce the details of a variant of LBFS [30] that allows arbitrary tie-breaking; later in the paper we impose specific tie-breaking mechanisms. We warn the reader that our LBFS ordering of the vertices of the graph may seem "backwards" compared to the ordering produced by other LBFS descriptions.

Procedure LBFS(G, u);
{Input: a graph $G = (V, E)$ and a distinguished vertex u of G;
Output: an ordering σ_u of the vertices of G}
begin
  label$(u) \leftarrow |V|$;
  for each vertex v in $V - \{u\}$ **do**
    label$(v) \leftarrow \Lambda$;
  for $i \leftarrow |V|$ **downto** 1 **do begin**
    pick any unnumbered vertex v with lexicographically the largest label; $(\star)$
    $\sigma_u(|V| + 1 - i) \leftarrow v$; {place v in σ_u; v is now considered to be numbered}
    for each unnumbered vertex w in $N(v)$ **do**
      append i to label(w)
    end
  end; {LBFS}

In an LBFS σ with two arbitrary vertices u and v, if vertex u is visited before v, i.e. $u <_\sigma v$ we say that u *occurs* before v in σ or that u is *visited before* v or that u is *to the left of* v. Note that if G is disconnected, then all vertices of the connected component containing u must be visited by the LBFS procedure before the next component is visited (see Claim 3.6). As mentioned above, this generic LBFS algorithm allows arbitrary choice of a vertex in step $(\star)$. We call a set of tied vertices encountered in step $(\star)$ a *slice*. We will also use the term slice to refer to the subgraph induced on the set of vertices that constitute a slice. Note that V itself is a slice, namely, the *universal slice*, since all vertices of G are considered tied before the first vertex is chosen. Unless explicitly stated otherwise, we will only be interested in nontrivial slices, that is, those containing at least two vertices. Note that all vertices of a slice with respect to LBFS σ appear consecutively in σ and all vertices to the left of slice S in σ are homogeneous to S. We will use square brackets to indicate the *slice structure* of an LBFS, that is, the collection of slices of the LBFS. Given two vertices u and v of an LBFS σ such that $u <_\sigma v$, $\Gamma^\sigma_{u,v}$ denotes the vertex-minimal slice with respect to σ that contains both u and v. Given slice S of LBFS σ, we say that slice

$S' \subset S$ is an *outermost* slice of S with respect to σ if there is no slice T of S with respect to σ such that $S' \subset T \subset S$. As an example of these concepts consider the graph in Figure 1 where the boxes indicate the nontrivial slices with respect to the LBFS σ (note that the vertices are numbered as visited by σ). The slice structure of LBFS σ is [1 [2 3] 4 [5 6 [7 8 9] 10] 11]. The nontrivial outermost slices of V with respect to σ are $\{2, 3\}$ and $\{5, 6, 7, 8, 9, 10\}$. $\Gamma^{\sigma}_{9,10}$ consists of $\{5, 6, 7, 8, 9, 10\}$. Occasionally we will refer to Γ^{σ}_{S} where S is a set of vertices; this is equivalent to $\Gamma^{\sigma}_{s,t}$ where $s \in S$ is the first vertex in σ of S and $t \in S$ is the last vertex in σ of S. The restriction of σ to S will be denoted σ_S.

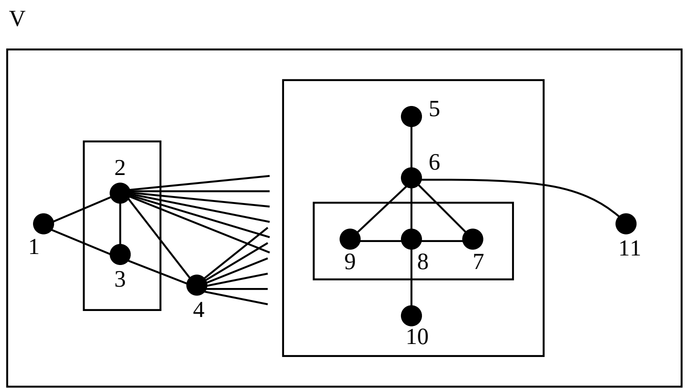

FIG. 1. *A graph with its LBFS slices*

We say that a vertex x is *good* or is an *end-vertex* if there is some LBFS that ends at x. Similarly, we say that a vertex is *bad* if no LBFS can end at x. Incidentally, for most families of graphs it is an interesting open question to characterize good vertices for that family; see [8] for a recent survey of results on end-vertices. In §4 we characterize good vertices for interval graphs. We say that an LBFS is *good* if every slice starts with a vertex that is good for the slice. We now describe a variant of LBFS that, for some graph classes, is guaranteed to produce a good LBFS. In particular, it uses a given ordering (in our case, an ordering produced by a previous LBFS) to break ties in step ($\star$). This variant has been independently investigated by Ma [26] and Simon [31].

Procedure LBFS+ (G, τ):
For this LBFS procedure, one previous LBFS ordering τ is needed. In the LBFS procedure at step ($\star$), let S be the set of vertices with the lexicographically largest label. Now v is chosen to be the vertex in S that appears *last* in τ.

As an example, LBFS+ when given the graph in Figure 1 and that LBFS, would produce the following order and slice structure: [11 6 [9 [8 [4 2]] 7 5] 10 3 1].

It is interesting to note that LBFS+ was used by Simon [31] in his interval graph recognition algorithm. In particular, he performed an arbitrary LBFS followed by three applications of LBFS+. He claimed that for an arbitrary interval graph G and for an arbitrary initial LBFS of G, the ordering resulting from the fourth sweep would exhibit a linear ordering of the cliques of G. Ma [26] however, showed that Simon's algorithm is flawed and that for any constant c, there is an interval graph, and an initial LBFS ordering such that after c applications of LBFS+, the linear ordering of cliques is still not apparent!

PART 1
THE LBFS STRUCTURE OF INTERVAL GRAPHS

3. LBFS structure of families containing interval graphs. We now present results on the LBFS structure of graph families that contain interval graphs. In particular, we present LBFS structural results for arbitrary graphs, chordal graphs and AT-free graphs. In §4, we study the LBFS structure of interval graphs themselves.

3.1. General graphs. We start with LBFS properties of arbitrary graphs. The first result characterizes an LBFS ordering with respect to every triple of vertices in the ordering.

THEOREM 3.1. *[13][15] An ordering $<$ of the vertices of an arbitrary graph $G = (V, E)$ is an LBFS ordering if and only if for all vertices a, b, c of G such that $ac \in E$ and $bc \notin E$, $c < b < a$ implies the existence of a vertex d in G, adjacent to b but not to a and such that $d < c$.*

As mentioned in §2, the concept of slice with respect to LBFS σ (namely a set of vertices that at some point in the execution of σ are unnumbered and have the highest label among the unnumbered vertices) is fundamental in the study of LBFS orderings.

To motivate the next concept, consider the graph in Figure 2 where the vertices are numbered as visited by LBFS σ, and some of the modules are indicated by boxes. The slice structure is: [1 2 3 [4 [5 [6 [7 8] 9] [10 11 12 13 14]] 15 [16 17] [18 19]]]. Now consider the P_4 induced on $\{4, 6, 9, 15\}$. Subsequently we will want to analyze the ordering of σ when restricted to such a set. If we examine the smallest slice containing these vertices, then we have $\Gamma^{\sigma}_{4,15} = \{4, \cdots, 19\}$ and thus we have ignored the fact that the vertices in $\Gamma^{\sigma}_{4,15} \setminus \{4, 6, 9, 15\}$ have no impact on how the vertices are chosen inside the P_4. This leads to the notion of an M-slice (module of a slice).

In the context of an LBFS on graph G, an *M-slice of slice* S is S itself or any nontrivial module of S. In the example of Figure 2, the M-slices of V are V, $\{4, \cdots, 19\}$ and the nontrivial modules of $\{4, \cdots, 19\}$. Subslice $\{6, 7, 8, 9\}$ is not an M-slice of V, but is an M-slice of the slice $\{5, \cdots, 14\}$. As can be seen in Figure 2, while an M-slice of S is a module of the graph induced on the vertices in slice S, it is not necessarily a module of G; note that $\{7, 8\}$ is a module of slice $\{6, 7, 8, 9\}$, but it is not a module of G since 16 is adjacent to 7 but not to 8. Nevertheless, as with slices, all vertices to the left of an M-slice are homogeneous to the M-slice.

We now present some results on the presence of modules in LBFS sweeps of arbitrary graph G. For X, a set of vertices in V and any LBFS σ, we let $sc_\sigma(X)$*, (the scope of X in σ)* denote the smallest contiguous set, with respect to the ordering of V imposed by σ, that contains X.

CLAIM 3.2. *Let M be a module of graph G. Then M is an M-slice with respect to every LBFS of G.*

Proof. Let σ be an LBFS of G and consider the point in σ just before the first vertex of M is chosen. Let S be the slice defined at this point; clearly $M \subseteq S$ and each vertex in $S \setminus M$ is either adjacent to all vertices in M or to no vertices in M. The result follows from the definition of M-slice. □

In fact, we can say something about which nodes can be in the scope of a module.

LEMMA 3.3. *Let σ be an arbitrary LBFS of graph $G = (V, E)$ that contains module M. Suppose $uw \in E$ where $u, w \in M$. If $v \notin M$ satisfies $u <_\sigma v <_\sigma w$, then*

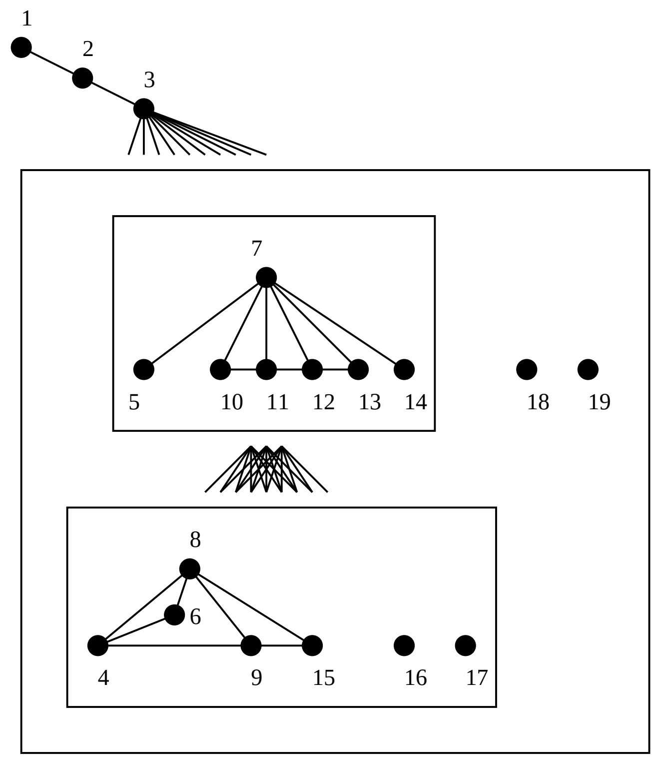

FIG. 2. *A graph with some of its modules indicated by boxes*

v is universal to M.

Proof. Suppose there is such a v that is not universal to M and without loss of generality, assume that v is the leftmost such vertex that is spanned by some edge of M. Applying Theorem 3.1 to the triple $\{u, v, w\}$ there is a vertex $x <_\sigma u$ such that $xv \in E, xw \notin E$. Without loss of generality, assume x is the leftmost such vertex. Since v is adjacent to x but not to vertices in M, $x \notin M$ and thus since $xw \notin E, xu \notin E$. Now applying Theorem 3.1 to the triple $\{x, u, v\}$, there is a vertex $y <_\sigma x$ such that $yu \in E, yv \notin E$. We conclude that $y \notin M$ since otherwise $yx \notin E$ and we would have contradicted the choice of v being the leftmost vertex not in M that is spanned by an edge of M. Since $yu \in E$, y is universal to M and thus $yw \in E$. Applying Theorem 3.1 to the triple $\{y, v, w\}$ there is a vertex $z <_\sigma y$ such that $zv \in E, zw \notin E$. But $z <_\sigma y <_\sigma x$ contradicting x being the leftmost vertex that is adjacent to v and not adjacent to w. ∎

For $u <_\sigma v$, we refer to any M-slice of $\Gamma^\sigma_{u,v}$ that contains u and v as an M-$\Gamma^\sigma_{u,v}$. In general, an M-Γ^σ_S is any M-slice of Γ^σ_S that contains S. We let $\widetilde{\Gamma}^\sigma_{u,v}$ denote the smallest M-$\Gamma^\sigma_{u,v}$. Note that the vertices of an M-slice of σ are not necessarily contiguous in σ. To illustrate these notions, consider the graph in Figure 2. $\Gamma^\sigma_{4,15} = \{4, \cdots, 19\}$, and $\widetilde{\Gamma}^\sigma_{4,15}$, its smallest M-slice containing $\{4, 15\}$ is $\{4, 6, 9, 15\}$. In addition, $\widetilde{\Gamma}^\sigma_{5,14}$ and $\widetilde{\Gamma}^\sigma_{6,9}$ are $\{5, 14\}$ and $\{6, 9\}$ respectively. Notice that $\{6, 9\}$ is an M-slice of slice $\{5, \cdots, 14\}$ but is not a module of G; also notice that 6 and 9 are not contiguous in σ.

When we refer to "any M-slice with respect to LBFS σ", without reference to a particular slice, we mean an M-slice of S where S is an arbitrary slice of σ.

The next lemma is fundamental in showing that M-slices inherit much of the structure of slices.

LEMMA 3.4. *Let S be an M-slice of a graph G with respect to an LBFS σ.*
(i) Then σ_S, the restriction of σ to S, is an LBFS of G_S.
(ii) Furthermore if σ is a good LBFS of G, then the first vertex of S in σ is a good vertex of S.

Proof. If S is a slice, results (i) and (ii) follow immediately. Otherwise, let S' be the smallest slice strictly containing S. First, assume claim (i) does not hold for S. By Theorem 3.1, since σ_S is not an LBFS, there exist vertices $c <_\sigma b <_\sigma a$ such that $ac \in E, bc \notin E$ but no vertex $d \in S$ such that $d <_\sigma c, db \in E$ and $da \notin E$. Since $\sigma_{S'}$ is an LBFS, such a vertex does exist in S'. But all vertices in $S' \setminus S$ adjacent to any vertex in S are universal to S, contradicting that fact that $da \notin E$.

To prove claim (ii), note that since every vertex of $S' \backslash S$ is adjacent to all or none of the vertices of S, all vertices of S, possibly together with some vertices of $S' \backslash S$, are tied when the first vertex of S is chosen in σ. Since σ is a good LBFS, the first vertex of S is thus a good vertex in the set of tied vertices and therefore a good vertex in S, by (i). Claim (ii) follows. ◻

The next five results present straightforward results concerning M-slices.

CLAIM 3.5. *Let S be an arbitrary M-slice of a graph G with respect to some LBFS σ and let t be the first vertex visited in S by σ. Every vertex of G that occurs before t in σ is either adjacent to all vertices of S or to none of them.*

Proof. If S is a slice, then this follows immediately from the fact that all the vertices in S had the same label when t was about to be numbered by σ. If S is a module of a slice, the proof follows from the above and the definition of M-slice. ◻

CLAIM 3.6. *Let S be an arbitrary disconnected M-slice of a graph G with respect to some LBFS σ and let C and D be different connected components of S. If some vertex of C occurs in σ before some vertex of D, then all vertices of C occur before all vertices of D.*

Proof. Let c be the first vertex numbered by σ in $C \cup D$ and assume that c belongs to C. We assume that $|C| \geq 2$, for otherwise there is nothing to prove. Let d be the first vertex of D visited by σ and let c' be any vertex of C that occurs after d. Since C is connected, there is a path P from c to c' that misses d and thus there exists edge $c_1c_2 \in P$ such that $c_1d, c_2d \notin E$ and $c_1 <_\sigma d <_\sigma c_2$. By Theorem 3.1, there exists vertex d' where $d' <_\sigma c_1, d'd \in E$ and $d'c_2 \notin E$. Now $d' \notin D$ by the choice of d, $d' \notin S \setminus D$ since it is adjacent to d and d' is not universal to S since $d'c_2 \notin E$. By Claim 3.5, we have a contradiction. ◻

CLAIM 3.7. *For graph G and LBFS σ:*
(i) let S be a slice of G with respect to σ and let x be a vertex that occurs after S in σ. Then some vertex y of G that occurs before S is adjacent to all the vertices in S and nonadjacent to x.
(ii) let S be an M-slice of G with respect to σ and let x be a vertex that occurs after S in σ where x is adjacent to some but not all vertices of S. Then some vertex y of G that occurs before S is adjacent to all the vertices in S and nonadjacent to x.

Proof. To prove (i), let t be the first vertex visited in S by σ. Since x does not belong to S, it must have had a smaller label than that of t at the moment when t was chosen by σ. In turn, this implies that there must exist a *numbered* vertex y in G adjacent to t but not to x. By Claim 3.5, y must be adjacent to all the vertices in S.

To prove (ii), suppose S is a proper module of slice T; by the definition of module, $x \notin T$ and thus the conclusion follows from (i). ◻

LEMMA 3.8. *Let S_1 and S_2 be arbitrary slices of a graph G with respect to some LBFS σ. If $S_1 \cap S_2 \neq \emptyset$, then $S_1 \subseteq S_2$ or $S_2 \subseteq S_1$.*

Proof. Suppose $x \in S_1 \backslash S_2$, $y \in S_1 \cap S_2$, and $z \in S_2 \backslash S_1$, where $x <_\sigma y <_\sigma z$. (Recall that all vertices in S_1 (respectively S_2) occur consecutively in σ.) By Claim 3.7(i), there exists a vertex $w <_\sigma x$ such that $wx, wy \in E$ but $wz \notin E$. But now, since w appears before all vertices in S_2, by Claim 3.5 we have contradicted the fact that S_2 is a slice. □

Recall that given two vertices u, v and an LBFS σ of G such that $u <_\sigma v$, $\Gamma^\sigma_{u,v}$ denotes the vertex-minimal slice with respect to σ that contains both u and v. From Lemma 3.8, if $u <_\sigma v <_\sigma w$, then either $\Gamma^\sigma_{u,v} \subseteq \Gamma^\sigma_{v,w}$ or vice versa. This implies that one of $\Gamma^\sigma_{u,v}$ and $\Gamma^\sigma_{v,w}$ matches $\Gamma^\sigma_{u,w}$.

CLAIM 3.9. *Let σ be an arbitrary LBFS of a graph G, let u and v be vertices of G satisfying $u <_\sigma v$, let T be an M-$\Gamma^\sigma_{u,v}$, and let t be the first vertex of the connected component of T that contains u. If t is distinct from u, then t cannot be adjacent to v.*

Proof. First we note that in every M-$\Gamma^\sigma_{u,v}$, the connected component containing u must start with the same vertex t. This follows from the fact that vertices of each of these sets are tied when the first vertex of the set is numbered, the definitions of slice and M-slice, and Claim 3.6. Suppose that t is adjacent to v. In particular, once t is visited by σ, v inherits t's label. Since all the vertices in $\Gamma^\sigma_{u,v}$ (and thus in T) were tied when t was numbered, this implies that all the vertices of $\Gamma^\sigma_{u,v}$ preceding v must also be adjacent to t. But now, there is a slice containing both u and v, namely $\Gamma^\sigma_{u,v} \cap N(t)$, strictly included in $\Gamma^\sigma_{u,v}$, contradicting the minimality of $\Gamma^\sigma_{u,v}$. □

The next theorem is fundamental in the study of LBFS on arbitrary graphs.

THEOREM 3.10. **(The Prior Path Theorem)** *Let σ be an arbitrary LBFS of a graph G, let u and v be vertices of G satisfying $u <_\sigma v$, let T be an M-$\Gamma^\sigma_{u,v}$, and let t be the first vertex of the connected component of T that contains u. There exists a t, u-path in T all of whose vertices, with the possible exception of u, are missed by v. Moreover, all vertices other than u on this path occur before u in σ. (Such a path is called a* prior path*).*

Proof. If t and u coincide, then there is nothing to prove. Therefore, we shall assume that t and u are distinct vertices. By Claim 3.9, t is not adjacent to v.

The minimality of $\Gamma^\sigma_{u,v}$ guarantees that u and v could not have been tied at the time when u was chosen by σ. In turn, this implies the existence of a vertex u_1, adjacent to u but not to v, that occurs prior to u in σ. By Claim 3.5, u_1 belongs to T and thus to C, the connected component of T that contains u. In fact, we choose u_1 to be the *earliest* vertex of C adjacent to u but not to v.

If u_1 and t coincide, the path consisting of u and u_1 is the desired path. Now, when u_1 was chosen, u_1 and v could not have been tied, otherwise u and v would have been tied, contradicting the minimality of $\Gamma^\sigma_{u,v}$. Let u_2 be the *earliest* vertex of G, adjacent to u_1 but not to v, that occurs before u_1 in σ. As before, Claim 3.5 guarantees that u_2 belongs to T and thus to C. Moreover, our choice of u_1 guarantees that u_2 and u are not adjacent.

If u_2 and t coincide, then u_2, u_1, u is the desired path. Otherwise, the process above continues inductively. The finiteness of G guarantees that for some subscript i, $(i \geq 1)$, u_i and t coincide. Now, $t = u_i, u_{i-1}, \ldots, u_1, u$ is the desired path. □

3.2. Chordal graphs. We now study LBFSs in chordal graphs. The first three results are either stated or are implicit in the seminal paper by Rose, Tarjan and Lueker [30].

THEOREM 3.11. *[30] Let σ be an LBFS of a chordal graph G and let v be an arbitrary vertex of G. Let W denote the set of vertices that occur before v in σ. Then v is simplicial in the subgraph of G induced by $W \cup \{v\}$.*

The following theorem is equivalent to Theorem 3.11, and will be used, again and again, in our subsequent arguments. We shall refer to it as the P_3 Rule.

THEOREM 3.12. **(The P_3 Rule)** *Let σ be an LBFS of a chordal graph G and let u, v, w be vertices of G with v adjacent to u and w and such that $u <_\sigma v$ and $w <_\sigma v$. Then vertices u and w must be adjacent.*

COROLLARY 3.13. *Let S be an M-slice with respect to LBFS σ, properly contained in chordal graph G. Then all vertices that occur before S and are adjacent to S form a clique.*

Such a clique with respect to S in σ is called *the clique of attachment of S with respect to σ* and is denoted $cl_\sigma(S)$. Note that $cl_\sigma(T) \supseteq cl_\sigma(\Gamma^\sigma_{u,v})$ for any M-$\Gamma^\sigma_{u,v}$, T. Let σ be an LBFS of a chordal graph G.

COROLLARY 3.14. *If vertex u occurs before v in σ, then in every induced u,v-path $\pi: \ u = u_1, u_2, \ldots, u_t = v$ in G there exists a unique subscript j, $(1 \leq j \leq t-1)$, such that, if $j \neq 1$, $u_j <_\sigma u_{j-1} <_\sigma \cdots <_\sigma u_1$ and $u_j <_\sigma u_{j+1} <_\sigma \cdots <_\sigma u_t$.*

Proof. If the statement if false, we find a subscript i, $(2 \leq i \leq t-1)$, such that both u_{i-1} and u_{i+1} occur in σ before u_i. Since the path π is chordless, u_{i-1} and u_{i+1} are not adjacent. However, this violates the P_3 Rule. □

COROLLARY 3.15. *If vertex u occurs before v in σ, then in every induced $u,v-$path π, the neighbour of v on π occurs before v.*

Proof. This follows immediately from Corollary 3.14. □

A path $\pi: v_1, v_2, \cdots, v_k$ is said to be *monotone with respect to LBFS σ* if $v_i <_\sigma v_{i+1}$ for all i, $1 \leq i < k$.

COROLLARY 3.16. *Let $\pi: w_k, w_{k-1}, \cdots, w_1$ be a monotone path with respect to σ and let x be a vertex outside π such that $w_k <_\sigma x <_\sigma w_{k-1}$. If $xw_1 \in E$, then x is universal to π.*

The next theorem illustrates the critical role that slices play in the study of different LBFSs of a chordal graph. This theorem is essential for the correctness of our multisweep LBFS interval graph recognition algorithm and is of significant independent interest.

THEOREM 3.17. **(The (Chordal) LBFS Theorem)** *Let G be a chordal graph and let S be an M-slice of an arbitrary LBFS ordering τ of G. Further let σ be an arbitrary LBFS ordering of G. The restriction of σ to S is an LBFS ordering of the graph induced by the vertices of S.*

Proof. By Lemma 3.4(i), τ_S is an LBFS of S. Thus we may assume that σ and τ are distinct for, otherwise, there is nothing to prove. Let σ_S denote the restriction of σ to S. Suppose that σ_S is not an LBFS ordering of $G[S]$ and let v_1 be the leftmost vertex of σ_S that is inconsistent with an LBFS of $G[S]$. Let w_0 be a vertex of S such that the subordering of σ_S up to but not including v_1 followed by w_0 is consistent with some LBFS of $G[S]$. (i.e. w_0 was a legitimate vertex to be chosen at the time that v_1 was chosen.) Clearly v_1 is before w_0 in σ.

The choice of v_1 and w_0 imply both the existence of vertex w_1 in S such that $w_1 <_\sigma v_1$, $w_1w_0 \in E$ and $w_1v_1 \notin E$ and the nonexistence of vertex x in S with $x <_\sigma w_1$, $xv_1 \in E$ and $xw_0 \notin E$.

By Theorem 3.1, the triple w_1, v_1, w_0 implies the existence of vertex $v_2 <_\sigma w_1$ (without loss of generality the leftmost such vertex) adjacent to v_1 and not adjacent to w_0; by the above argument, $v_2 \notin S$.

We now show that $v_1 <_\tau v_2$. Since v_2 is adjacent to v_1 but not to w_0, v_2 cannot be before S in τ. If S is a slice in τ, then since $v_2 \notin S$, v_2 must be after S in τ. If S is a proper module of slice T in τ, then by the definition of a module, $v_2 \notin T$ and v_2 must be after T in τ and thus also after S. Since $v_1 \in S$, the claim follows.

The edge $v_2 w_1 \notin E$ since otherwise τ has a P_3: v_1, v_2, w_1 contradicting the P_3 Rule. Since $w_1 <_\sigma v_1$, there exists a vertex to the left (in σ) of v_2 that is adjacent to w_1 and not adjacent to v_1. Let w_2 be the leftmost such vertex. We claim that

$$\text{If } v_2 <_\tau w_2, \text{then } w_1 <_\tau w_2. \tag{3.1}$$

To see this, note that since v_2 is after S in τ, w_2 is also after S. Since $w_1 \in S$, (3.1) follows.

We claim that $w_2 v_2 \notin E$. Suppose not. If $w_2 <_\tau v_2$, then the P_3: v_1, v_2, w_2 contradicts the P_3 Rule. If $v_2 <_\tau w_2$, then the P_3: v_2, w_2, w_1 contradicts the P_3 Rule (by (3.1)).

Now suppose that we have a sequence of vertices in σ: $w_k <_\sigma v_k <_\sigma w_{k-1} <_\sigma v_{k-1} \cdots <_\sigma w_1 <_\sigma v_1 <_\sigma w_0$ where $k \geq 2$ such that for all $i, 2 \leq i \leq k$, all of the following hold:

$$\left\{\begin{array}{l} \bullet\ v_i \text{ is the leftmost (in } \sigma) \text{ vertex of } G \text{ with } v_i v_{i-1} \in E \text{ and } v_i w_{i-2} \notin E \\ \bullet\ w_i \text{ is the leftmost (in } \sigma) \text{ vertex of } G \text{ with } w_i w_{i-1} \in E \text{ and } w_i v_{i-1} \notin E \\ \bullet\ v_{i-1} <_\tau v_i \\ \bullet\ v_i w_{i-1} \notin E \\ \bullet\ v_i <_\tau w_i \Rightarrow w_{i-1} <_\tau w_i \\ \bullet\ w_i v_i \notin E \end{array}\right. \tag{3.2}$$

Note that the previously defined sequence $w_2 <_\sigma v_2 <_\sigma w_1 <_\sigma v_1 <_\sigma w_0$ satisfies these properties. We will show how to extend any such sequence to a longer one satisfying the same properties, thereby contradicting the finiteness of G. Since $v_k <_\sigma w_{k-1}$ there is a vertex before w_k in σ, adjacent to v_k and not adjacent to w_{k-1}. Let v_{k+1} be the leftmost (in σ) such vertex. We claim that

$$v_k <_\tau v_{k+1}. \tag{3.3}$$

Suppose $v_{k+1} <_\tau v_k$. We know $v_{k-1} <_\tau v_k$ by (3.2). Therefore $v_{k-1} v_{k+1} \in E$ else the P_3: v_{k-1}, v_k, v_{k+1} contradicts the P_3 Rule in τ. Also $v_{k+1} w_{k-1} \notin E$ by the definition of v_{k+1}. Therefore $v_{k+1} w_{k-2} \notin E$ else the P_3: $v_{k+1}, w_{k-2}, w_{k-1}$ exists in σ. But now the choice of v_k is contradicted (v_{k+1} should have been chosen instead). We claim that

$$v_{k+1} w_k \notin E. \tag{3.4}$$

Suppose $v_{k+1} w_k \in E$. If $w_k <_\tau v_{k+1}$, then τ has the P_3: w_k, v_{k+1}, v_k ($v_k <_\tau v_{k+1}$ by (3.3); $w_k v_k \notin E$ by (3.2). If $v_{k+1} <_\tau w_k$ then τ has the P_3: w_{k-1}, w_k, v_{k+1} ($v_k <_\tau w_k$ by (3.3) and transitivity; $w_{k-1} <_\tau w_k$ by (3.2).

Since $w_k <_\sigma v_k$ there is a vertex to the left of v_{k+1} (in σ) that is adjacent to w_k and not adjacent to v_k. Let w_{k+1} be the leftmost such vertex.

$$\text{If } v_{k+1} <_\tau w_{k+1}, \text{then } w_k <_\tau w_{k+1}. \tag{3.5}$$

Assume $w_{k+1} <_\tau w_k$. Thus $v_{k+1} <_\tau w_k$ and $v_k <_\tau w_k$ (by (3.3)). Now $w_{k-1} <_\tau w_k$ by (3.2) and $w_{k-1} w_{k+1} \in E$, since otherwise τ has the P_3: w_{k-1}, w_k, w_{k+1}. Also

$w_{k+1}v_k \notin E$ and thus $w_{k+1}v_{k-1} \notin E$ (otherwise σ has the P_3: w_{k+1}, v_k, v_{k-1}). Since $w_{k+1} <_\sigma w_k$ this contradicts the choice of w_k.

To complete the proof we show that $w_{k+1}v_{k+1} \notin E$. Suppose $w_{k+1}v_{k+1} \in E$. If $w_{k+1} <_\tau v_{k+1}$ then τ has the P_3: v_k, v_{k+1}, w_{k+1} ($v_kv_{k+1} \in E$ by the definition of v_{k+1}; $v_k <_\tau v_{k+1}$ by (3.3); $w_{k+1}v_k \notin E$ by the definition of w_{k+1}). If $v_{k+1} <_\tau w_{k+1}$ then $w_k <_\tau w_{k+1}$ by (3.5) and $w_kv_{k+1} \notin E$ by (3.4). But now τ has the P_3: v_{k+1}, w_{k+1}, w_k. □

Note that the proof of the LBFS Theorem does not require τ to be an LBFS; it requires only that τ satisfy the P_3 Rule and that all vertices to the left of S in τ be homogeneous to S.

To put the LBFS Theorem in perspective, it is important to note that the theorem does not hold for S being an arbitrary subset of vertices of a chordal (or even interval) graph G. For example, consider the interval graph shown in Figure 3. The numbering of the vertices indicates a legitimate LBFS ordering; however when vertex 1 is removed the restriction of this ordering to the remaining subset is not a legitimate LBFS ordering of the subset. Also, as shown in Figure 4, the theorem does not hold for AT-free graphs. $S = \{2, 3, 4\}$ is a slice of the LBFS: 1 2 3 4 5. Now consider an arbitrary LBFS starting at 5. Vertex 3 occurs after 2 and 4, which cannot occur in an LBFS of S.

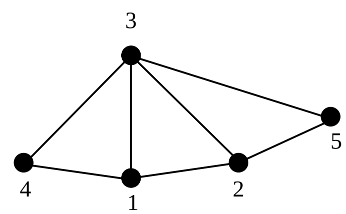

FIG. 3. *The LBFS Theorem does not hold for arbitrary subsets of vertices of interval graphs*

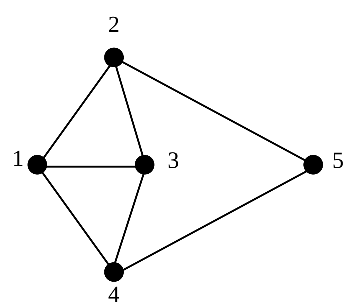

FIG. 4. *The LBFS Theorem does not hold for AT-free graphs*

We now present three corollaries of the LBFS Theorem. Recall that for X, a set of vertices in G we let $sc_\tau(X)$, *(the scope of X in τ)* denote the smallest contiguous set, with respect to the ordering of V imposed by τ, that contains X.

COROLLARY 3.18. *Let S, with components $C_1, C_2 \cdots, C_k, (k \geq 2)$, be a disconnected M-slice of chordal graph G with respect to some LBFS. Then for any LBFS τ of G, $sc_\tau(C_i) \cap sc_\tau(C_j) = \emptyset$, for all $i \neq j$.*

We now turn our attention to good LBFSs in chordal graphs. Recall that in a good LBFS every M-slice (see Lemma 3.4(ii)) S (including V itself) must start with a vertex that can end an LBFS of S.

COROLLARY 3.19. *LBFS+ produces a good LBFS for chordal graphs.*

Proof. Let G be a chordal graph, and let τ^+ be the LBFS produced by LBFS+(τ) where τ is an arbitrary LBFS of G. Let S be any M-slice in τ^+. By the LBFS Theorem, x, the last vertex of $sc_\tau(S)$, is a good vertex of S thereby showing that in τ^+, S starts with a good vertex. □

Note that, for arbitrary graphs, LBFS+ is not guaranteed to produce a good LBFS. For the graph in Figure 4, if τ is the LBFS: 5 4 2 3 1, then τ^+ is 1 [3 2 4] 5, but the slice [3 2 4] does not start with a good vertex.

COROLLARY 3.20. *Let S be an M-slice of chordal graph G with respect to LBFS τ. Let T, a superset of S, be an M-slice with respect to LBFS τ. Then any vertex $x \in S$ that is good with respect to T is also good with respect to S.*

Proof. Let τ' be an LBFS of T that places x last. By the LBFS Theorem, τ'_S is an LBFS of S. Since τ'_S places x last, x is good in S. □

The next lemma is analogous to the LBFS Theorem insofar as it describes the effect that good LBFSs have on M-slices from another good LBFS. Recall that a vertex is called bad in S, where S is an M-slice, if it is not good in S. We say that a set of vertices $X \subset Y$ is *pulled with respect to Y in LBFS τ* if there exists a vertex $z <_\tau sc_\tau(Y)$ such that $zx \in E$ for all $x \in X$ and $zy \notin E$ for all $y \in Y \backslash X$. X is referred to as a *pulled set* and a vertex is said to be *pulled* if it belongs to some pulled set.

LEMMA 3.21. *Let σ and τ be good LBFSs of a chordal graph G, and let S be an M-slice of σ. If τ_S is not a good LBFS of S then s_1, the first vertex of $sc_\tau(S)$, is a bad vertex with respect to S and is pulled in τ.*

Proof. Assume τ_S is not a good LBFS. Look at the point in τ where s_1 was chosen. First assume that the set of tied vertices at this point includes s_2, the last vertex of S in τ. Since τ is a good LBFS, s_1 is a good vertex of S in a superslice of S and thus, by Corollary 3.20, in S too. If the set of tied vertices at the time s_1 was chosen by τ does not include s_2, then there is a vertex $x <_\tau s_1$ such that $xs_1 \in E$ but $xs_2 \notin E$. Note that x is providing a pull on S. If s_1 is bad with respect to S, we are done so we assume s_1 is a good vertex in S.

By the LBFS Theorem (Theorem 3.17), τ_S is an LBFS of S but, by assumption, is not a good LBFS. Thus there is a slice S' of S with respect to τ_S such that y_1, the first vertex of S' is not good in S'. At the time y_1 was chosen by τ, y_2 (the last vertex of $sc_\tau(S')$) could not have been tied since otherwise, since τ is a good LBFS, y_1 would be good in a superslice of S' and thus by Corollary 3.20 in S' too. Thus there exists vertex z, $z <_\tau y_1$ such that $zy_1 \in E$ and $zy_2 \notin E$. Now try to place z in σ: z is not universal to S' and thus is not universal to S, thus $z \notin cl_\sigma(S)$; $z \notin S$ since τ_S is an LBFS of S and all vertices in S' have to be homogeneous with respect to all vertices of S that appear before S' (Claim 3.5). Thus, whether S is a slice or a module of a slice, $z \notin sc_\sigma(S)$. Therefore z is after S in σ and by Claim 3.7 there exists $z' \in cl_\sigma(S)$ such that $zz' \notin E$ (note that z is adjacent to some but not all vertices of S since $zy_2 \notin E$). In τ, z' cannot be before y_1 since otherwise there is a P_3: z, y_1, z'. Since S' is not a clique (otherwise all LBFSs are good), there exists $w_1, w_2 \in S'$ such that $w_1w_2 \notin E$. Since $s_1 <_\tau y_1 <_\tau z'$, $s_1y_1 \in E$ (otherwise there is the P_3: s_1, z', y_1; note $z's_1 \in E$ since $z' \in cl_\sigma(S)$). Thus $s_1 \in cl_\tau(S')$ and $s_1w_1, s_1w_2 \in E$; but now we have a P_3: w_1, s_1, w_2 contradicting s_1 being good in S. □

We end this section with another property of good LBFSs of chordal graphs.

Claim 3.22. *Let σ be a good LBFS of chordal graph G and let S be a nonclique M-slice of σ where $cl_\sigma(S) \neq \emptyset$. Then there is a vertex before the last vertex of $cl_\sigma(S)$ that is not in $cl_\sigma(S)$ and is universal to $cl_\sigma(S)$.*

Proof. Let x be the rightmost vertex in $cl_\sigma(S)$. If x has no neighbour to the left that is not in $cl_\sigma(S)$, then x starts a slice including S, but x is not simplicial in this slice since S is not a clique, thereby contradicting σ being good. Thus x has a neighbour y, where $y <_\sigma x$, that is not in $cl_\sigma(S)$. By the P_3 Rule, y is universal to all vertices of $cl_\sigma(S)$ as required. □

3.3. AT-free graphs. In this subsection we list three important theorems from [11] and prove a new claim regarding separators.

Theorem 3.23. *[11] Let σ be an LBFS of an AT-free graph G and let v be an arbitrary vertex of G. Let W denote the set of vertices w that occur before v in σ. Then v is admissible in the subgraph of G induced by $W \cup \{v\}$.*

Note the similarity between this theorem and Theorem 3.11.

Theorem 3.24. *[11] Let u be an arbitrary vertex of a connected AT-free graph G. Let a be the vertex numbered last by LBFS(u) and let b be the vertex numbered last by LBFS(a). Then (a, b) is a pokable dominating pair in G.*

Theorem 3.25. *[11] Let $G = (V, E)$ be a connected AT-free graph and suppose that G contains no vertices unrelated with respect to vertex x of G. Consider the vertex ordering produced by an LBFS σ starting at x. Then, for all vertices u, v in V with $u <_\sigma v$, we have $u \in D(v, x)$.*

Let A be a minimal separator of graph G. A connected component C_i of $G \setminus A$ is called *shallow* if $\widetilde{C_i} = C_i$ ($\widetilde{C_i}$ denotes the vertices in C_i that are universal to A) and is called *deep* otherwise. As the following claim shows, any AT-free graph has at most two deep components.

Claim 3.26. *Let A be a minimal separator of AT-free graph G with $\{C_i, 1 \leq i \leq k\}$ the deep components of $G \setminus A$. Then $k \leq 2$.*

Proof. Suppose $k \geq 3$ and for $1 \leq i \leq 3$ consider vertex $c_i \in C_i \setminus \widetilde{C_i}$ where c_i is not adjacent to $a_i \in A$. Now, between any pair of $\{c_i\}$ there is a path that avoids the neighbourhood of the third. For example, consider a path from c_1 to a vertex in $\widetilde{C_1}$ to a_3 to a vertex in $\widetilde{C_2}$ to c_2; clearly this path misses the neighbourhood of c_3. □

4. LBFS properties of interval graphs. In this section we present new LBFS structure results on interval graphs. Let S be an M-slice of an interval graph G with respect to some LBFS. A vertex u of S is said to be *S-valid* (or simply *valid* if no confusion is possible) if, in S, u is admissible and is the midpoint of no P_3, i.e. simplicial.

We note that since G itself is a slice with respect to any of its LBFSs, it makes sense to talk about valid and good vertices in G. Theorem 3.12 (the P_3 Rule) and Theorem 3.23 combined, imply that, in any LBFS of an interval graph, every good vertex in some M-slice must also be valid in that M-slice. As it turns out, the converse is also true (Theorem 4.5).

Observation 4.1. *If v is S-valid then it is S'-valid for every M-slice $S' \subset S$ where $v \in S'$.*

The following claim has most likely been observed many times. Note that the 3-sun graph (the left hand graph in Figure 5 with vertex y removed) and C_4 show that it does not hold for either chordal graphs or AT-free graphs respectively.

Claim 4.2. *A nonclique interval graph has diameter two if and only if it contains a universal clique.*

Proof. Any nonclique graph that contains a universal clique has diameter two. Consider an interval representation for nonclique interval graph G of diameter two and let u and v be the vertices represented by the intervals having leftmost right endpoint p and rightmost left endpoint q, respectively. Since G is not a clique, $u \neq v$ and p is to the left of q. Now, the set of vertices corresponding to intervals having left endpoint to the left of p and right endpoint to the right of q forms a universal clique, and at least one such vertex exists since $d(u, v) = 2$. ◻

We now consider the structure of interval graphs in more detail. Specifically, we define the *type* of an interval graph, based on the diameter of the graph after the removal of all universal vertices. Note that there is a correspondence between type 1 and type 2 interval graphs as defined below, and the Q nodes and P nodes, respectively, of [2]. For any nonclique interval graph G we let K_G denote the maximal universal clique of G; by convention, $K_G = \emptyset$ if G does not have a universal clique. Since G is not a clique, only vertices in $G\backslash K_G$ can be good or valid in G. We let $C_1, \cdots, C_k$ $(k \geq 1)$ denote the connected components of $G\backslash K_G$. We say that G is of *type 1* if $G\backslash K_G$ is connected (i.e. $k = 1$, which implies $diam(G\backslash K_G) \geq 3$ by Claim 4.2 and the maximality of K_G) and is of *type 2* if $G\backslash K_G$ is disconnected (i.e. $k \geq 2$). As we shall see, type 2 interval graphs are of limited interest from the LBFS perspective. Neither universal vertices nor disconnected components have any effect on the LBFS restricted to a component. Henceforth when we state that $diam(H) \geq 3$ for some graph H, we imply that H is connected.

The following claim summarizes some facts about LBFSs for interval graphs with diameter two. These facts follow from Theorem 3.17, Corollary 3.18, Corollary 3.20, Claim 4.2, and the observation that universal vertices do not affect the ordering of non-universal vertices produced by the LBFS algorithm.

CLAIM 4.3. *Let G be a nonclique interval graph with a maximal universal clique K_G and let $C_1, C_2, \ldots, C_k$, $(k \geq 1)$, be the connected components of $G \setminus K_G$. Let σ be any LBFS of G in which the first vertex of $G\backslash K_G$ visited is vertex x in C_1. Let σ' be the restriction of σ to $G\backslash K_G$.*

(i) *All vertices in $C_1 \cup K_G$ are visited by σ before any vertices in $C_2 \cup \cdots \cup C_k$.*

(ii) *For any j, $(2 \leq j \leq k)$, the vertices of C_j occur consecutively in σ.*

(iii) *The ordering σ' is an LBFS of $G\backslash K_G$; furthermore, if σ is a good LBFS of G then σ' is a good LBFS of $G\backslash K_G$.*

(iv) *If vertex v is good (respectively valid) in $G\backslash K_G$ then v is good (respectively valid) in G.*

(v) *For $u, v \notin K_G$, $u <_\sigma v$: $\Gamma^{\sigma}_{u,v}\backslash K_G = \Gamma^{\sigma'}_{u,v}$.*

We now state and prove a fundamental result of LBFS in interval graphs that will be key in many of our subsequent arguments.

LEMMA 4.4. **(The Flipping Lemma)** *Let G be an interval graph and let y and z be valid vertices of G. If there is an LBFS of G that starts at y and ends at z, then some LBFS of G starts at z and ends at y.*

Proof. Suppose the statement is true for all interval graphs with fewer vertices than G. If y and z are adjacent, then G is a clique and there is nothing to prove. If G is disconnected, enumerate its connected components as $C_1, C_2, \ldots, C_k$, $(k \geq 2)$, such that $z \in C_1$ and $y \in C_k$ (note that by Claim 3.6, y and z cannot be in the same component of G). Consider an LBFS of C_k starting at y and ending at a vertex y'. Trivially, y' is a valid vertex of C_k. Moreover, by Observation 4.1, since y is valid in G it is valid in C_k. By the induction hypothesis, some LBFS τ of C_k starts at y' and ends at y. Now, there exists an LBFS of G that starts at z and, having exhausted C_1, visits in order $C_2, \ldots, C_{k-1}$ and then, visiting C_k in the same order as τ, ends at y.

Therefore, from now on, we shall assume that y and z are nonadjacent and that G is connected. Consider an LBFS σ of G that starts at z and ends at a vertex w distinct from y. We claim that

$$\Gamma^{\sigma}_{y,w} \subset V. \tag{4.1}$$

To justify (4.1), observe that if $\Gamma^{\sigma}_{y,w} = V$ then, since G is connected, Theorem 3.10 guarantees the existence of a prior z, y-path all of whose internal vertices are missed by w. Since (y, z) is a dominating pair in G, it must be the case that $yw \in E$. However, now y is the midpoint of a P_3, contradicting that y is valid. Thus, (4.1) must hold.

Consider an arbitrary LBFS of $\Gamma^{\sigma}_{y,w}$ that starts at y and ends at some vertex w'. By the induction hypothesis (which can be applied by (4.1)), there exists an LBFS τ of $\Gamma^{\sigma}_{y,w}$ that starts at w' and ends at y. But now, by combining σ and τ in the obvious way, we obtain an LBFS of G that starts at z and ends at y, completing the proof. ▯

Note that the Flipping Lemma does not extend to either chordal or AT-free graphs. For the graphs in Figure 5 with the given initial LBFSs that end at vertex x, there is an LBFS from x that ends at y. Thus both x and y are good vertices but there is no LBFS that starts at y and ends at x. (In fact all LBFSs from y must end at z.)

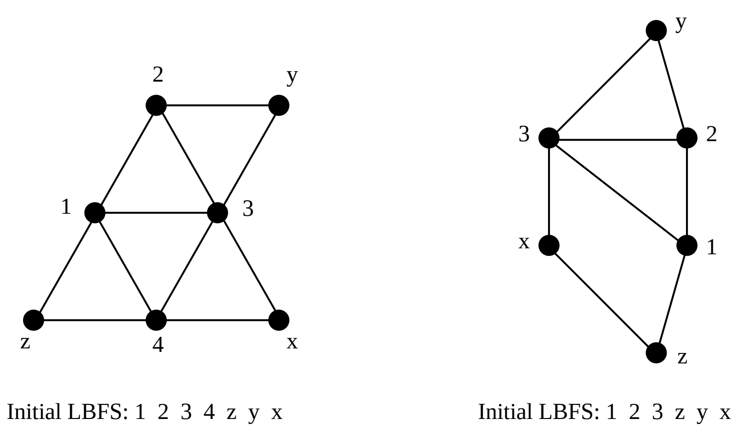

FIG. 5. *The Flipping Lemma does not extend to chordal graphs or to AT-free graphs*

The Flipping Lemma is related to the following results, the first of which shows the equivalence between valid and good vertices in M-slices of interval graphs. As mentioned in §2, to the best of our knowledge this is the first characterization of good vertices for any nontrivial family of graphs.

THEOREM 4.5. *Let S be an M-slice of an interval graph with respect to some LBFS. A vertex of S is good if and only if it is valid, i.e., simplicial and admissible.*

Proof. The "only if" part follows from Theorem 3.12 (the P_3 Rule) and Theorem 3.23. To prove the "if" part, let u be a valid vertex in S and let σ be an LBFS of S starting at u. Let v be the vertex of S occurring last in σ. Since v is valid, by the Flipping Lemma (Lemma 4.4) there exists an LBFS of S that starts at v and places u last, confirming that u is a good vertex. ▯

By Theorem 4.5, when referring to an M-slice S, good and valid vertices are synonyms. In the remainder of this work we shall use *good* to stand for both valid and good vertices. The exact property that we have in mind will be clear from the context. Note that one of the implications of Theorem 4.5 holds in general, specifically for an arbitrary graph, if a vertex is valid in M-slice S then it is also good in S [8].

The next theorem shows that LBFS+(G, σ) "flips" an LBFS σ and produces a good LBFS provided σ starts with a good vertex.

THEOREM 4.6. *Let G be an interval graph and let y and z be good vertices of G. If σ is an arbitrary LBFS of G that starts at y and ends at z, then LBFS+(G, σ) is a good LBFS that starts at z and ends at y.*

Proof. Let σ^+ be the ordering produced by LBFS+(G, σ). By Corollary 3.19, σ^+ is a good LBFS. Suppose that σ^+ ends at a vertex w distinct from y. Let C be the connected component of $\Gamma^{\sigma^+}_{y,w}$ containing y, and let v be the first vertex of C visited by σ^+. Observe that v and y are distinct vertices; this is because y has the lowest LBFS number in σ among all the vertices in G and, therefore, cannot be chosen first by σ^+.

We note, further, that in σ, v must occur after w. By Theorem 3.10 applied to σ^+, there exists a v, y-path P in C, all of whose internal vertices are missed by w. By Theorem 3.25, applied to σ, $w \in D(v, y)$, forcing w to be adjacent to y. But now, y is the midpoint of the P_3 with endpoints w and y's neighbour on P, contradicting y being good. ∎

Let S be an M-slice of an interval graph G with respect to some good LBFS. Two good vertices u and v of S are said to be *antipodal in S* if there exists a good LBFS of S that starts at u and ends at v.

CLAIM 4.7. *Let C be a connected M-slice of an interval graph G with respect to some good LBFS and let u and v be antipodal vertices in C. Then $d_C(u, v) = diam(C)$.*

Proof. Dragan et al. [13] have shown that the last vertex of any LBFS of a connected interval graph G has eccentricity equal to the diameter of G. Thus $ecc_C(v) = diam(C)$. Since there is an LBFS starting at v that ends at u (by the Flipping Lemma), u is at maximum distance from v and thus $d_C(u, v) = diam(C)$. ∎

The next lemma characterizes type 1 M-slices (i.e. M-slices S where $diam(S \backslash K_S) \geq 3$).

LEMMA 4.8. *Let S be a nonclique M-slice in a good LBFS σ of interval graph G. S is of type 1 if and only if there exists in S an induced path of length greater than two from the first vertex of S in σ to the last vertex of S in σ.*

Proof. Since σ is good, neither the first nor the last vertex of σ_S belongs to K_S. If S is of type 1 then $S \backslash K_S$ is connected and has $diam(S \backslash K_S) \geq 3$. Thus, by Claim 4.7, a shortest path in $S \backslash K_S$ from the first to the last vertex is the required path. To prove the "if" direction, we note that an induced path of length greater than two between the first and last vertices of $sc_\sigma(S)$ cannot contain a vertex of K_S. Thus $S \backslash K_S$ is connected and has diameter greater than two by Claim 4.2 and the maximality of K_S. ∎

We now prove a claim that will be used often in the proof of correctness of the interval graph recognition algorithm.

CLAIM 4.9. *Let S be an M-slice of an interval graph G with respect to some LBFS. Let σ be a good LBFS of S and let b, c, d be vertices of S such that*

- *$b <_\sigma c$ and $c <_\sigma d$,*
- *c misses a b, d-path π in S.*

Then $d \notin \Gamma^\sigma_{b,c}$.

Proof. Suppose that $d \in \Gamma^{\sigma}_{b,c}$. We claim that

$$\text{every vertex of } \pi \text{ belongs to } \Gamma^{\sigma}_{b,c}. \tag{4.2}$$

If (4.2) is false, then by Corollary 3.14 there must exist vertices of π that occur in σ prior to $\Gamma^{\sigma}_{b,c}$. Let u be the first such vertex encountered while traversing π in the direction from b to d. The choice of u guarantees that the previous vertex on π belongs to $\Gamma^{\sigma}_{b,c}$. By Claim 3.5, u must be adjacent to all the vertices in $\Gamma^{\sigma}_{b,c}$ and, in particular, to c, a contradiction. Thus, (4.2) must hold.

Let C be the connected component of $\Gamma^{\sigma}_{b,c}$ containing b. By (4.2), all vertices of π must also belong to C and by Claim 3.6 $c \in C$. Let t be the first vertex visited by σ in $\Gamma^{\sigma}_{b,c}$. By the definition of $\Gamma^{\sigma}_{b,c}$, t is also the first vertex of C. By Theorem 3.10, there exists a prior t, b-path π' in $\Gamma^{\sigma}_{b,c}$ all of whose internal vertices are missed by c. But now, $\pi \cup \pi'$ contains a t, d-path in $\Gamma^{\sigma}_{b,c}$ missed by c, contradicting Theorem 3.25. ∎

CLAIM 4.10. *Let C be a connected M-slice of an interval graph with respect to some LBFS and let u and v be antipodal good vertices in C. Let τ be an LBFS of C that starts at v and ends at u, and let σ be an LBFS of C that starts at u and ends at v. Then, for any good vertex w of C, at least one of the conditions $\Gamma^{\sigma}_{w,v} \subset C$ or $\Gamma^{\tau}_{w,u} \subset C$ must be satisfied.*

Proof. If w coincides with one of the vertices u or v, then there is nothing to prove. Similarly, if u and v are adjacent, then C is a clique and the conclusion is immediate.

We shall, therefore, assume that w is distinct from both u and v and that u and v are nonadjacent. If the statement is false, we have $\Gamma^{\sigma}_{w,v} = \Gamma^{\tau}_{w,u} = C$. Theorem 3.10 applied to σ guarantees the existence of a prior u, w-path π in C all of whose internal vertices are missed by v. In fact, v misses π completely. To see this, observe that v cannot be adjacent to w, for otherwise w would be the midpoint of a P_3 consisting of v and w's neighbour on π; a contradiction to w being good in C.

Further, Theorem 3.10 applied to τ guarantees the existence of a prior v, w-path π' in C all of whose internal vertices are missed by u. It is easy to see that, in fact, u misses the path π' completely. But now, the two paths π and π' confirm that u and v are unrelated vertices with respect to w, contradicting that w is good, in particular, it is not admissible. ∎

Our next result specializes Claim 4.10 to M-slices of type 1.

THEOREM 4.11. *Let C be a type 1 M-slice of an interval graph G with respect to some LBFS and let u and v be antipodal good vertices in C. Let τ be an LBFS of C that starts at v and ends at u, and let σ be an LBFS of C that starts at u and ends at v. Then, exactly one of the conditions $\Gamma^{\sigma}_{w,v} \subset C$ or $\Gamma^{\tau}_{w,u} \subset C$ must be satisfied.*

Proof. By Claim 4.10 at least one of the conditions $\Gamma^{\sigma}_{w,v} \subset C$ or $\Gamma^{\tau}_{w,u} \subset C$ must hold. Suppose that both of the conditions are satisfied. Let σ' and τ' be the restrictions of σ and τ respectively to $C \setminus K_C$. Vertices u, v, and w belong to $C \backslash K_C$ since they are good in C (note that τ ends at u). Thus, Claim 4.3(v) applies and $\Gamma^{\sigma}_{w,v} \backslash K_C = \Gamma^{\sigma'}_{w,v} \subset C \backslash K_C$ and $\Gamma^{\tau}_{w,u} \backslash K_C = \Gamma^{\tau'}_{w,u} \subset C \backslash K_C$. Thus we may turn our attention to σ' and τ'.

First, since $\Gamma^{\sigma'}_{w,v} \subset C \backslash K_C$, the connectedness of $C \backslash K_C$ guarantees the existence of a vertex x in $C \setminus (\Gamma^{\sigma'}_{w,v} \cup K_C)$ which, by Claim 3.5, is adjacent to all vertices in $\Gamma^{\sigma'}_{w,v}$. But now, we have reached a contradiction: If $x \in C \setminus (\Gamma^{\tau'}_{w,u} \cup K_C)$ then, by Claim 3.5, x is adjacent to u (since x is adjacent to w), implying that $d_{C \backslash K_C}(u, v) \leq 2$ which implies $diam(C \setminus K_C) \leq 2$ by Claim 4.7. If $x \in \Gamma^{\tau}_{w,u}$ then, by Claim 3.5, $d(u,v) = 1$ and $C \setminus K_C$ is a clique. □

CLAIM 4.12. *Let S be an M-slice of an interval graph G with respect to some good LBFS. Let u and v be arbitrary good vertices of S and let σ be a good LBFS of S that places v last. Then, u and v are antipodal vertices in $\Gamma^{\sigma}_{u,v}$.*

Proof. We proceed by induction on $|\Gamma^{\sigma}_{u,v}|$. If $\Gamma^{\sigma}_{u,v}$ is a clique, then there is nothing to prove. If $\Gamma^{\sigma}_{u,v}$ is disconnected, enumerate its connected components as $C_1, C_2, \ldots, C_k$, $(k \geq 2)$, such that $u \in C_1$ and $v \in C_k$. Now, u and v are antipodal as confirmed by a good LBFS of $\Gamma^{\sigma}_{u,v}$ that starts at u and, having exhausted C_1, visits in order $C_2, \ldots, C_{k-1}$ and then visits C_k in the same order as σ.

Therefore, from now on, we assume that $\Gamma^{\sigma}_{u,v}$ is connected but not a clique. Suppose $\Gamma^{\sigma}_{u,v}$ starts with w and consider a good LBFS τ of $\Gamma^{\sigma}_{u,v}$ that starts at v and ends at w. By Claim 4.10, $\Gamma^{\tau}_{u,w} \subset \Gamma^{\sigma}_{u,v}$ and by Corollary 3.20, u is a good vertex in $\Gamma^{\tau}_{u,w}$. By the induction hypothesis applied to $\Gamma^{\tau}_{u,w}$, u and w are antipodal vertices in $\Gamma^{\tau}_{u,w}$. In other words, there exists a good LBFS θ of $\Gamma^{\tau}_{u,w}$ that starts at w and ends at u. But now, by combining τ and θ in the obvious way we obtain a good LBFS of $\Gamma^{\sigma}_{u,v}$ that starts at v and ends at u confirming that u and v are antipodal. □

Claim 4.12 can be extended as follows.

CLAIM 4.13. *Let S be an M-slice of an interval graph G with respect to some good LBFS. Let u and v be arbitrary good vertices of S and let σ be a good LBFS of S that places u before v. Then u and v are antipodal vertices in $\Gamma^{\sigma}_{u,v}$.*

Proof. Let t be the last vertex in $\Gamma^{\sigma}_{u,v}$. If t and v coincide, the conclusion is implied by Claim 4.12. Therefore, we assume that v and t are distinct. By Corollary 3.20, v is a good vertex in $\Gamma^{\sigma}_{v,t}$. Claim 4.12 guarantees that v and t are antipodal vertices in $\Gamma^{\sigma}_{v,t}$. Thus, there exists a good LBFS τ of $\Gamma^{\sigma}_{v,t}$ that starts at t and places v last. We can now apply Claim 4.12 to the good LBFS of $\Gamma^{\sigma}_{u,v}$ obtained by combining σ and τ in the obvious way, confirming that u and v are antipodal vertices in $\Gamma^{\sigma}_{u,v}$. □

Claims 4.12 and 4.13 can be specialized to slices of type 1 as follows.

CLAIM 4.14. *Let C be a type 1 M-slice of an interval graph G with respect to some good LBFS. Let u and v be good vertices of C and let σ be a good LBFS of C that places u before v. Then u and v are antipodal vertices in C if and only if $\Gamma^{\sigma}_{u,v} \supseteq C \backslash K_C$.*

Proof. Notice that, by Claim 4.13, u and v are antipodal vertices in $\Gamma^{\sigma}_{u,v}$. If $\Gamma^{\sigma}_{u,v} \supseteq C \backslash K_C$, then u and v are antipodal vertices in C, since vertices of K_C are not good in C.

Conversely, suppose that u and v are antipodal vertices in C. By Claim 4.7, $diam(C \backslash K_C) \geq 3$ guarantees that $d_{C \backslash K_C}(u,v) \geq 3$. We now look at σ', the restriction of σ to $C \setminus K_C$.

Let a be the first vertex visited by σ' in $C \setminus K_C$. If $a \in \Gamma^{\sigma'}_{u,v}$, then $\Gamma^{\sigma'}_{u,v} = C \setminus K_C$ which in turn, by Claim 4.3(v) implies $\Gamma^{\sigma}_{u,v} \supseteq C \backslash K_C$. Thus assume $a \notin \Gamma^{\sigma'}_{u,v}$. The connectedness of $C \setminus K_C$ guarantees that some vertex y of $\Gamma^{\sigma'}_{u,v}$ is adjacent to some vertex x in $C \setminus (\Gamma^{\sigma}_{u,v} \cup K_C)$ with $x <_{\sigma'} y$. By Claim 3.5, x is adjacent to all the vertices in $\Gamma^{\sigma'}_{u,v}$. But now, $d_{C \backslash K_C}(u,v) = 2$, a contradiction. □

Let S be an M-slice of an interval graph G with respect to some good LBFS. Two vertices u and v of S are said to be *clones in S with respect to some vertex w of S* if there exist two good LBFSs of S starting at w and ending, respectively, at u and v. We will also say that u and v are clones in S without specifying w. Our next result shows that good vertices in an M-slice of an interval graph are related to one another by the "antipodal" relation or by the "clone" relation, or both. Good vertices can be both clones and antipodal in M-slice S when S is a clique of size greater than two or a type 2 interval graph in which $S \setminus K_S$ has more than two connected components.

CLAIM 4.15. *Let S be an M-slice of an interval graph G with respect to some good LBFS and let u and v be two good vertices in S. Then u and v are antipodal or clones in S.*

Proof. Let $C_1, C_2, \ldots, C_k$, $(k \geq 1)$, be the connected components of S and assume without loss of generality that $u \in C_i$ and $v \in C_j$, for some $1 \leq i, j \leq k$. Notice that, by Corollary 3.20, u and v are good vertices of C_i, and C_j, respectively. We claim that

$$\text{if } i \neq j \text{ then } u \text{ and } v \text{ are antipodal in } S. \tag{4.3}$$

To justify (4.3), we only need construct a good LBFS of S that starts at u and ends at v. Such an LBFS starts at u and, having exhausted C_i, visits in some order $S \setminus (C_i \cup C_j)$ and then visits C_j starting at an antipodal vertex of v and ends at v. Thus, (4.3) must hold.

Next, we claim that

$$\text{if } i = j \text{ and } S \text{ is disconnected then } u \text{ and } v \text{ must be clones.} \tag{4.4}$$

To see this, let w be a good vertex in some component C_t distinct from C_i. Consider the good LBFS of S that starts at w and, after exhausting C_t, proceeds with $S \setminus (C_i \cup C_t)$ and then visits C_i starting at an antipodal vertex of u and ends at u. A good LBFS starting at w and ending at v is constructed similarly, confirming that u and v are clones. Thus, (4.4) must hold.

Finally, to complete the proof, we show that

$$\text{if } S \text{ is connected then } u \text{ and } v \text{ are either clones or antipodal vertices.} \tag{4.5}$$

Consider a good LBFS σ of S starting at u and ending at a vertex w. If v and w coincide, then u and v are antipodal. Otherwise, consider $\Gamma^{\sigma}_{v,w}$. If $\Gamma^{\sigma}_{v,w} \subset S$ then by Claim 4.12 v and w are antipodal in $\Gamma^{\sigma}_{v,w}$, confirming that u and v are antipodal in S.

Thus, from now on, we assume that $\Gamma^{\sigma}_{v,w} = S$. By Theorem 4.6, there exists a good LBFS τ of S that starts at w and ends at u. By Claim 4.10, $\Gamma^{\tau}_{v,u} \subset S$ and by Claim 4.12, u and v are antipodal in $\Gamma^{\tau}_{v,u}$ confirming that there exists a good LBFS of S starting at w and ending at v. Thus, in this case, u and v are clones. $\square$

CLAIM 4.16. *Let S be an M-slice of an interval graph G with respect to some good LBFS and let u, v, w be arbitrary good vertices in S. At least one of the pairs (u, v), (u, w), (v, w) is a pair of clones in S.*

Proof. If u and v are clones in S, then there is nothing to prove. Otherwise, by Claim 4.15, u and v must be antipodal in S.

If S is disconnected, then by Claim 3.6, u and v belong to distinct components C and C' of S. Proceeding as in the proof of Claim 4.15, it is easy to see that in this case either u and w, or v and w are clones.

Now, S is connected. Let σ be an LBFS of S that starts at u and ends at v. If $\Gamma^{\sigma}_{w,v} \subset S$ then, by Claim 4.12, v and w are antipodal vertices in $\Gamma^{\sigma}_{w,v}$, confirming that v and w are clones in S. We assume, therefore, that $\Gamma^{\sigma}_{w,v} = S$. By Theorem 4.6, there exists a good LBFS τ that starts at v and ends at u. By Claim 4.10 it must be the case that $\Gamma^{\tau}_{w,u} \subset S$. In this case, Claim 4.12 guarantees that u and w are antipodal in $\Gamma^{\tau}_{w,u}$, confirming that they are clones in S. $\square$

CLAIM 4.17. *Let C be a type 1 M-slice of an interval graph G with respect to a good LBFS. Any two clones in C are distance at most two apart in $C \setminus K_C$.*

Proof. Let u and v be clones in C. The fact that u and v are clones is confirmed by a vertex a of C and two good LBFSs σ and τ starting at a and ending, respectively, at u and v. We examine σ' and τ', the restrictions of σ and τ to $C \setminus K_C$.

If $\Gamma^{\sigma'}_{v,u} \subset C \setminus K_C$, then the connectedness of $C \setminus K_C$ along with Claim 3.5 guarantees the existence of a vertex w in $C \setminus (\Gamma^{\sigma'}_{v,u} \cup K_C)$ adjacent to both u and v confirming that $d_{C \setminus K_C}(u, v) \leq 2$.

We shall assume, therefore, that $\Gamma^{\sigma'}_{v,u} = C \setminus K_C$. By Theorem 3.10, there exists a chordless prior a, v-path π all of whose internal vertices are missed by u. In fact, u misses π entirely, for otherwise, $uv \in E$ and we're done. But now, we have reached a contradiction: in τ, u misses π contradicting the fact that (a, v) is a dominating pair, that is, contradicting Theorem 3.24. □

Our next result provides a characterization of clones in connected components of slices where the component is of type 1.

CLAIM 4.18. *Let C be a type 1 M-slice of an interval graph G with respect to a good LBFS. Two good vertices are clones of C if and only if in $C \setminus K_C$ they are distance at most two apart.*

Proof. One implication follows immediately from Claim 4.17. To prove the other implication, let w and v be good vertices of C with $d_{C \setminus K_C}(w, v) \leq 2$ and let σ be a good LBFS of C starting at some vertex u and ending at v; since v is a good vertex, such an LBFS must exist. We now look at σ', the restriction of σ to $C \setminus K_C$. By Claim 4.7, $diam(C \setminus K_C) \geq 3$ implies that $d_{C \setminus K_C}(u, v) \geq 3$. By Claim 4.12, w and v are antipodal vertices in $\Gamma^{\sigma'}_{w,v}$. Thus there is an LBFS τ of $\Gamma^{\sigma'}_{w,v}$ that starts at w and ends at v. Since $d_{C \setminus K_C}(u, v) \geq 3$ and $d_{C \setminus K_C}(w, v) \leq 2$, $w \neq u$ and thus $\Gamma^{\sigma'}_{w,v} \neq C \setminus K_C$ and $\Gamma^{\sigma}_{w,v} \neq C$, establishing that $\Gamma^{\sigma}_{w,v} \subset C$. But now, by Claims 4.14 and 4.15 w and v are clones, as claimed. □

CLAIM 4.19. *Let C be a type 1 M-slice of an interval graph G with respect to a good LBFS σ, let u and v be, respectively, the first and last vertex of C visited by σ, and let w be an arbitrary good vertex in C. Then either u and w are clones or else v and w are clones in C, but not both.*

Proof. By Claim 4.7, $d_{C \setminus K_C}(u, v) \geq 3$ and thus u and v are not clones by Claim 4.18. By Claim 4.16, u and w are clones or v and w are clones (and possibly both). Consider good LBFS τ of C that starts at v and ends at u, guaranteed by Theorem 4.6. By Theorem 4.11, exactly one of the conditions $\Gamma^{\sigma}_{w,v} \subset C$ or $\Gamma^{\tau}_{w,u} \subset C$ holds. Now Claim 4.12 guarantees that u, w or v, w are antipodal in C and, therefore, by Claims 4.7 and 4.18, not clones. □

We now examine type 1 interval graphs in more detail.

THEOREM 4.20. *Let C be an M-slice in a good LBFS of an interval graph. If C is of type 1 (i.e. $diam(C \setminus K_C) \geq 3$) then:*

(i) *The set of good vertices in C partitions into disjoint sets X_1 and X_2 such that*
- *good vertices u and v are antipodal in C if and only if they belong to distinct X_is;*
- *good vertices u and v are clones in C if and only if they belong to the same X_i.*

(ii) *There exist disjoint cliques K_1 and K_2 in $C \setminus K_C$, universal to X_1 and X_2, respectively, where K_1 and K_2 are minimal separators in $C \setminus K_C$ between X_1 and X_2.*

(iii) *K_1 and K_2 are universal to X_1' and X_2', respectively, where $X_1' \supseteq X_1$ (respectively $X_2' \supseteq X_2$) is the set of all vertices that K_1 (respectively K_2) separates from X_2 (respectively X_1).*

See Figure 6.

Proof. Let u and v be the first and last vertex of C, respectively, visited by a good LBFS σ of C. First, let X_1 be the set of all vertices of C that are clones of u, and let X_2 be the set of vertices of C that are clones of v. By Claim 4.19 every good vertex in C is either a clone of u or a clone of v but not both, confirming that

$$X_1 \cap X_2 = \emptyset. \tag{4.6}$$

Let w be the earliest vertex in X_2 visited by σ (possibly $w = v$). If $w \neq v$, then since w and v are clones, $d_{C \setminus K_C}(w, v) \leq 2$ by Claim 4.18. Therefore, w and v are not antipodal in C by Claim 4.7. Thus, $\Gamma^{\sigma}_{w,v} \subset C$, by Claim 4.14. Note that if $w = v$, then $\Gamma^{\sigma}_{w,v} = \{v\}$. Now consider σ', the restriction of σ to $C \setminus K_C$. The connectedness of $C \setminus K_C$ along with Claim 3.5 guarantees that some vertex a of $C \setminus K_C$ is universal to $\Gamma^{\sigma'}_{w,v}$ and therefore to $\Gamma^{\sigma}_{w,v}$ as well since $\Gamma^{\sigma}_{w,v} \setminus \Gamma^{\sigma'}_{w,v}$ contains only vertices of K_C. Thus, a is universal to X_2. By Corollary 3.13, the set of vertices in $C \setminus \Gamma^{\sigma'}_{w,v}$ that are universal to $\Gamma^{\sigma'}_{w,v}$ (and thus to X_2) forms a clique, denoted K_2; clearly $K_2 \cup K_C$ separates $\Gamma^{\sigma}_{w,v}$ from the rest of C. We now show that K_2 is a minimal separator in $C \setminus K_C$ between $\Gamma^{\sigma'}_{w,v}$ and $C \setminus (K_2 \cup \Gamma^{\sigma'}_{w,v})$. Suppose not and let $K_2' \subset K_2$ be such a minimal separator with $x \in K_2 \setminus K_2'$. Since K_2' is a minimal separator in $C \setminus K_C$ between $\Gamma^{\sigma'}_{w,v}$ and $C \setminus (K_2' \cup \Gamma^{\sigma'}_{w,v})$, x has no neighbours in $C \setminus (K_2 \cup K_C \cup \Gamma^{\sigma'}_{w,v})$. Thus x is the first vertex of a slice that contains $\Gamma^{\sigma'}_{w,v}$. If the slice is not a clique, then x is not simplicial in the slice and we have contradicted the fact that σ is a good LBFS. If the slice is a clique (possibly just containing x and v), then x is good but it is also a good vertex of C, and thus in X_2, contradicting the choice of w being the earliest vertex in X_2 visited by σ. We now have $X_2' = \Gamma^{\sigma}_{w,v}$ and thus K_2 is universal to X_2'.

A similar argument applied to a good LBFS of C starting at v and ending at u establishes the existence of clique K_1 that is universal to X_1 and is a minimal separator in $C \setminus K_C$ between X_1 and X_2 and is universal to X_1'.

The assumption $diam(C \setminus K_C) \geq 3$ ensures that K_1 and K_2 are disjoint, thereby settling (ii) and (iii). In turn, this guarantees that every two vertices in X_2 are distance at most two apart in $C \setminus K_C$. By Claim 4.18, all vertices in X_2 are clones. A similar argument shows that all vertices in X_1 are clones. Our discussion in conjunction with (4.6) guarantees that two good vertices of C are clones if and only if they belong to the same X_i. Finally, let x_1 and x_2 be arbitrary vertices in X_1 and X_2, respectively. By the above argument, $d_{C \setminus K_C}(x_1, x_2) \geq 3$ and thus by Claims 4.15 and 4.18, x_1 and x_2 are antipodal and not clones. □

CLAIM 4.21. *Let C be a type 1 M-slice of an interval graph G with respect to a good LBFS σ. In every LBFS of C all vertices of X_1' occur before all vertices of X_2' or vice versa.*

Proof. By consideration of distances in $C \setminus K_C$, if some LBFS of C starts at some vertex in $X_1' \cup K_1$ then, trivially, X_1' will be visited before X_2'. Similarly, if some LBFS starts at a vertex in $X_2' \cup K_2$, then all vertices in X_2' will be visited before any vertex in X_1'.

Suppose, therefore, that some LBFS starts at a vertex x outside $X_1' \cup K_1 \cup X_2' \cup K_2$. Let z be the first vertex of $X_1' \cup K_1 \cup X_2' \cup K_2$ visited by this LBFS. Since, as observed above, both K_1 and K_2 are cutsets of $C \setminus K_C$, it must be the case that $z \in K_1 \cup K_2$.

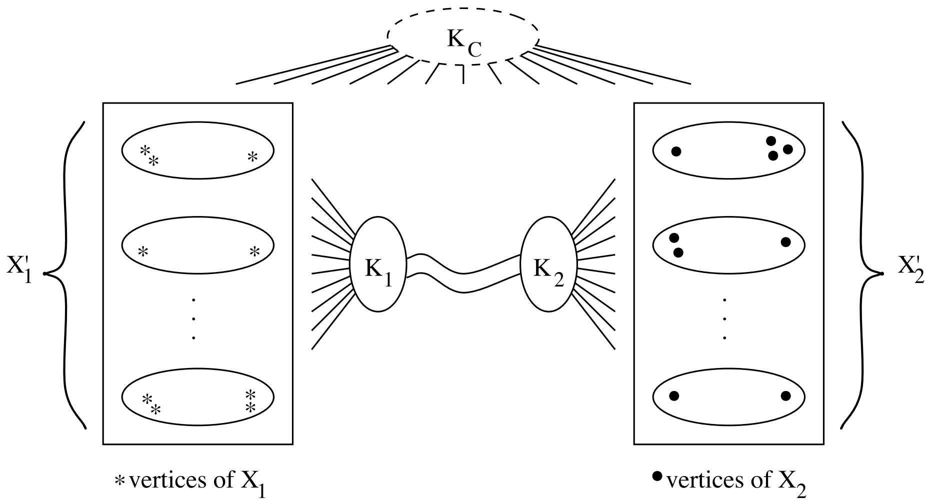

FIG. 6. *The structure of type 1 interval graphs as described in Theorem 4.20*

Since $K_1 \cap K_2 = \emptyset$, we may assume, without loss of generality, that $z \in K_1$. But now, all the vertices in X'_1, and no vertices in X'_2, inherit the label of z. The conclusion follows. ▯

CLAIM 4.22. *Let u and v be nonadjacent good vertices of interval graph G and let σ be a good LBFS of G that ends at v. Consider $A = \widetilde{\Gamma}^{\sigma}_{u,v}$. (Note that u and v are antipodal in A, by Claim 4.12.) Now let A_1 (respectively A_2) be the clones of u (respectively v) in A. Then, $A_1 \cap A_2 = \emptyset$ and, in every good LBFS of G, either all vertices of A_1 are before all vertices of A_2 or vice versa.*

Proof. The proof follows from elementary properties of LBFS, the LBFS Theorem and previous results. Since G is not a clique and A is an M-slice defined by two nonadjacent good vertices, there is no universal clique in A and $diam(A) \geq 3$ or A is disconnected. If $diam(A) \geq 3$, the claim follows from Claim 4.21. If A is disconnected, it follows from Corollary 3.18. ▯

Let S be a subset of the vertices of an interval graph. A vertex $w \in S$ is called *an S-flyer with respect to LBFS σ* (or *a flyer of S in σ*) and is said to *fly with respect to S*, if w is adjacent to a vertex x that occurs, in σ, after all of the vertices in S. We let F_S denote the set of S-flyers in S.

CLAIM 4.23. *Let C be an M-slice of an interval graph G with respect to some LBFS σ of G. If a vertex w of C is a C-flyer with respect to σ, then either w is a good vertex in C or else it is adjacent to a good vertex in C.*

Proof. Let w be a counterexample: w is not good in C and none of its neighbours is good in C, yet it flies to a vertex x that occurs after C in σ. Let $\mathcal{L}$ be the set of all LBFS orderings of C. From $\mathcal{L}$ select an LBFS ρ such that for some suitably chosen vertex z of C, z is the last vertex of C visited by ρ and $|\,\Gamma^{\rho}_{w,z}|$ is as small as possible. We note that since w cannot be the last vertex of any LBFS in $\mathcal{L}$, such a vertex $z \neq w$ must exist. We claim that

$$\Gamma^{\rho}_{w,z} \text{ is connected.} \tag{4.7}$$

To see that this must be the case, let C' be the connected component of $\Gamma^{\rho}_{w,z}$ containing

w and let z' be the last vertex in C'. Note that if w were a good vertex in C' it would also be a good vertex in C. Therefore, w cannot be good in C', and so w and z' are distinct vertices. Since $\Gamma^{\rho}_{w,z'} \subseteq C' \subseteq \Gamma^{\rho}_{w,z}$, our choice of ρ guarantees that $\Gamma^{\rho}_{w,z'} = \Gamma^{\rho}_{w,z} = C'$. Thus, (4.7) must hold.

Let τ be an LBFS of $\Gamma^{\rho}_{w,z}$ starting at z and ending at some vertex a. Since w is adjacent to no good vertices in C, $aw,\ wz \notin E$. The minimality of $|\,\Gamma^{\rho}_{w,z}|$ guarantees that $\Gamma^{\rho}_{w,z} = \Gamma^{\tau}_{w,a}$. We claim that x is not admissible in the subgraph of G induced by $W \cup \{x\}$ where W is the set of vertices that appear before x in σ, thereby contradicting Theorem 3.23. Note that $xz,\ xa \notin E$ by the P_3 Rule (Theorem 3.12) applied to σ. Consider, by the Flipping Lemma (Lemma 4.4), an LBFS θ of $\Gamma^{\rho}_{w,z}$ from a to z. Again by the minimality assumption $\Gamma^{\rho}_{w,z} = \Gamma^{\theta}_{w,z}$. Now, the following paths demonstrate that x is not admissible: a prior path in $\Gamma^{\tau}_{w,a}$ from z to w (guaranteed by Theorem 3.10) and the edge wx (missed by a), and a prior path in $\Gamma^{\theta}_{w,z}$ from a to w (guaranteed by Theorem 3.10) and the edge wx (missed by z). ◻

Claim 4.24. *Let S be a connected M-slice of interval graph G with respect to good LBFS σ. For all x', y' that occur after S in σ and have neighbours in S, $N(x') \cap S \subseteq N(y') \cap S$ or $N(y') \cap S \subseteq N(x') \cap S$.*

Proof. If not then there exist $x, y \in S$ and $x', y' \notin S$ such that $xx', yy' \in E$ and $xy', yx' \notin E$. If $x'y' \in E$, then the P_3 Rule has been violated (assume without loss of generality that $x' <_{\sigma} y'$ and consider the $P_3 : yy'x'$). Thus $x'y' \notin E$ but now G has an AT: x', y', and a vertex z that is universal to $cl_{\sigma}(S)$ but not in $cl_{\sigma}(S)$, guaranteed by Claim 3.22. Note that z is not adjacent to x' or y' or else the P_3 Rule is violated in σ by $zx'x$ or $zy'y$. Paths demonstrating the AT are induced on: $z, \tilde{y}, x, x'$; $z, \tilde{x}, y, y'$; and x', x, an x, y-path in S, y, y'; where $\tilde{x} \in cl_{\sigma}(S)$ is not adjacent to x' and $\tilde{y} \in cl_{\sigma}(S)$ is not adjacent to y' (guaranteed by Claim 3.7). ◻

Corollary 4.25. *Let S be a connected M-slice of interval graph G with respect to good LBFS σ. The flyers of S with respect to σ induce a clique in G.*

Proof. From Claim 4.24, there exists a vertex w after S in σ that is universal to F_S. Therefore, by the P_3 Rule, F_S induces a clique. ◻

Claim 4.26. *Let C be an M-slice of type 1 of an interval graph G with respect to some good LBFS σ. Let X_1 and X_2 be the sets of good vertices of C and let $\widetilde{X_1}$ (respectively $\widetilde{X_2}$) denote $X_1 \cup (N(X_1) \cap (C \setminus K_C))$ (respectively $X_2 \cup (N(X_2) \cap (C \setminus K_C))$). Then at most one of $\widetilde{X_1}$ and $\widetilde{X_2}$ can contain vertices that fly, with respect to σ, out of C.*

Proof. Suppose to the contrary that $x \in F_C \cap \widetilde{X_1}$ and $y \in F_C \cap \widetilde{X_2}$. Then, by Claim 4.24, x and y fly to a common vertex z and, therefore, $xy \in E$ by the P_3 Rule. Thus $x \in K_1$ and $y \in K_2$ (Theorem 4.20). But now there is an AT on $\{x', z, y'\}$ where x' and y' are arbitrary vertices in X_1 and X_2 respectively. (The paths are: z, x, x'; z, y, y'; and x', z', y' where z', guaranteed by Claim 3.7, is in the clique of attachment of C and $zz' \notin E$.) ◻

We now turn our attention to type 2 M-slices (i.e. M-slices C where $C \setminus K_C$ is disconnected).

Claim 4.27. *Let C be a connected M-slice of type 2 of an interval graph G with respect to some good LBFS σ. Let the connected components of $C \setminus K_C$ be $C_1, \ldots, C_k$ $(k \geq 2)$ with X_i being the good vertices of $C_i, (1 \leq i \leq k)$. Let $\widetilde{X_i}$ denote $X_i \cup (N(X_i) \cap C_i)$. Then at most one $\widetilde{X_i}$ set can contain vertices that fly with respect to σ out of C.*

Proof. This follows immediately from Corollary 4.25. ◻

We now describe the structure of a connected type 2 M-slice when the given interval graph is module free.

COROLLARY 4.28. *Let G be a module free interval graph with good LBFS σ. For any connected type 2 M-slice C in σ, $C \backslash K_C$ has exactly two connected components: one contains a C-flyer with respect to σ and the other consists of a single vertex that does not fly out of C.*

Proof. If no component $C_1, C_2, \ldots, C_k$ of $C \setminus K_C$ flies, then any union of components forms a module in G; thus by Claim 4.27, exactly one component contains vertices that fly. G being module free requires $k = 2$ and the non-flying component to consist of a single vertex. ◻

Note that Claim 4.27 and Corollary 4.28 require the type 2 M-slice C to be connected. We now indicate what happens when $K_C = \emptyset$.

CLAIM 4.29. *Let C be a disconnected M-slice of type 2 of an interval graph G with respect to some good LBFS σ. Let the connected components of C be $C_1, \ldots, C_k$ ($k \geq 2$) with X_i being the good vertices of C_i ($1 \leq i \leq k$). Let $\widetilde{X_i}$ denote $X_i \cup (N(X_i) \cap C_i)$. Then at most two $\widetilde{X_i}$ sets can contain vertices that fly with respect to σ out of C. If $x \in C_i$ flies to $\tilde{x}$ and $y \in C_j$ ($j \neq i$) flies to $\tilde{y}$ where $\tilde{x} <_\sigma \tilde{y}$, then T, the connected component of $\Gamma^\sigma_{C,\tilde{x}}$ that contains C and $\tilde{x}$ is strictly contained in T', the connected component of $\Gamma^\sigma_{C,\tilde{y}}$ that contains C and $\tilde{y}$.*

Proof. First we note that, by the P_3 Rule, $x\tilde{y}, y\tilde{x}, \tilde{x}\tilde{y} \notin E$. We prove the second part first. Suppose that two components C_i and C_j, $j \neq i$ contain flyers x and y, as stated above, but $T = T'$. We now claim that t, the first vertex of T is inadmissible contradicting the fact that σ is good. To see this, let x', y' be vertices in the clique of attachment of $C \cap T$ that miss $\tilde{x}, \tilde{y}$ respectively; possibly $x' = y'$. The paths are: $t, P, x', y, \tilde{y}$ and $t, P', y', x, \tilde{x}$ where P is an arbitrary prior path from t to x' and P' is an arbitrary prior path from t to y'. Thus $T \subset T'$.

Now suppose that three different connected components, C_h, C_i, C_j contain vertices x, y, z respectively, that fly to $\tilde{x}, \tilde{y}, \tilde{z}$, respectively where $\tilde{x} <_\sigma \tilde{y} <_\sigma \tilde{z}$. By the result above, we know that $T \subset T' \subset T''$, where T'' is the connected component of $\Gamma^\sigma_{C,\tilde{z}}$ that contains C and $\tilde{z}$. Now let x' be a vertex in the clique of attachment of $C \cap T$ that misses $\tilde{x}$ and let y' be a vertex in the clique of attachment of $T \cap T'$ that misses $\tilde{y}$. By examining the paths induced on $\tilde{y}, y, x', z, \tilde{z}$ and $\tilde{x}, y', z, \tilde{z}$, we see that $\tilde{z}$ is not admissible in the graph induced on $\tilde{z}$ and the vertices before it in σ, contradicting Theorem 3.23. ◻

Note that there is a similar Corollary to Claim 4.29 as Corollary 4.28 is to Claim 4.27.

We now examine the behaviour of LBFS with respect to minimal separators of interval graph M-slices. We let S be an M-slice of interval graph G with respect to LBFS σ. Let A be a minimal separator of S (since G is chordal, A is a clique by Dirac's Theorem [12]) with $C_1 \cup C_2 \cup \cdots \cup C_k, (k \geq 2)$ the connected components of $S \setminus A$. We let $\widetilde{C_i}$ denote the vertices in C_i that are universal to A. (For chordal graphs, it is well known (for example, see Exercise 12, page 101 of [15]) that $\widetilde{C_i} \neq \emptyset$). In fact we can say something stronger.

CLAIM 4.30. *Let A be a minimal separator of chordal graph G with $\widetilde{C_i}$ the vertices in C_i (a connected component of $G \backslash A$) that are universal to A. Then $\widetilde{C_i}$ is connected.*

Proof. Suppose not and let c_1 and c_2 be in different connected components of $\widetilde{C_i}$. Since C_i is connected, consider an induced path $P : c_1 = p_0, p_1, \cdots, p_h = c_2)$ in C_i. Without loss of generality assume that $p_1 \notin \widetilde{C_i}$ and suppose p_1 is not adjacent to $a' \in A$. Let p_j, $2 \leq j \leq h$, be the first vertex on P after p_1 that is adjacent to a'.

Now the subgraph induced on $a', c_1, p_1, \cdots, p_j$ forms a cycle of size > 3, contradicting G being chordal. □

LEMMA 4.31. *Let S be a connected M-slice of interval graph G with respect to good LBFS σ where S has a minimal separator A and connected components $C_1, C_2, \cdots, C_k$ $(k \geq 2)$ of $S \setminus A$. In any LBFS τ of G with the first vertex of $S \setminus A$ in C_i, all vertices of A appear before any vertices of $C_j, j \neq i$, and every $\widetilde{C_j}, j \neq i$ appears as an M-slice in τ.*

Proof. By the LBFS Theorem, τ_S is an LBFS of S. Since A is a minimal separator of S, no vertex of $C_j, j \neq i$, can appear until at least one vertex of A has been visited. Let x be the first vertex of τ_S. Let a be the last vertex of A, u be the first vertex of C_i and v be the first vertex of $C_j, j \neq i$, visited by τ.

We now show that no vertex of $C_j, j \neq i$, can be before any vertex of $A \cup cl_\sigma(S)$ in τ. Suppose $v <_{\tau_S} a$. By the choice of u, $u <_{\tau_S} v$. Consider an arbitrary induced u, a-path P in $C_i \cup \{a\}$, and an arbitrary induced v, a-path Q in $C_j \cup \{a\}$. By Corollary 3.15, u', v', the neighbours of a on P, Q respectively must be before a. But now the P_3 Rule is violated since $u' \in C_i, v' \in C_j, j \neq i$. If v appears before a vertex of $cl_\sigma(S)$ in τ then the P_3 Rule is violated in τ.

By Corollary 4.25, at most one connected component of $S \backslash A$ has an S-flyer with respect to σ. If no $C_j, j \neq i$, has an S-flyer then, since all of $A \cup cl_\sigma(S)$ appears before C_j, the slice starting at v consists of exactly the $\widetilde{C_j}$'s, $j \neq i$, or contains all of the $\widetilde{C_j}$'s, $j \neq i$, as modules, and thus every $\widetilde{C_j}, j \neq i$ appears as an M-slice, as required.

Otherwise, some vertex of $C_j, j \neq i$, has an S-flyer with respect to σ. But then C_i has no such vertex. Consider the moment just before u, the first vertex of C_i, is visited in τ. Since C_i cannot be pulled, the only already visited neighbours of u at this point are in $A' \cup cl_\sigma(S)$, where A' is the set of already visited vertices of A. (It is possible that $A' = \emptyset$ or the subset of $cl_\sigma(S)$ before $u = \emptyset$.) Vertex u must be universal to A' because there is such a vertex in C_i, for example the vertices of $\widetilde{C_i}$, and therefore LBFS τ must choose such a vertex. At this moment, the slice T of τ consists of vertices universal to both A' and the already visited $cl_\sigma(S)$ vertices. Slice T includes $A \backslash A'$ (A is a clique, since G is a chordal graph), the unvisited vertices of $cl_\sigma(S)$, all of the $\widetilde{C_j}$'s, $j \neq i$, possibly some other vertices of $S \backslash A$, and possibly some other vertices not in S. The subslice of T starting immediately after u contains all unvisited vertices of $cl_\sigma(S)$, no vertices that appear after S in σ, and no vertices of $C_j, j \neq i$. Thus, in slice T, all vertices of $cl_\sigma(S)$ are visited before any vertices not in S except possibly vertices that are universal to S. This, in turn, causes all vertices of $S \cap T$ to be visited before any vertex of $T \backslash S$ that is not universal to S. Since, in addition, all of $A \cup cl_\sigma(S)$ is visited before v, the subslice of T starting at v consists of exactly the $\widetilde{C_j}$'s, $j \neq i$, or contains all of the $\widetilde{C_j}$'s, $j \neq i$, as modules, and thus, each $\widetilde{C_j}$ appears as an M-slice, as required. □

The next claim shows conditions where an M-slice in one good LBFS is also an M-slice in another good LBFS.

CLAIM 4.32. *Let C be a connected nonclique M-slice of an interval graph $G = (V, E)$ with respect to some good LBFS σ. Let $C_1, C_2, \ldots, C_k$, $k \geq 1$, be the connected components of $C \setminus K_C$. If C is of type 1 ($k = 1$), let X_1, X_2, K_1, K_2, $\widetilde{X_1}$, $\widetilde{X_2}$ be as defined in Theorem 4.20 and Claim 4.26 and, if C is of type 2 ($k \geq 2$), let $X_1, X_2, \ldots, X_k$, $\widetilde{X_1}, \widetilde{X_2}, \ldots, \widetilde{X_k}$ be as defined in Claim 4.27. Suppose that some vertex in $\widetilde{X_1}$ flies out of C in σ. Then for any good LBFS τ of G where the last C vertex is in X_1, C must appear as an M-slice.*

Proof. Note that in any LBFS, C and $C \setminus K_C$ are both M-slices or neither is. Assume for contradiction that C and $C \setminus K_C$ are not M-slices in good LBFS τ. By Claims 4.26 and 4.27, no vertex in X_i, $i \neq 1$, flies or neighbour flies out of C, unless it neighbour flies via a vertex of K_C. By Lemma 3.21, if $\tau_{C \setminus K_C}$ is not a good LBFS of $C \setminus K_C$, then the first vertex of $sc_\tau(C \setminus K_C)$ must be pulled and thus is in $\widetilde{X_1}$, contradicting Claim 4.3 or Claim 4.21. Thus v, the first vertex of $sc_\tau(C \setminus K_C)$ is good and must be in X_i, $i \neq 1$. We let S be the smallest M-slice of τ that starts at v and contains all of $C \setminus K_C$. Since $C \setminus K_C$ is not an M-slice of S, there are some vertices in $S \setminus C$ adjacent to some, but not all, vertices of $C \setminus K_C$ and these vertices must occur after C in σ. We let y (adjacent to x in $C \setminus K_C$) be such a vertex and denote by Y the subset of $cl_\sigma(C)$ that is not adjacent to y (by Claim 3.7, $Y \neq \emptyset$ since in σ, $y \notin C$). In τ, no vertex in Y can be before v, since otherwise $y \notin S$. No vertex of Y can be after S in τ; otherwise the P_3 Rule is violated since C is not a clique. Thus in τ, all vertices in Y are in $N(v) \cap S$.

Now consider σ and let C' be the smallest connected M-slice that strictly contains C and some vertices before C (in σ) and denote $(cl_\sigma(C) \cup K_C) \cap C'$ by A. In C', A is a minimal separator between C and C'_1, the connected component of $C' \setminus A$ that precedes A ($C'_1 \neq \emptyset$ since C' cannot start with a vertex that is universal to C since such a vertex is not simplicial). By Lemma 4.31, in τ, if any vertex of C'_1 precedes v then all of $cl_\sigma(C)$, including Y, is before v, a contradiction. Thus $\tilde{c}$, an arbitrary vertex of $\widetilde{C'_1}$ (vertices of C'_1 that are universal to A) must be after v in τ, and $\tilde{c}$ must be in S since it is universal to $cl_\sigma(C)$. But now, looking at S, the paths $v, y', \tilde{c}$ and v, P, x, y (where y' is an arbitrary vertex in Y and P is an arbitrary path between v and x in C) show that v is not admissible in S, contradicting τ being a good LBFS. (Note: $\tilde{c}v, \tilde{c}x \notin E$ since A separates $\tilde{c}$ from C in C'; $yy' \notin E$ by the definition of Y; $\underset{\sim}{y\tilde{c}} \notin E$ since otherwise the P_3 Rule is violated in σ; and $yv \notin E$ since only vertices in $\widetilde{X_1}$ can fly to y and $v \notin \widetilde{X_1}$.) $\square$

We now consider two good LBFSs operating on a type 1 interval graph, both starting from good vertices in the same clone set. Note that the notation follows that presented in Figure 6 and that we are dealing with "outermost" slices; recall that slice $S' \subset S$ is an *outermost* slice of S with respect to σ if there is no slice T of S such that $S' \subset T \subset S$. We define an *M-outermost slice* of S to be an outermost slice of S, a nontrivial module of an outermost slice of S, or a nontrivial module of S.

LEMMA 4.33. *Let S be a type 1 M-slice of an interval graph G with respect to a good LBFS and let X_1, X_2, K_1, K_2, X'_1, X'_2 be as defined in Theorem 4.20. Let u and v be arbitrary good vertices in X_1 and let σ and τ be arbitrary good LBFSs of S starting at u and v respectively. Then $S' \subseteq S \setminus (K_1 \cup K_S \cup X'_1)$ is an M-outermost slice of S with respect to σ if and only if S' is an M-outermost slice of S with respect to τ. Furthermore every connected component of X'_1 is an M-outermost slice with respect to both σ and τ.*

Proof. To simplify the argument we will assume that $K_S = \emptyset$. First we note that if a subgraph H of S appears as an M-outermost slice in both σ and τ then, by the definition of a module, all modules of H are M-outermost slices in both σ and τ. Let $C_1, \cdots, C_k$ be the connected components of X'_1 where $u \in C_1$ and $v \in C_i, 1 \leq i \leq k$. Since X'_1 is a module of S, each $C_i, 1 \leq i \leq k$ is a module of S and thus an M-outermost slice of S with respect to both σ and τ.

We let Z denote the vertices in $S \setminus (X'_1 \cup K_1)$ that are universal to K_1; Z is guaranteed by Theorem 4.20 and Lemma 4.31. Any LBFS starting at u will next visit $N_{C_1}(u) \cup K_1$ followed by the rest of C_1 followed by an outermost slice T containing the components $C_2, \cdots, C_k, Z$ in arbitrary order. Since Z is a module of T, it is an M-outermost slice of S with respect to σ. A similar argument shows that Z is also an M-outermost slice of S with respect to τ. The slice (possibly a trivial slice consisting of a single vertex) immediately following Z in both σ and τ is determined strictly by its neighbourhood in $Z \cup K_1$ and is not affected by the ordering of $\sigma_{Z \cup K_1}$ and $\tau_{Z \cup K_1}$. This is because, for all $y, z \in S \backslash (X_1' \cup K_1 \cup Z)$, $N(y) \cap (K_1 \cup Z) \subseteq N(z) \cap (K_1 \cup Z)$ or $N(y) \cap (K_1 \cup Z) \supseteq N(z) \cap (K_1 \cup Z)$; otherwise, we either violate the P_3 Rule or find vertices unrelated with respect to u, contradicting the admissibility of u. Again the orderings of this slice do not affect the determination of the next slice and the process continues until K_2 is encountered. The last outermost slice (again possibly trivial) is X_2', which is a module of S. □

PART 2
THE MULTI-SWEEP LBFS
INTERVAL GRAPH RECOGNITION ALGORITHM

5. The interval graph recognition algorithm.

Our initial algorithm reported in [10] uses four LBFS sweeps and, as shown in §5.2, this algorithm is flawed since a particular early sweep must be good (i.e. the result of an LBFS+). In the following, a six sweep algorithm is presented and proved correct. We strongly believe that the five sweep version is also correct; however, the proof of that algorithm is more complicated than the proof of correctness of the six sweep algorithm. Note that the difference amongst these various algorithms is the number of LBFS+ sweeps that follow an arbitrary preprocessing LBFS. After the last LBFS+ sweep, the algorithms employ a new LBFS, called LBFS* that chooses the first vertex of a slice based on the *two* previous LBFS+ sweeps. Before specifying the full algorithm we describe LBFS*. Note that the correctness of the six sweep algorithm is presented in §6 together with suggestions for the proof of correctness of the five sweep algorithm. The linear time implementation details are presented in §7.

5.1. LBFS*.

Given an M-slice S recall that vertex x in S *flies (F) with respect to S* if there is a vertex y such that y occurs after S and $xy \in E$. Further x is said to *neighbour fly (NF) with respect to S* if it does not fly but it has a neighbour in S that flies. Finally, vertex x is said to be *OK with respect to S* if it neither flies nor neighbour flies with respect to S. As before, reference to S will be omitted when the context makes it obvious.

Procedure LBFS* (G, τ_1, τ_2):
This variant of LBFS needs *two* previous LBFS sweeps, τ_1 and τ_2. Given a slice S (i.e. as identified in step $(\star)$ of LBFS), we select two vertices α and β where α is the last S-vertex in τ_1 and β is the last S-vertex in τ_2. LBFS* chooses between α and β by referring to the following decision table.

		β		
		F	NF	OK
	F	β	β	β
α	NF	α	β	β
	OK	α	α	β

Table 1: Tie-breaking rules for LBFS*

For example, if α neighbour flies and β flies, then α is chosen. When α is chosen we say that the *α-Rule has been followed* (similarly for the *β-Rule*). Implementation details of this algorithm are provided in §7.

5.2. Overview of the algorithm. The interval graph recognition algorithm is as follows:

1. Do an arbitrary LBFS π'.
2. LBFS+ (G, π') yielding sweep π.
3. LBFS+ (G, π) yielding sweep σ.
4. LBFS+ (G, σ) yielding sweep σ^+.
5. LBFS+ (G, σ^+) yielding sweep σ^{++}.
6. LBFS* $(G, \sigma^+, \sigma^{++})$ yielding sweep σ^*.
7. If σ^* is an I-ordering then conclude that G is an interval graph; else, conclude that G is not an interval graph.

The algorithm does six LBFS sweeps. The four sweep version of this algorithm does an arbitrary LBFS followed by two LBFS+ sweeps followed by LBFS* (in particular steps 2 and 3 are omitted and the LBFS+ in step 4 is with respect to π' instead of σ). As mentioned previously, the four sweep algorithm is flawed since it allows the sweep preceding σ^+ to be arbitrary; however, as shown below, this sweep must be good (as guaranteed by LBFS+).

To illustrate the various features of the algorithm, we assume that the five sweep algorithm is run on the graph of Figure 7. Thus we have an initial sweep π' that is followed by three LBFS+ sweeps, followed by an LBFS* sweep (in particular step 2 is omitted and the LBFS+ in step 3 is with respect to π' rather than π). For the given π', the algorithm will produce the following sweeps $\sigma, \sigma^+, \sigma^{++}$ and σ^*. Square brackets indicate the slice structure of each sweep.

π': [20 [8 [2 4] 21] [16 [15 [9 12]] 13 [11 14] 17 10 18 7 19 6] 5 3 1 22]
σ: [22 4 [3 2 [5 6 [7 8] 9 [18 17 12 [14 13 [11 15] 16] 10] 19 20] 21] 1]
σ^+: [1 2 [20 [8 4] [19 18 [17 9] 12 [16 15 13 [11 14]] 10 7 6] 5 3] 21 22]
σ^{++}: [22 4 [21 20 [8 2] [6 7 9 [10 11 [13 12] [14 15] 16 17 18] 19] 5 3] 1]
σ^*: [1 2 [3 4 [5 6 [7 8] 9 [10 11 [12 13] [14 15] 16 17 18] 19 20]] 21 22]

There are a few interesting facts illustrated by this example. First it shows many of the cases of tie-breaking illustrated in Table 1. To illustrate the LBFS* algorithm, we will examine each slice in turn and show how the algorithm made its choice in

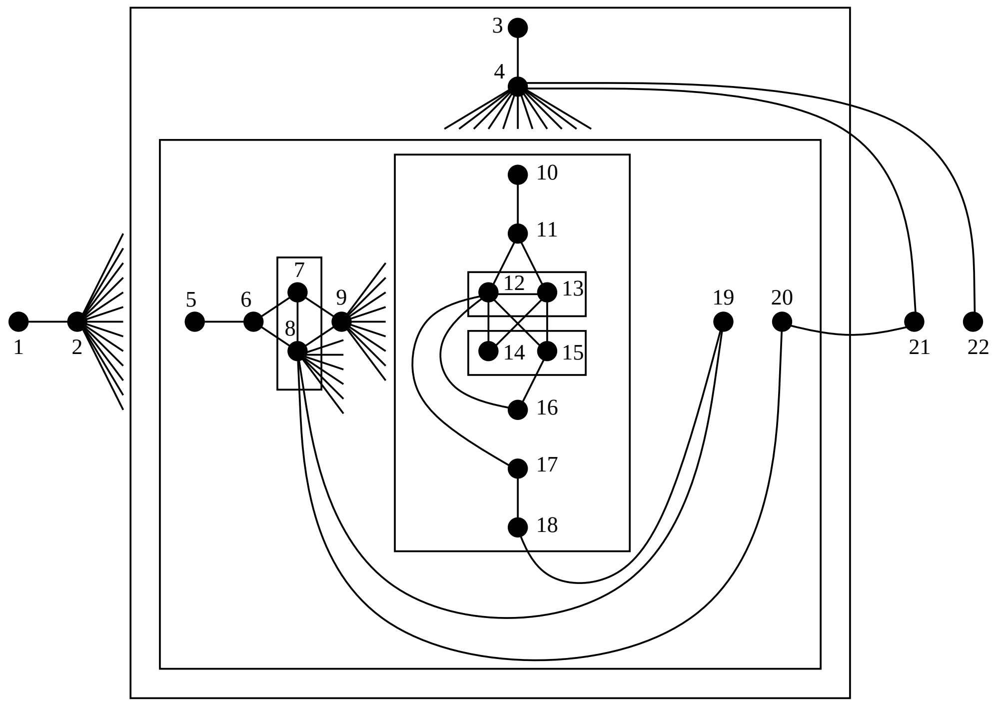

FIG. 7. *Illustrating the interval graph recognition algorithm: the slice structure of σ^**

producing σ^*. The first slice is V itself. Here $\alpha = 22$ and $\beta = 1$. Since there are no vertices outside this slice, both α and β are OK and thus β is chosen. The next slice includes vertices 3 to 20; here $\alpha = \beta = 3$; both α and β neighbour fly (because of the edge from 4 to 21) and β is chosen. The next slice includes vertices 5 to 20; now $\alpha = \beta = 5$ and both α and β are OK so β is chosen. The clique $\{7, 8\}$ is next and $\alpha = \beta = 7$ (both α and β fly because of the edge from 7 to 9 and β is chosen). Now $S = \{10, 11, 12, 13, 14, 15, 16, 17, 18\}$ is encountered; $\alpha = 10$ and $\beta = 18$. Since β flies and α is OK, α is chosen. $\{12, 13\}$ is next; $\alpha = 13$ and $\beta = 12$. Both fly and β is chosen. Finally, when $\{14, 15\}$ is encountered, $\alpha = 14$ and $\beta = 15$; here, β flies and α is OK and α is chosen. In both cases where α is chosen, if an ordinary LBFS+ were used, the β vertex would have been chosen resulting in umbrellas. Thus, in addition to the examples produced by Ma [26], this example shows how the Simon approach [31] fails even when extended to five sweeps. (Note that in σ^{++} there are a number of umbrellas, adding to Ma's list of counter-examples to Simon's four sweep algorithm.)

Secondly, this example shows the necessity of the preprocessing sweep that was omitted from our initial four sweep algorithm reported in [10]. In our example, assume that the four sweeps are π', σ, σ^+, and an LBFS* done on (σ, σ^+); this produces the ordering: [22 4 [21 20 [8 2] [19 18 [9 17] 12 [14 13 [11 15] 16] 10 7 6] 5 3] 1] which has the umbrellas 11 15 10 and 11 16 10.

The next claim shows that, in the six sweep algorithm, the sweeps σ^+, σ^{++} and σ^* "flip" the first and last vertices of the previous sweep.

CLAIM 5.1. *In the six sweep algorithm, suppose that σ starts at vertex y and ends at vertex z. Then σ^+ starts at z and ends at y, σ^{++} starts at y and ends at z, and σ^* starts at z and ends at y.*

Proof. The claims for σ^+ and σ^{++} follow from Theorem 4.6 and Corollary 3.19. Turning our attention to σ^* we immediately see that it starts at z since $V(G)$ has no flyers and thus the β Rule is invoked. To show that σ^* ends at y we follow the proof of Theorem 4.6 and note that v must be a β vertex since $\Gamma^*_{y,w}$ has no vertices after it in σ^* and thus no flyers. □

5.3. Overview of the proof of correctness. The first step (§6.1) in the proof assumes that the algorithm fails to produce an umbrella free ordering σ^* for some interval graph G with initial LBFS π'. A specific umbrella uw over v is identified and the structure of $S = \Gamma^{\sigma^*}_{u,v}$ is examined in σ^*, σ^{++}, σ^+ and σ. In §6.2, the set W is defined to be the set of vertices that occur after S (in σ^*) and are reachable by a monotone path from vertices in $F_S \setminus K_S$ (note that $w \in W$) where F_S is the set of flyers from S and K_S is the universal clique of S (recall that it is possible that $K_S = \emptyset$). It is shown that in σ^{++}, σ^+ and σ, all vertices of S must occur before the vertices of W. The proof concludes in §6.3 where the structure of G outside $S \cup W$ is examined. By considering an M-slice in a fictitious LBFS of G, we show that its behaviour in at least one of σ, σ^+ or σ^{++} must require W to precede S.

6. Correctness of the algorithm. Before studying the correctness of the algorithm itself we present some notation. We will use Γ^+, Γ^{++} and Γ^* to denote Γ^{σ^+}, $\Gamma^{\sigma^{++}}$ and Γ^{σ^*} respectively. Similarly $<_+$, $<_{++}$, $<_*$, cl_+, cl_{++} and cl_* will denote $<_{\sigma^+}$, $<_{\sigma^{++}}$, $<_{\sigma^*}$, cl_{σ^+}, $cl_{\sigma^{++}}$ and cl_{σ^*}.

6.1. Choice of the umbrella and definition of S. We now examine the algorithm itself. First, we assume that G is an interval graph, but for initial LBFS π', the resulting σ^* is not umbrella free. We also assume, without loss of generality, that $diam(G) \geq 3$. Clearly we may assume that any counterexample is connected; if $diam(G) = 2$, then there is a universal clique whose removal does not affect the orderings of the algorithm. We now choose a particular umbrella in σ^*.

Let uw ($u <_* w$) be an umbrella of G flying over a vertex v, such that $\Gamma^*_{u,w}$ is as outermost as possible (i.e. no slice strictly containing $\Gamma^*_{u,w}$ has an umbrella that is not contained in a subslice) for all umbrella edges in G with respect to σ^*. Note that by the P_3 Rule $wv \notin E$. Amongst the tied umbrellas choose one such that v is as rightmost as possible. We also assume without loss of generality that G has no nontrivial modules. To justify this assumption, note that a problem could occur only if exactly two of $\{u, v, w\}$ were in a module M. By the definition of module, the pairs $\{u, v\}$ and $\{v, w\}$ are ruled out; the exclusion of pair $\{u, w\}$ follows from Lemma 3.3.

Let $S = \Gamma^*_{u,v}$ with first vertex s. It is possible, for now, that $u = s$; later (Claim 6.8), we will show that this cannot happen. Note that, by the choice of v, u is adjacent to all vertices between v and w (in σ^*). Since there is a u, w-path missing v, Claim 4.9 ensures $w \notin S$ and thus $S \subset \Gamma^*_{u,w}$. Clearly, v is the last vertex of S. (If there is a vertex v' after v in S, with $uv' \in E$, then by the P_3 Rule, $vv' \notin E$. Now there is a u, v'-path missing v which by Claim 4.9 implies $v' \notin \Gamma^*_{u,v} = S$. If $uv' \notin E$, then, as before, we have contradicted the choice of v.)

We now examine various forms that S may have and define sets L, R, $\widetilde{L}$ and $\widetilde{R}$. Clearly $diam(S) \neq 1$, since otherwise S is a clique and $uv \in E$. If S is of type 1 (i.e. $diam(S \setminus K_S) \geq 3$), then we define L to be the clones of s in S and R to be the clones of v in S. $\widetilde{L}$ and $\widetilde{R}$ denote $L \cup (N(L) \cap S \setminus K_S)$ and $R \cup (N(R) \cap S \setminus K_S)$ respectively. Recall the definitions of $X_1, X_2, K_1, K_2, \widetilde{X_1}$ and $\widetilde{X_2}$ from Theorem 4.20 and Claim 4.26. If $s \in X_1$ then $L = X_1$, $R = X_2$, $\widetilde{L} = \widetilde{X_1}$ and $\widetilde{R} = \widetilde{X_2}$. Now, by Theorem 4.20, K_1 and K_2 are minimal separating cliques universal to L and R respectively and, thus, $K_1 \subseteq \widetilde{L} \backslash L$ and $K_2 \subseteq \widetilde{R} \backslash R$.

Suppose S is of type 2 (i.e. $S \setminus K_S$ is disconnected) with connected components $C_1, C_2, \ldots, C_k$ ($k \geq 2$) of $S \setminus K_S$, as ordered by σ^*. If S is connected then, by the definition of $\Gamma^*_{u,v}$, Theorem 3.10, Claim 4.3, and Corollary 4.28, $k = 2$, $\{u, s\} \subseteq C_1$, $C_2 = \{v\}$ and v does not fly. Now we define L to be the good vertices in C_1 and R to be $\{v\}$. Here $\widetilde{L}$ denotes $L \cup (N(L) \cap C_1)$ and $\widetilde{R} = \{v\}$. We now examine the case where S is disconnected; as before, $s, u \in C_1$ and $v \in C_k$. By Claim 4.29, at most one other component C_j, $j \neq 1$, of S may contain flyers, where C_j contains vertex x that flies to $\tilde{x}$; then by Claim 4.29 and the choice of u, v, w, $\tilde{x} <_* w$. But now we have contradicted the choice of v since $v <_* \tilde{x}$ and $u\tilde{x}, w\tilde{x} \notin E$ by the P_3 Rule. Thus we have exactly the same situation as when S is connected, namely, $k = 2$, $\{u, s\} \subset C_1$, $C_2 = \{v\}$ and v does not fly. The L and R sets are the same as above.

We now study the structure of S and its projections in σ^+ and σ^{++}. Later, we will examine its projections in σ. First, some results on vertex u.

CLAIM 6.1. *$u \in \widetilde{L}$.*

Proof. For each type of S, this is a straightforward observation. First, by Claim 4.23, u is either good or adjacent to a good vertex. If S is of type 1, then if $u \in \widetilde{R} \backslash K_2$ then $\Gamma^*_{u,v} \subset S$, contradicting the definition of S, and if $u \in K_2$, then $uv \in E$. If S is of type 2 then, as noted above, $u \in C_1$ and thus u must be in $\widetilde{L}$. □

If $diam(S) = 2$, some flying vertices of S may be in K_S, the universal clique. Of course, such edges cannot form an umbrella and, thus, have no effect on the algorithm. We let F_u denote the set of flying vertices in $S \setminus K_S$. Since u is not in K_S, F_u is not empty.

CLAIM 6.2. *F_u induces a clique in G.*

Proof. If S is connected then this follows immediately from Corollary 4.25. Otherwise, F_u is a subset of C_1 and the result follows by applying Corollary 4.25 to C_1. □

We now examine which vertices in S may fly out of S.

CLAIM 6.3. *No vertex outside $\widetilde{L}$ may fly with respect to S, except for a vertex in K_S, the universal clique, if $diam(S) = 2$.*

Proof. This follows immediately from Claims 4.26, 4.27, 6.1, and 6.2. □

COROLLARY 6.4. *No vertex in R may fly with respect to S; a vertex in R may neighbour fly only if $diam(S) = 2$ and a vertex in the universal clique K_S flies with respect to S.*

CLAIM 6.5. *In σ, σ^+ and σ^{++} either all vertices of L are before all vertices of R or vice versa.*

Proof. This follows from Claim 4.22. □

CLAIM 6.6. *In σ^{++}, $sc_{++}(S)$ must have all the vertices in R occur before the vertices in L.*

Proof. This follows immediately from Claim 6.3, Corollary 6.4, Claim 6.5 and the algorithm. In particular, if $sc_{++}(S)$ ended with an R vertex, then it would have to be chosen, contradicting the fact that $s \in L$. □

CLAIM 6.7. *In σ^+, $sc_+(S)$ must have all the vertices in L occur before the vertices in R.*

Proof. Let a_1 be the first vertex of $sc_{++}(S)$. By Lemma 3.21, either a_1 is a bad vertex that has been pulled, or a_1 is a good vertex and σ_S^{++} is a good LBFS of S. In the latter case, $a_1 \in R$ by Claim 6.6. Now look at the set of vertices tied with a_1 as it is visited by σ^{++}. If this set does not include a_2, the last vertex of $sc_{++}(S)$, then there is a vertex $x \notin S$ such that x is adjacent to $a_1 \in R$ but not adjacent to $a_2 \in L$. This contradicts Corollary 6.4. Thus, the set of tied vertices contains a_2. Since a_1 was chosen by the algorithm, a_1 is the last vertex of $sc_+(S)$. Thus, by Claim 6.5, in σ^+, the L vertices occur before the R vertices in σ^+.

Thus a_1 is a bad vertex and is pulled. Let a'_1 be the first vertex of $S \setminus K_S$ in σ^{++} (note it is possible that $a_1 \in K_S$ and is pulled). Again, consider the set of vertices tied with a'_1 as it is visited by σ^{++}. If this set does not include a_2, then a'_1 has been pulled and thus, by Claim 4.23, a'_1 is good or adjacent to a good vertex in $S \setminus K_S$. We claim that a'_1 is not good in $S \setminus K_S$. This is clearly true if $a_1 = a'_1$. Now, proceeding as above, we see that $a'_1 \in R$ has been pulled but $a_2 \in L$ has not, contradicting Corollary 6.4. Thus $a'_1 \in \widetilde{L} \setminus L$ or $a'_1 \in \widetilde{R} \setminus R$. If S is of type 2, $a'_1 \notin \widetilde{L} \setminus L$ by Corollary 3.18 and $a'_1 \notin \widetilde{R} \setminus R$ since $\widetilde{R} \setminus R = \emptyset$. If S is of type 1, then $a'_1 \notin \widetilde{L} \setminus L$ by Claim 4.21 (note that σ_S^{++} would have a vertex in $\widetilde{L} \setminus L$, namely a'_1, followed by vertices in R, followed by vertices in L as guaranteed by Claim 6.6). Thus, $a'_1 \in \widetilde{R} \setminus R$ but this contradicts Corollary 6.4. Now the set of tied vertices includes a_2; thus since $a'_1 \in S$ it is the last vertex of $sc_+(S)$ and thus is good in S. This implies, by Claim 6.6, that $a'_1 \in R$ thereby completing the proof. ∎

Claim 6.8. *Vertex s is a β vertex and does not fly with respect to S; s neighbour flies only if $diam(S) = 2$ and a vertex in the universal clique K_S flies with respect to S.*

Proof. The fact that s is the β vertex follows immediately from Claims 6.6 and 6.7. Now examine the possible situations where β is chosen, according to the decision table for the LBFS* algorithm. By Table 1, if β flies, then α must fly, contradicting Corollary 6.4. If β neighbour flies, then α either flies (contradicting Corollary 6.4) or neighbour flies. Again, by Corollary 6.4, this implies that $diam(S) = 2$ and a universal vertex flies. ∎

Corollary 6.9. *Vertex s is adjacent to no vertices in F_u.*

Proof. The result follows from Claim 6.8 and the definitions of s and F_u. ∎

6.2. The definition of W and its placement in σ^*, σ^{++}, σ^+ and σ. Now define W, with last vertex $\tilde{w}$, to be the set of vertices that occur after S in σ^* and are reachable by a monotone path from F_u. The rest of this subsection discusses basic properties of W in σ^*, σ^{++}, σ^+ and σ. In particular, in all three sweeps, all vertices of S occur before the vertices of W.

Lemma 6.10. *In σ^*, the vertices in W are consecutive and occur immediately after S.*

Proof. Suppose they are not consecutive and consider the path P in $F_u \cup W$: $u, \cdots, w_1, w_2$ such that $w_1 <_* p <_* w_2$ where $p \notin W$. (Note that by Claim 6.2 the neighbour of u in P possibly is a vertex in F_u.) Such a path exists by the existence of $\tilde{w}$. By the choice of v, u is adjacent to all vertices after v up to and including w and thus $w <_* p$. From the definition of slices it follows that $\Gamma^*_{u,w} \subseteq \Gamma^*_{u,p} \subseteq \Gamma^*_{u,w_2}$. Since p misses P, by Claim 4.9, $w_2 \notin \Gamma^*_{u,p}$ and thus $\Gamma^*_{u,w} \subset \Gamma^*_{u,w_2}$. Since $p \notin W$, p is not adjacent to any vertex of P up to and including w_1. Also, $pw_2 \notin E$ or else we have the $P_3 : w_1 w_2 p$. The existence of P implies that $u \in \Gamma^*_{w_1,w_2}$ and thus $\Gamma^*_{w_1,w_2} = \Gamma^*_{u,w_2} \supset \Gamma^*_{u,w}$. But now, the umbrella w_1, w_2 over p contradicts our choice of u, w (i.e. $\Gamma^*_{u,w}$ is not outermost). ∎

The structure of S and W in sweep σ^* is depicted in Figure 8.

Lemma 6.11. *In σ^{++}, all vertices of W occur after $sc_{++}(S)$.*

Proof. Let x' be an arbitrary vertex in W that is adjacent to some vertex x in F_u; by Claim 6.3, $x \in \widetilde{L}$.

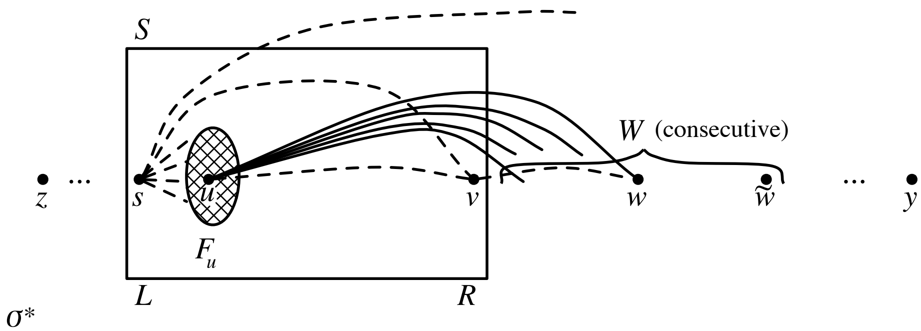

FIG. 8. *The structure of S and W in σ^**

First, we claim that x' must occur after s in σ^{++}. Recall that s is the last vertex of $sc_{++}(S)$. We assume that $x' <_{++} s$ and will show that x' and v are unrelated with respect to s, contradicting Theorem 3.23. If S is of type 1, we assume without loss of generality that $s \in X_1$ and $v \in X_2$. Thus, $L = X_1$, $\widetilde{L} \subseteq X_1' \cup K_1$, and $R = X_2$. Now we note that $F_u \cap K_1 = \emptyset$ by Claim 6.8 and thus $x \notin K_1$. Now the path s, K_1, x, x' and an s,v-path through K_1 and K_2, show that x' and v are unrelated with respect to s. If S is of type 2, let $\tilde{x}$ be an arbitrary vertex in $cl_*(S)$ such that $\tilde{x}x' \notin E$ (such a vertex exists by Claim 3.7). Note that, by the P_3 Rule, in σ^{++}, $\tilde{x}$ is before the first nonadjacent pair of S vertices and thus is before s. Now consider the paths $s, \tilde{x}, v$ and an s,x-path in C_1 together with the edge xx'. Again, this shows that x' and v are unrelated with respect to s.

Having shown $s <_{++} x'$ we finish the proof by noting that all paths from F_u to vertices in W must be monotone. To see this, let P be a path from vertex $x \in F_u$ to a vertex $w' \in W$ where w' is the first vertex after x on P such that $w' <_{++} s$. Since x', the neighbour of x on P is after s in σ^{++}, we have contradicted Corollary 3.14. ▯

This structure in σ^{++} is illustrated in Figure 9.

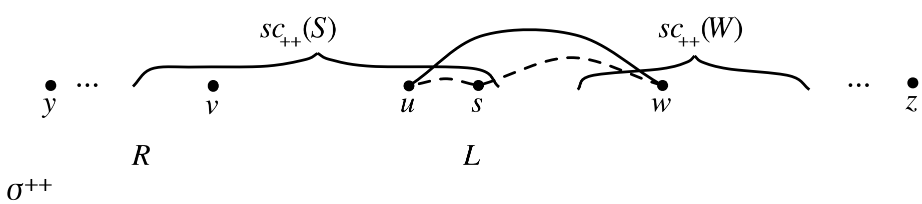

FIG. 9. *The structure of S and W in σ^{++}*

LEMMA 6.12. *In σ^+, all vertices of W occur after $sc_+(S)$.*

Proof. Consider σ^+ at the point that it visits s_1, the first vertex of $sc_+(S \setminus K_S)$. If the set of vertices tied with s_1 does not include s_2, the last vertex of $sc_+(S)$, then there is some vertex x before s_1 in σ^+ that is adjacent to s_1 but not adjacent to s_2 (Claim 3.7(i)). Thus x is not universal to S and therefore x is in W and $s_1 \in F_u$. Now, by Claim 6.8, since s does not fly, $s_1 \neq s$, and by Corollary 6.9, $s_1 s \notin E$.

Since s is the last S-vertex in σ_{++} (Claim 6.8), $s_1 <_{++} s$. Now look at $\Gamma^{++}_{s_1,s}$ with first vertex a. Clearly, $a \neq s_1$ since s_1 could never have been chosen as the first vertex (since s_1 is before s in σ^+). In σ^+ we have $s_1 <_+ s <_+ a$. By Theorem 3.10 applied to σ^{++}, there is a prior path P from s_1 to a that is missed by s. Now try to place a in σ^*. Because of P, a does not occur before S in σ^*; since W is after $sc(S)$ in σ^{++} (by Lemma 6.11) $a \notin W$. Thus $a \in S$.

Now look at $\Gamma^+_{s_1,s}$ and let b be the first vertex ($b \neq s_1$ since $x <_+ s_1$, $s_1x \in E$ and $sx \notin E$). By Claim 4.9, the path P forces $a \notin \Gamma^+_{s_1,s}$. Thus, there is a vertex $c \in cl_+(\Gamma^+_{s_1,s})$ such that $ca \notin E$. Trying to place c in σ^*, we easily see that $c \in S \setminus K_S$ ($c \notin cl_*(S) \cup K_S$ since $ca \notin E$; c is not after S since s does not fly with respect to S). But now, we have contradicted the fact that s_1 is the first vertex of $sc_+(S \setminus K_S)$.

The above argument shows that s_1, the first vertex of $sc_+(S \setminus K_S)$ cannot be pulled and thus is good. Therefore, X, the set of vertices tied with s_1 at the moment when σ^+ visits s_1, includes s_2, the last vertex of $sc_+(S)$.

We first claim that no vertex of W occurs before s_1 in σ^+. To see this, let w_0 be any vertex in W that occurs before s_1 and consider any shortest path P from w_0 to a vertex in F_u. On P there must be an edge w_1w_2 such that $w_1 <_+ s_1 <_+ w_2$. Since w_1 must be homogeneous to X, $S\backslash K_S \subseteq F_u$, contradicting Claim 6.3.

We now show that a vertex $w_1 \in W$ (without loss of generality, let w_1 be the leftmost W vertex in σ^+) can be inside $sc_+(S)$ only if it is in $N(s_1)$. To see this, let w_1' be a vertex of $cl_*(S)$ that is not adjacent to w_1. Since $w_1 \in X$, $w_1' \notin cl_+(X)$ and thus $s_1 <_+ w_1'$. No vertex of $S\backslash N(s_1)$ can occur before w_1' in σ^+ since this would violate the P_3 Rule. Once w_1' (which is universal to S) has been visited, all other vertices of S must occur in σ^+ before w_1. Otherwise, by Theorem 3.1, w_1 is adjacent to a neighbour a of s_1 that occurs before w_1' in σ^+ and is not adjacent to s_2, but then w_1 and s_2 are unrelated in X with respect to s_1, contradicting that s_1 is good.

Consider a good LBFS θ that is identical to σ^* until slice S is encountered, and then S is visited in θ in the order of a good LBFS of S that begins at v and ends at s. (Theorem 4.6 guarantees the existence of such an ordering of S.)

To complete the proof, suppose that $w_1 \in N(s_1)$ is in $sc_+(S)$. Then, again, since $s_1 \in F_u$, $s_1 \neq s$ and $s_1s \notin E$ by Claim 6.8 and Corollary 6.9. Now, let $A = \widetilde{\Gamma}^{\theta}_{s_1,s}$ and let A_1 be the clones of s_1 in A and A_2 the clones of s in A. By Claim 4.22, $A_1 \cap A_2 = \emptyset$. $s_1 \in A_1$ and thus from Claim 4.22 we conclude that $A_1 <_+ A_2$. Now look at $sc_{++}(A)$ with first vertex a. The last A vertex in σ^{++} is s; thus by Claim 4.22 $A_1 <_{++} A_2$. In σ^{++}, a is either pulled or good. By Lemma 6.11, all W vertices are after $sc_{++}(S)$ and thus there are no vertices before a in σ^{++} that are adjacent to a but not to s. Now, when σ^{++} reaches a, there is a nontrivial slice but we have reached a contradiction since $a \in A_1$ but $A_1 <_+ A_2$. □

See Figure 10 for a depiction of the structure in σ^+.

Now we turn our attention to σ and examine the relative positions of L and R as well as W and S.

Claim 6.13. *In $sc_\sigma(S)$, the vertices in R occur before the vertices in L.*

Proof. The proof is similar to that of Claim 6.7. Let a_1 be the first vertex of $sc_+(S)$. By Lemma 3.21, either a_1 is a bad vertex that has been pulled or a_1 is a good vertex and σ^+_S is a good LBFS of S. In the latter case, by Claim 6.7, $a_1 \in L$. Consider the set of vertices tied with a_1 as it is visited by σ^+. If this set does not include a_2, the last vertex of $sc_+(S)$, then there exists vertex $x \notin S$, such that x is adjacent to $a_1 \in L$ and not adjacent to $a_2 \in R$. Thus, $x \in W$ contradicting Lemma 6.12. Now, the set of tied vertices includes a_2. Since a_1 was chosen by the LBFS+

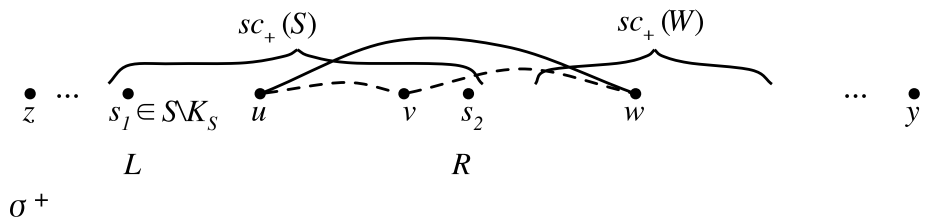

FIG. 10. *The structure of S and W in σ^+*

algorithm, a_1 is the last vertex of $sc_\sigma(S)$ and, thus, by Claim 6.5, the vertices in R occur before the vertices in L.

Thus, a_1 is a bad vertex and is pulled. Let a_1' be the first vertex of $S \setminus K_S$ in σ^+. Again we look at the set of vertices tied with a_1' as it is visited by σ^+. If this set does not include a_2, then a_1' has been pulled and thus following the argument in the proof of Claim 6.7, $a_1' \in \widetilde{L} \setminus L$ or $a_1' \in \widetilde{R} \setminus R$. $a_1' \notin \widetilde{R} \setminus R$ by Corollary 6.4; thus $a_1' \in \widetilde{L} \setminus L$ and the pulling vertex is in W contradicting Lemma 6.12. Thus, the set of tied vertices includes a_2 and as above, a_1' is good and is the last vertex of $sc_\sigma(S)$. Therefore, $a_1' \in L$ and the proof is complete. □

LEMMA 6.14. *In σ, all vertices of W occur after $sc_\sigma(S)$.*

Proof. Let s' be the last vertex of $sc_\sigma(S)$. By Claim 6.13, $s' \in L$, and thus $s' \neq v$ and $s'v \notin E$. Suppose xx' is a flying edge where $x \in F_u$, $x' \in W$ and $x' <_\sigma s'$. Note there is a vertex $\tilde{x} \in cl_*(S)$ such that $\tilde{x}x' \notin E$; $\tilde{x}$ must be before s' in σ, otherwise it would form a P_3 with any nonadjacent pair of vertices of S. If $x = s'$, $x' <_\sigma x$ and we have contradicted the P_3 Rule. Thus $x \neq s'$. Similarly, by the P_3 Rule, $s'x' \notin E$. If $s'x \in E$, then we immediately see that x' and v are unrelated with respect to s' contradicting Theorem 3.23 (consider the paths s', x, x' and $s', \tilde{x}, v$). Finally, we examine the case where s' is not adjacent to any neighbour of x' in F_u. (Having fixed the adjacencies of s' to x and x', the proof proceeds exactly as in Lemma 6.11 where s' now plays the role of s.) If S is of type 1 (where we assume $X_1 = L$), consider the paths s', K_1, x, x' and $s', K_1, \cdots, K_2, v$; if S is of type 2, consider the paths $s', \tilde{x}, v$ and s' followed by a path in C_1 to x and then x'. As with Lemma 6.11, the proof follows from the fact that the paths starting with the edge xx' must be monotone. □

See Figure 11.

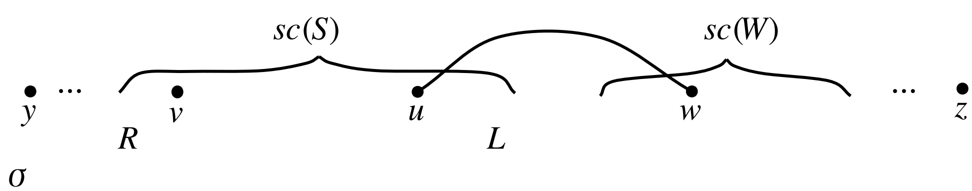

FIG. 11. *The structure of S and W in σ*

To summarize the results of this subsection, Lemmas 6.14, 6.12 and 6.11, combined, guarantee that in all three sweeps, σ, σ^+ and σ^{++}, all vertices of S precede all vertices of W.

6.3. Completion of the proof. To complete the proof of correctness of the algorithm, we will show that at least one of σ, σ^+, and σ^{++} must have violated Lemma 6.14, Lemma 6.12, or Lemma 6.11, respectively. We do this by examining a sequence of super slices of S (called the S-hierarchy), as seen in σ^*, and seeing the structure of such slices when they are visited from the "other direction", i.e. the opposite direction from σ^*. The motivation behind the structure results we now present is to understand the effects of the parts of G that precede S and follow W. Clearly these parts of the graph will play a role in the various good sweeps of the algorithm. We now define the *S-hierarchy*, a nesting of slices between S and G, with respect to σ^*. S_0 is defined to be S, and for $i > 0$, S_i is the smallest connected slice strictly containing S_{i-1}. Let $G = S_g$ and let s_i denote the first vertex (with respect to σ^*) of S_i (note $s_0 = s$). Since G is module free, Corollary 4.28 implies that if S_i, $i > 0$, is of type 2 then it consists of a (possibly empty) universal clique, one component that flies and one component of a single vertex that does not fly. Such a single vertex has no impact on subsequent arguments completing the correctness proof of the algorithm.

To illustrate these and subsequent definitions, we consider the graph in Figure 12, together with the sweeps that illustrated the counter-example to the four sweep algorithm; in particular, consider the sweeps:

π': [20 [8 [2 4] 21] [16 [15 [9 12]] 13 [11 14] 17 10 18 7 19 6] 5 3 1 22]
σ: [22 4 [3 2 [5 6 [7 8] 9 [18 17 12 [14 13 [11 15] 16] 10] 19 20] 21] 1]
σ^+: [1 2 [20 [8 4] [19 18 [17 9] 12 [16 15 13 [11 14]] 10 7 6] 5 3] 21 22]
σ^*: [22 4 [21 20 [8 2] [19 18 [9 17] 12 [14 13 [11 15] 16] 10 7 6] 5 3] 1],
where $\sigma^* = \text{LBFS*}(\sigma, \sigma^+)$.

As mentioned previously, this σ^* has the umbrellas 11 15 10 and 11 16 10. By our choice of $\{u, v, w\}$, $u = 11, v = 16$ and $w = 10$. (See Figure 12.) The S-hierarchy for this example is:

$S_0(= S)$: [14 13 [11 15] 16]
S_1: [19 18 [9 17] 12 [14 13 [11 15] 16] 10 7 6]
S_2: [21 20 [8 2] [19 18 [9 17] 12 [14 13 [11 15] 16] 10 7 6] 5 3]
$S_3(= G)$.

For $0 \le i < g$ we define $cl'_*(S_i) = cl_*(S_i) \cap S_{i+1}$. For the example in Figure 12, $cl'_*(S_1) = \{2, 8\}$, whereas $cl_*(S_1) = \{2, 4, 8\}$. We let $S_j, j \ge 1$ be the innermost slice in the S-hierarchy that contains all of W. In our example $j = 1$; subsequently (Claim 6.18) we will show that $S_0, \ldots, S_{j-1}$ all end at v. For $0 \le i \le g$, we let L_i denote the good vertices of S_i that are clones of s_i, the first S_i vertex in σ^*. Later (Claim 6.15), we will see that the vertices in L_i cannot fly with respect to S_i. We let R_i denote the good vertices of S_i that are clones of the last S_i vertex in σ^*. Vertices of L_i (respectively R_i) will be termed the L or left (respectively R or right) end of S_i. When we say that a slice occurs in the L-to-R direction in an LBFS sweep, we mean that the vertices of L_i occur before the vertices of R_i in the sweep. A slice occurring in the R-to-L direction has $R_i < L_i$. (Recall Claim 4.22.)

In Claim 6.3 and Corollary 6.4 we saw that the only vertices in S_0 that fly with respect to S_0 must be in $F_u \cup K_{S_0}$. We now look at flyers in other slices in the S-hierarchy.

CLAIM 6.15. *The only vertices in S_i ($i > 0$) that occur in σ^* before S_{i-1} and fly with respect to S_i, are in $cl'_*(S_{i-1})$.*

Proof. Suppose $a \in S_i$, a is before S_{i-1} in σ^*, a flies to vertex b and $a \notin cl'_*(S_{i-1})$. Thus a is not adjacent to any vertices in S_{i-1}. Now, any such b must occur after W

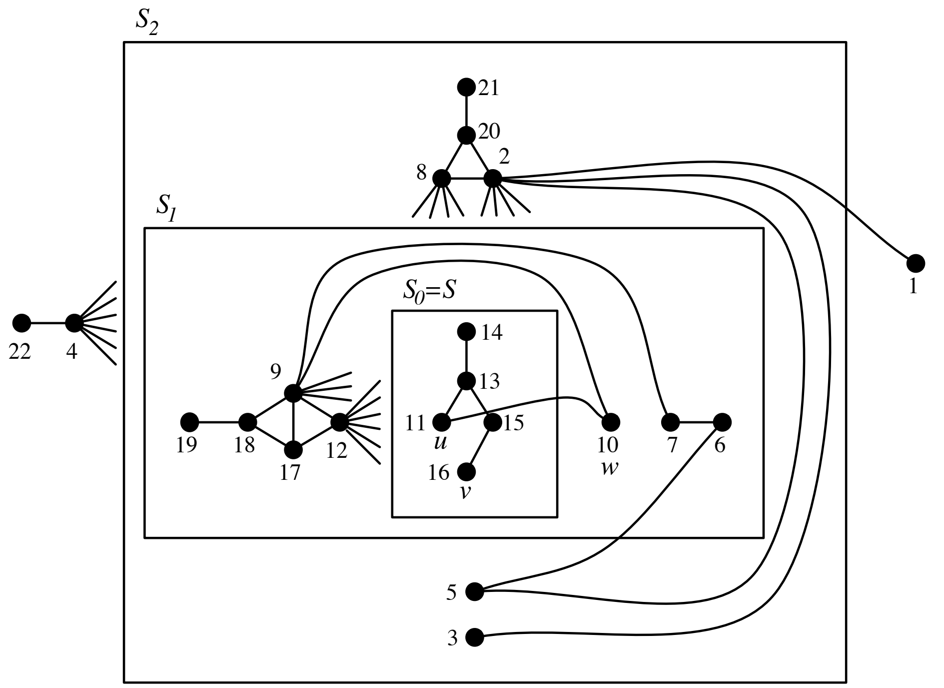

FIG. 12. *G, as seen by LBFS* on σ, σ^+*

in σ^*; otherwise, repeated applications of the P_3 Rule would contradict the fact that a is adjacent to no vertex of S_{i-1}. Since a is not adjacent to w and edge ab is an umbrella over w, either $\Gamma^*_{a,b} \supset \Gamma^*_{u,w}$ or $\Gamma^*_{a,b} = \Gamma^*_{u,w}$ and $v <_* w$; in either case the choice of u, v, and w is contradicted. Thus, the only possible flyers in S_i before S_{i-1} are in $cl'_*(S_{i-1})$. □

We now show how type 2 elements of the S-hierarchy are arranged in σ^* and, as a consequence, conclude that S_j must be of type 1.

CLAIM 6.16. *If S_i, $i > 0$, is of type 2 then the single vertex component of $S_i \backslash K_{S_i}$ is the first vertex of S_i in σ^*.*

Proof. By Corollary 4.28 and the definition of the S-hierarchy, $K_{S_i} \neq \emptyset$ and $S_i \backslash K_{S_i}$ has exactly two connected components: C_1 which contains vertices that fly out of S_i with respect to σ^*, and C_2 which consists of a single vertex that does not fly. Suppose the claim is false. S_i cannot start with a vertex of K_{S_i} since such a vertex is not good; thus, by Lemma 4.3, the single vertex of C_2 is the last vertex of $\sigma^*_{S_i}$, and $s_i \in C_1$.

If $C_2 = \{v\}$ then u flies out of S_i and therefore $u \in C_1$. Now, any s_i, u-path in C_1 contains a vertex of $cl_*(S_0) \backslash K_{S_i}$. But now s_i and v are connected by a similar path in $S_i \backslash K_{S_i}$ contradicting that $C_2 = \{v\}$.

If $C_2 = \{x\}$, $x \neq v$, then $v <_* x$ and, by Corollary 4.28, an edge from a flying vertex of C_1 forms an umbrella over x. But now the choice of u, v, w is contradicted. □

COROLLARY 6.17. *S_j is of type 1.*

Proof. Otherwise, by Claim 6.16, $S_j \backslash (K_{S_j} \cup \{s_j\})$ is an element of the S-hierarchy that contradicts the choice of j. □

Having examined the structure of the vertices of G that precede v in σ^*, we now turn our attention to the vertices of G that follow W in σ^*. Let Z be the set of vertices that occur in σ^* after $\tilde{w}$ with connected components $Z_1, \ldots, Z_p$ having first vertices $z_1, \ldots, z_p$, respectively. Note that $Z \neq \emptyset$ since otherwise, by Claim 5.1, both σ and σ^{++} would start with $\tilde{w}$, contradicting Lemmas 6.14 and 6.11. For any $z \in Z$, we let $N'(z)$ denote the neighbourhood of z that occurs before Z in σ^*; $N'(z)$ is a clique by the P_3 Rule.

For the graph in Figure 12, $\{7, 6, 5\}$ is the Z_1 set and $N'(z_1) = \{4, 8, 2, 9\}$, $Z_2 = \{3\}$ with $N'(z_2) = \{4, 2\}$ and $Z_3 = \{1\}$ with $N'(z_3) = \{2\}$. We let $ind(N'(z_i)), 1 \leq i \leq p$ denote the minimum k such that $N'(z_i) \cap S_k \neq \emptyset$. In subsequent arguments we denote $ind(N'(z_i))$ by i'. Now by Claim 6.15 and the definitions of W and the Z sets, $N'(z_i) \subseteq cl_*(S_{i'-1})$ if $i' > 0$ and $N'(z_i) \subseteq K_{S_0} \cup cl_*(S_0)$ if $i' = 0$.

We now show that all $S_0, \ldots, S_{j-1}$ end at v and that S_j must intersect Z.

Claim 6.18. *No slice in the S-hierarchy may contain an element of W unless it intersects Z. Thus in σ^*, all $S_0, \ldots, S_{j-1}$ end at v.*

Proof. Suppose to the contrary that S' is the lowest indexed element of the S-hierarchy that contains part or all of W but none of Z ($S' \neq S_0$ by the definition of S_0), and let S'' be the largest element of the S-hierarchy strictly contained in S'. Let s' be the first vertex of S' in σ^* and let w_1 and $\hat{w}$, respectively, be the first and last vertices of $S' \cap W$ in σ^*.

First we prove that S' is of type 1. Suppose not. Then, by Claim 6.16, s' must be the single vertex component of $S' \backslash K_{S'}$. Now, by Lemma 4.3, vertices of $K_{S'}$ immediately follow s' in σ^*. Thus, $S' \backslash (\{s'\} \cup K_{S'})$ is a connected component of $S' \backslash K_{S'}$ and is therefore a slice in the S-hierarchy that contradicts our choice of S'.

Next, we prove that all clones of $\hat{w}$ in S' are in W. Let X_1, X_2, K_1, K_2, X_1', X_2' be defined as in Theorem 4.20, where $\hat{w} \in X_2$. Since $\hat{w} \notin S''$, Claim 3.7 implies the existence of $w' \in cl_*'(S'')$ such that $w'\hat{w} \notin E$. The distance in $S' \backslash K_{S'}$ from s' to w' is strictly less than the distance from s' to $\hat{w}$ since $cl_*'(S'')$ minimally separates $S'' \cup W$ from s'. Now, since $K_2 \subseteq N(\hat{w})$, and by the above distance observation, $w' \notin K_2 \cup X_2'$. Any vertex in $S' \backslash W$ either has a path to s' that is shorter than $dist(s', \hat{w})$, or is adjacent to w'. Thus no such vertex can be in X_2' and $X_2' \subseteq W$. Therefore, all clones of $\hat{w}$ in S' are in W.

Finally, since S' is a slice in σ^*, s' must have been last, and a clone of $\hat{w}$ must have been first, in $sc_+(S')$ or $sc_{++}(S')$ (depending on whether s' is α or β), thus contradicting either Lemma 6.11 or Lemma 6.12. □

Now we know that S_j is the lowest slice in the S-hierarchy that contains an element of W and that it intersects Z; in particular, we assume that the last vertex of S_j is in Z_h. In subsequent arguments we will examine the appearance of S_j in sweeps when its first vertex is in Z_h. First we establish some claims about the Z sets.

Claim 6.19. *The neighbourhoods of vertices in $N'(z_i), 1 \leq i \leq p$ include the following sets:*
(i) $\tilde{z} \in N'(z_i) \cap S_0$ implies that $\tilde{z}$ is universal to all vertices other than itself that occur between s_0 and z_i, as well as all vertices of $cl_*(S_0)$.
(ii) $\tilde{z} \in N'(z_i) \cap (S_k \setminus S_{k-1}), k > 0$ implies that $\tilde{z}$ is universal to all vertices that occur between s_{k-1} and z_i, as well as all vertices other than itself of $cl_*(S_{k-1})$.

Proof. In case (i), by Claim 6.3 and Corollary 6.4, $\tilde{z} \in K_{S_0}$ and thus is universal to all vertices other than itself between s_0 and v, and all vertices of $cl_*(S_0)$. In case (ii), by Claim 6.15, $\tilde{z} \in cl_*'(S_{k-1})$ and thus is universal to all vertices between s_{k-1} and

v, as well as other vertices in S_{k-1} that follow v, and all other vertices of $cl_*(S_{k-1})$. Now suppose $\tilde{z}$ misses $y \in W \cup Z_q$ $(q < i)$ (i.e. $\tilde{z} <_* v <_* y <_* z_i$); either $\Gamma^*_{\tilde{z},z_i} \supset \Gamma^*_{u,w}$ or they are equal and $v < y$, thereby contradicting the choice of u, v, w. □

Claim 6.20. *If $z, z' \in Z$ with $z <_* z'$, then $N'(z) \supseteq N'(z')$.*

Proof. Suppose $\tilde{z} \in N'(z') \setminus N'(z)$; by the P_3 Rule. $zz' \notin E$ and $\tilde{z}z'$ is an umbrella over z. Using the previous argument, $\tilde{z}, z, z'$ would contradict our choice of u, v, w. □

Henceforth, for $i > 1$, we assume that $N'(z_i) \subset N'(z_{i-1})$. Otherwise, $N'(z_i) = N'(z_{i-1})$ and either there is a module in G or $|Z_{i-1}| = 1$. As indicated previously, such single vertex Z sets will be ignored.

Claim 6.21. *Each Z_i $(1 \le i \le p)$ is consecutive in σ^*.*

Proof. Suppose Z_i is not consecutive and consider the path in Z_i, $z_i, \cdots, z'$, z'', such that $z' <_* y <_* z''$ and $y \notin Z_i$. Note that $y \notin W$ by Lemma 6.10, and that the path is guaranteed since Z_i is connected. Again using previous arguments, it is clear that z', y, z'' contradict the choice of u, v, w. □

Since $S_j \setminus Z_h$ does not contain a module of G, $N'(z_h) \cap S_j \neq \emptyset$. Let Z_k be the rightmost Z set such that $N'(z_k) \cap S_j \neq \emptyset$. By Claim 6.20, $h \le k$. If $h = k$, then $Z_h \not\subseteq S_j$ since otherwise, S_j is a module in G.

In order to complete the proof (i.e., show that at least one of Lemmas 6.14, 6.12 or 6.11 has been violated) we use the previous technique of studying a fictitious LBFS in order to define the *T-hierarchy* which corresponds to the S-hierarchy and to identify T_1, a specific M-slice in this hierarchy. In particular, our fictitious LBFS is $\rho = $ LBFS+(σ^*). For the example in Figure 12:

ρ: [1 2 [3 4 [5 6 [7 8] 9 [10 11 [13 12] [15 14] 16 17 18] 19 20]] 21 22]

which is the same as the sweep illustrated in Figure 7, with the reversal of 12 with 13 and 14 with 15. By convention T_{p+1} is G. Sweep ρ starts in Z_p which is minimally separated from $G \setminus Z_p$ by the clique $N'(z_p)$. Let T'_p be the slice of ρ consisting of the vertices of $G \setminus Z_p$ that are universal to $N'(z_p)$. Then T_p is defined to be the connected component of T'_p that contains $S \cup W$. Recall that $i' = ind(N'(z_i))$ is the smallest integer such that $N'(z_i) \cap S_{i'} \neq \emptyset$. Claims 6.15 and 6.19 imply that T_p has the following composition. (Note: The first case can occur for any value of p'; the second case can occur only when $p' > 0$.) If $N'(z_p) \cap K_{S_{p'}} \neq \emptyset$, then T_p consists of $Z_1 \cup \ldots \cup Z_{p-1} \cup W \cup (S_{p'} \cup cl_*(S_{p'})) \setminus N'(z_p)$. Otherwise, T_p consists of $Z_1 \cup \ldots \cup Z_{p-1} \cup W \cup S_{p'-1} \cup cl_*(S_{p'-1}) \setminus N'(z_p) \cup$ vertices in $S_{p'} \setminus cl'_*(S_{p'-1})$ that are universal to $cl'_*(S_{p'-1}) \cap N'(z_p)$ and occur before $S_{p'-1}$. The good vertices of T_p in this set are denoted by $\widetilde{s_{p'}}$. Such vertices are guaranteed to exist by Claim 3.22, and are guaranteed to be in T_p since $cl_*(S_{p'-1}) \not\subseteq N'(z_p)$ by Claim 3.7. For example, in the graph in Figure 12, $p = 3$, $Z_3 = \{1\}$, $N'(z_3) = \{2\}$, $p' = 2$, and $\widetilde{s_2} = \{20\}$. T_3 is thus the slice of ρ on the vertices $\{3, 4, 5, 6, 7, 8, 9, 10, 11, 13, 12, 15, 14, 16, 17, 18, 19, 20\}$ (see Figure 12). Once $N'(z_p) \cup Z_p$ has been visited in ρ, by Lemma 4.31, slice T_p is encountered and since ρ is an LBFS+, the first T_p vertex visited in ρ is the vertex preceding z_p in σ^*, namely the last vertex of Z_{p-1}, assuming $p > 1$, otherwise it is the last vertex of W. The process continues where T_q is the M-slice identified after visiting all vertices in Z_q, $1 \le q < p$. The configuration of T_q, an element of the T-hierarchy in which $N'(z_q) \cap K_{S_{q'}} = \emptyset$, is illustrated in Figure 13.

Consider when $q = 1$ and M-slice T_1 is encountered. As before, $1'$ denotes $ind(N'(z_1))$. T_1 has the following composition. (The first case can occur for any value of $1'$ and the second case can occur only when $1' > 0$.) If $N'(z_1) \cap K_{S_{1'}} \neq \emptyset$, then T_1 consists of $W \cup (S_{1'} \cup cl_*(S_{1'})) \setminus N'(z_1)$. Otherwise, T_1 consists of $W \cup S_{1'-1} \cup cl_*(S_{1'-1}) \setminus N'(z_1) \cup$ vertices in $S_{1'} \setminus cl'_*(S_{1'-1})$ that are universal to $cl'_*(S_{1'-1}) \cap N'(z_1)$

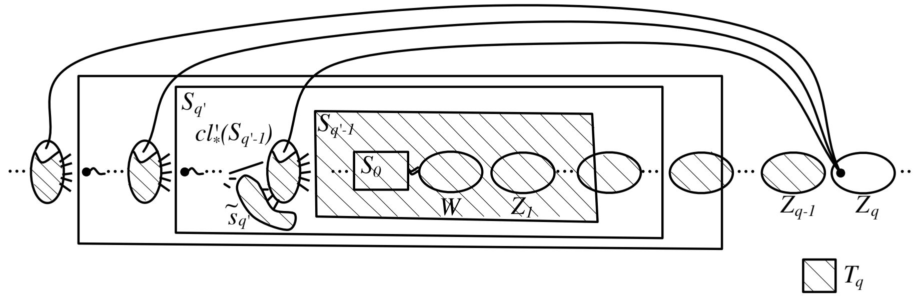

FIG. 13. T_q *shown in* σ^*

and occur before $S_{1'-1}$. Such vertices are guaranteed by Claim 3.22. Note that $\tilde{w} \in W$ is the last vertex (from the perspective of σ^*) in T_1. For our example in Figure 7:

$T_2 = \{5, 6, 7, 8, 9, 10, 11, 13, 12, 15, 14, 16, 17, 18, 19, 20\}$
$T_1 = \{10, 11, 13, 12, 15, 14, 16, 17, 18\}$.

We will prove the correctness of the algorithm by showing that, in at least one of σ, σ^+, σ^{++}, T_1 is "twisted" so that $W < S$. As with the S-hierarchy, we identify a left and right "end" of each T_q ($1 \leq q \leq p$) with respect to its ordering in σ^*. When we refer to the "ends" of T_q, we are consistent with the locations of the last vertices in T_q, as visited by σ^*. In particular, we define the R (right) end to be the clones of the last vertex in Z_{q-1} if $q > 1$ and to be the clones of $\tilde{w}$ otherwise.

The next two claims examine some properties of the elements of the T-hierarchy in good LBFSs.

CLAIM 6.22. *In $\rho = LBFS+(\sigma^*)$, for all $1 \leq i \leq p$, the L end (with respect to σ^*) of T_i flies or neighbour flies via a vertex not in K_{T_i} and the R end does not fly or neighbour fly except possibly through a vertex of K_{T_i}.*

Proof. Since G is module free and $T_i \backslash K_{T_i}$ is an M-slice in ρ, some vertex of $T_i \backslash K_{T_i}$ must fly in ρ. By Claim 4.23, every such flyer is good or adjacent to a good vertex. We show that the good vertices in the R end cannot fly or neighbour fly, thereby implying that the L end must fly or neighbour fly. The right good vertices of T_i are in Z_{i-1} or in W if $i = 1$. Thus, the right good vertices and their neighbours in $T_i \backslash K_{T_i}$ are in $Z_{i-1} \cup N'(z_{i-1}) \backslash K_{T_i}$ or in $W \cup F_u$ if $i = 1$. By Claim 6.15, $N'(z_{i-1}) \backslash K_{T_i} \subseteq cl_*(S_{i'-1}) \cup S_{i'-1}$. No vertex of $F_u \cup W \cup Z_1 \cup \ldots \cup Z_{i-1} \cup N'(z_{i-1}) \backslash K_{T_i}$ has an unvisited neighbour outside of T_i at the moment when T_i is encountered as an M-slice in ρ, since all of $cl_*(S_{i'-1})$ is already visited or in T_i and since $S_{i'-1} \cup W \subset T_i$. Therefore, the R end of T_i cannot fly or neighbour fly except via K_{T_i} and, hence, the L end must fly or neighbour fly. ▯

Note that, for all $j \leq i \leq g$, only the R end of S_i flies or neighbour flies with respect to σ^*, by Claims 4.26 and 4.27, since all flying vertices must be adjacent to the rightmost vertex of S_i by Claim 6.19.

We say that $W, Z_1, \ldots, Z_q, (1 \leq q \leq p)$ are "stacked" in a sweep if $W < Z_1 < \ldots < Z_q$ in the sweep.

CLAIM 6.23. *In any good LBFS θ, if T_ℓ, $h \leq \ell \leq k$, is in the L-to-R direction then $W, Z_1, \ldots, Z_{\ell-1}$ are stacked.*

Proof. Since $\ell \leq k$, $\ell' \leq j$, where ℓ' is the index of z_ℓ. Consider θ on T_ℓ. Since T_ℓ is L-to-R in θ, the first good vertex of θ_{T_ℓ} is a vertex of $\widetilde{s_{\ell'}}$ (if z_ℓ is adjacent to a vertex of $S_{\ell'} \backslash K_{S_{\ell'}}$) or a vertex of $L_{\ell'}$ (the left good vertices of $S_{\ell'}$; this happens if z_ℓ is adjacent only to vertices of $K_{S_{\ell'}}$ in $S_{\ell'}$). In T_ℓ, $cl_*(S_{\ell'-1}) \cap T_\ell$ minimally separates $(\widetilde{s_{\ell'}} \cup L_{\ell'}) \cap T_\ell$ from $S_{\ell'-1} \cup W \cup Z_1 \cup \ldots \cup Z_{\ell-1}$. Thus, by Lemma 4.31, all of $(\widetilde{s_{\ell'}} \cup L_{\ell'} \cup cl_*(S_{\ell'-1})) \cap T_\ell$ is visited before any vertices in $S_{\ell'-1} \cup W \cup Z_1 \cup \ldots \cup Z_{\ell-1}$ in θ. Since W is not in $S_{\ell'-1}$ and $W \subseteq S_j$, there is some vertex x in $cl_*(S_{\ell'-1}) \cap S_j$ that is not adjacent to the first vertex of W. $x \notin N'(z_\ell)$ since this would imply the existence of an umbrella over the first vertex of W, contradicting the choice of u, v, w. Thus, $x \in T_\ell$. Thus, $S_{\ell'-1}$ is visited before any vertices of $W \cup Z_1 \cup \ldots \cup Z_{\ell-1}$. At this point, all neighbours of vertices in $W \cup Z_1 \cup \ldots \cup Z_{\ell-1}$ have been visited and therefore, by the definition of W and Claim 6.20, $W <_\theta Z_1 <_\theta \ldots <_\theta Z_{\ell-1}$. □

In addition, S-hierarchy slices in the L-to-R direction sometimes imply stacking. In particular, if S_j is in the L-to-R direction in a good LBFS, and thus is an M-slice by Claim 4.32, then by the neighbourhood containment of the Z sets, $W, Z_1, \ldots, Z_k$ are "stacked" in this sweep. The importance of stacking is captured in the next claim.

CLAIM 6.24. *Let θ be an LBFS of G with sets $W, Z_1, \cdots, Z_q$ ($q \leq p$) "stacked" where the last T_{q+1} vertex is in Z_q. If T_{q+1} ($q < p$) appears as an M-slice in the subsequent LBFS+ sweep, θ^+, then T_1 must appear as an M-slice in θ^+ and $W <_{\theta^+} S$.*

Proof. Since T_{q+1} is an M-slice in θ^+, the first vertex of T_{q+1} is in Z_q (i.e. T_{q+1} is in the R-to-L direction) which implies that T_q is encountered as an M-slice with its first vertex in Z_{q-1}. Thus T_q is also in the R-to-L direction. Clearly all $T_i, 1 \leq i \leq q$ are in the R-to-L direction and in particular since T_1 is in the R-to-L direction $W <_{\theta^+} S$, as required. □

Having defined the T-hierarchy and considered the behaviour of the S-hierarchy and the T-hierarchy with respect to flying and stacking, we now return to the algorithm and examine how LBFS* chose S_j to be in the L-to-R direction in σ^*. In particular, we let τ_3 be the sweep (either σ^+ or σ^{++}) that was chosen by the algorithm to orient S_j in σ^* (τ_1 and τ_2 are the two sweeps that precede τ_3). Thus S_j is in the R-to-L direction in τ_3.

We now prove that some element of the T-hierarchy appears as a slice in τ_3, and then show that, in one of $\sigma, \sigma^+, \sigma^{++}$, we must have $W < S$ thereby proving the correctness of the algorithm.

CLAIM 6.25. *T_ℓ occurs as an M-slice in τ_3, for some $h \leq \ell \leq k$.*

Proof. Let x be the first vertex of S_j and t the first vertex of Z_h in τ_3.

Case 1. $t \leq x$.

Let U be the slice of τ_3 starting at t. Vertices in $cl_{\tau_3}(U)$ are adjacent to $t \in Z_h$ but are not in Z_h since t is the first vertex of Z_h, and are not in S_j since x is the first vertex of S_j. Therefore, $cl_{\tau_3}(U) \subseteq cl_*(S_j)$ and all of S_j is in U. Thus all of $Z_h \cap S_j$ and perhaps more of Z_h is in U, and possibly some of Z_h is after U. Note that U could be S_i for some $i \geq j$, in the R-to-L direction. As previously mentioned, $h' \leq j$ since otherwise $S_j \backslash Z_h$ is a module in G. This observation, combined with the fact that $z_h \in S_j$, implies that $cl_*(S_j) \subset N'(z_h)$ and therefore $T_h \subset S_j \subset U$.

We now look at Claim 4.31 applied to U. All of $(Z_h \cap U) \cup (N'(z_h) \cap U)$ will be visited in τ_3 before the rest of U. In particular, $N'(z_h) \cap U$ is a separator and T_h is an M-slice of the vertices in the rest of U universal to $N'(z_h) \cap U$.

Case 2. $x < t$.

We first claim that there exists a unique index ℓ, $h < \ell \leq k$, such that there is a vertex $z \in Z_\ell$ that is before x and adjacent to x. Suppose no such index exists and consider the slice U starting at x. $cl_{\tau_3}(U)$ contains only vertices in $cl_*(S_j)$ and thus all of S_j is in U. Since S_j is R-to-L and $x \notin Z_h$, x must be in $S_j \cap N'(z_h)$ and therefore adjacent to u and v by Claim 6.19. But now x is not simplicial in U contradicting the fact that τ_3 is a good LBFS. Thus such an index exists and, furthermore, it is unique since otherwise the P_3 Rule would be violated by x and neighbours of x from two distinct Z sets occurring before x.

Let z' be the leftmost vertex in Z_ℓ in τ_3 and consider U, the slice of τ_3 starting at z'. It is possible that $z' = z$ but, in general, z' may or may not be adjacent to x. Since $cl_{\tau_3}(U)$ contains only vertices in $cl_*(S_j)$, all of S_j is in U. By Claim 6.20, since $z' \in U$, all vertices of $Z_h \cup \ldots \cup Z_{\ell-1}$ are also in U. (Vertices of $Z_h \cup \ldots \cup Z_{\ell-1}$ cannot occur before U since they are all adjacent to x by Claim 6.20 and therefore would violate the P_3 Rule with x and z.) Note, however, that some vertices of Z_ℓ may occur after U. This will occur, for example, if U is some R-to-L S set that cuts Z_ℓ.

Now, as in Case 1, consider Claim 4.31 applied to U. All of $(Z_\ell \cap U) \cup (N'(z_\ell) \cap U)$ is before the rest of U in τ_3. In particular, $N'(z_\ell) \cap U$ is a separator and T_ℓ is an M-slice of the vertices in the rest of U universal to $N'(z_\ell) \cap U$. □

THEOREM 6.26. *The algorithm is correct.*

Proof. By Claim 6.25, there is a T_ℓ, $h \leq \ell \leq k$, M-slice in τ_3.

Case 1. T_ℓ is L-to-R in τ_3.

Since T_ℓ is a L-to-R M-slice in τ_3, it begins with a good vertex in its L end, which must be the last vertex of $sc_{\tau_2}(T_\ell)$. Thus T_ℓ is R-to-L in τ_2. Thus, by Claims 6.22 and 4.32, T_ℓ is an M-slice in τ_2. Therefore, T_ℓ is L-to-R in τ_1 and, by Claim 6.23, $W, Z_1, \ldots, Z_{\ell-1}$ are stacked in τ_1. But now $W <_{\tau_2} S$ by Claim 6.24, contradicting Lemma 6.12 or Lemma 6.14.

Case 2. T_ℓ is R-to-L in τ_3.

In this case, T_ℓ must be L-to-R in τ_2 and therefore, by Claim 6.23, $W, Z_1, \ldots, Z_{\ell-1}$ are stacked in τ_2. Now we have $W <_{\tau_3} S$ by Claim 6.24, contradicting Lemma 6.11 or Lemma 6.12. □

Note that we fully believe that the five sweep version is correct. In the proof of correctness of the six sweep algorithm all that we require of τ_1 is that it be good. This allows us to conclude, in Case 1 of Theorem 6.26, that we have the "stacking" required. Note that if $\tau_3 = \sigma^{++}$, (i.e. the β-Rule was invoked for S_j), then $\tau_1 = \sigma$ which is guaranteed to be good in both the five and six sweep versions. Only when $\tau_3 = \sigma^{+}$, (i.e. the α-Rule was invoked for S_j) do we require $\tau_1 = \pi$ to be good. To prove the correctness of the five sweep version, there are two possible approaches. The first would examine exactly how T_ℓ could be L-to-R in τ_3 while S_j is R-to-L in τ_3. There may be extra structure results that show this cannot happen. The second would show that the properties required of the various S and T slices in τ_3 $(= \sigma^{+})$ would contradict the structure of σ^{++}. In particular, for the algorithm to have chosen σ^{+} to orient S_j in σ^*, S_j must be in the L-to-R direction in σ^{++}, and thus a slice. Possibly the structure of superM-slices of S_j would not allow that.

7. Linear time implementation of the algorithm. We now show how the algorithm can be implemented in linear time using elementary data structures. Linear time implementations of LBFS itself are described in [30, 15, 16]. In our discussion we will follow the implementation presented in [16], namely, one that follows the paradigm of "partioning". In this scheme, we start with all vertices in the same cell (i.e. slice) and choose an arbitrary vertex, in particular, for reasons that will become clear, the first vertex in the cell. As we will see during our discussion of the implementation of LBFS+ and LBFS*, it will be very advantageous for us to assume that the vertices already have some order. When a vertex is chosen as the *pivot*, it is placed in its own cell and invokes a partitioning of all cells that follow it in the ordering. Under this partitioning of a cell, vertices that are adjacent to the pivot form a new cell that precedes the cell containing the vertices not adjacent to the pivot. After this partitioning is complete, a new pivot is chosen from the cell immediately following the old pivot and the process of refinement continues. We refer the reader to Figure 14 for an example of a few steps of partitioning on the graph of Figure 1.

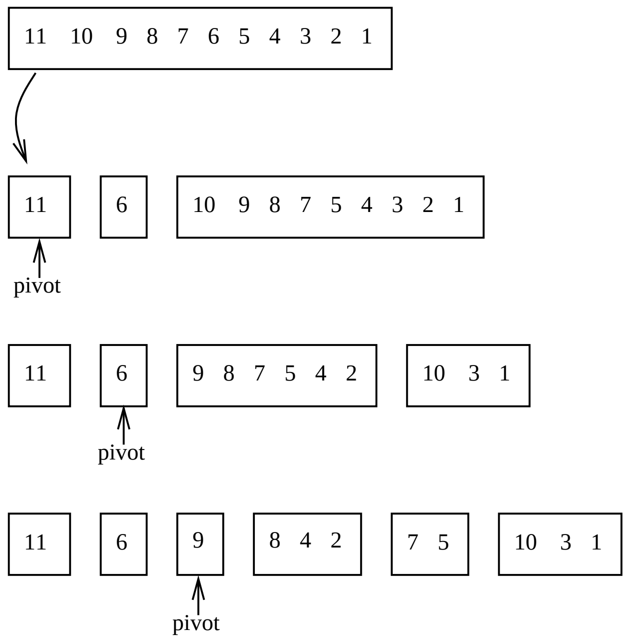

Fig. 14. *The first few steps of a partitioning*

As pointed out by Lanlignel [24], one of the advantages of using this paradigm is that we immediately have an implementation of LBFS+. Once our initial LBFS has terminated, we merely reverse the ordering of the vertices produced by the first LBFS and run the algorithm again. Every time a slice is encountered, the last vertex from the previous LBFS is automatically the vertex at the front of the list. The example in Figure 14 is the LBFS+ for the output of the LBFS indicated by the vertex numbering in Figure 1.

We now turn to the issue of implementing LBFS*. The first problem is to identify α and β easily. To do this, we follow the idea used to implement LBFS+. This time we run two parallel LBFSs, one on the reverse of the σ^+ ordering and the other on the reverse of the σ^{++} ordering. As seen above, once a slice is identified, the α vertex is the first vertex in the σ^+ ordering, whereas the β vertex is the first vertex in the σ^{++} ordering. Once we have chosen which vertex should be first in the σ^* ordering, we place it first in both orderings and continue with the partitioning process. For example, if the β vertex is chosen, the LBFS that is operating on the σ^{++} ordering is not touched whereas the LBFS that is operating on the σ^+ ordering removes β from the slice and places it first before continuing. When the sweep is completed, both of the parallel LBFSs contain σ^*.

Thus, the only remaining problem is to implement efficiently the choice between α and β. To do this, we first introduce the *neighbour index value*. For any LBFS τ

and vertex x, the neighbour index value of x in τ, denoted $i_\tau(x)$ is the index, with respect to τ, of the rightmost vertex of $N[x]$. As previously, we let $i_+(x), i_{++}(x)$ and $i_*(x)$ denote $i_{\sigma^+}(x), i_{\sigma^{++}}(x)$, and $i_{\sigma^*}(x)$, respectively. Thus $i_+(x) > x$ if and only if x is adjacent to some vertex to the right of it in σ^+ and $i_+(x) = x$ if and only if x is not adjacent to any vertex to the right of it. As described above, we can easily find the last vertex of $sc_+(S)$ and $sc_{++}(S)$, i.e. α and β respectively for any slice S. Hereafter, all flying and neighbour flying are assumed to be in σ^* and with respect to slice S.

CLAIM 7.1. *Let S be a nonclique slice in σ^* and let τ be any LBFS of G where x is the last vertex of $sc_\tau(S)$. In σ^*, x is a flyer if and only if $i_\tau(x) > x$.*

Proof. For the "only if" part, suppose xy is a flying edge in σ^*. In σ^* there is a vertex $y' \in cl_*(S)$ such that y' is universal to S but is not adjacent to y. By the P_3 Rule, and since S is not a clique, $y' <_\tau x$. If $y <_\tau x$ then we have a P_3: y', x, y. Thus, y must occur after x in τ and $i_\tau(x) > x$.

For the "if" part, suppose $i_\tau(x) > x$ as witnessed by vertex y. Since S is not a clique, x is not universal to S (since x is last in $sc_\tau(S)$ and by the P_3 Rule). Thus there is a vertex $v \in S$ such that $xv \notin E$. Since y is after x in τ, then by the P_3 Rule, y is not adjacent to v and thus y is not universal to S. Thus, in σ^*, y must be after S by Claim 3.5. ∎

COROLLARY 7.2. *To check to see if α (respectively β) flies (when S is not a clique) we just have to check whether $i_+(\alpha) > \alpha$ (respectively $i_{++}(\beta) > \beta$).*

CLAIM 7.3. *Let S be a nonclique slice in the σ^* sweep and let τ be any LBFS of G with x the last vertex in $sc_\tau(S)$. If $i_\tau(x) = x$, then we have the following:*

1. *If $diam(S) \neq 2$:*
   x neighbour flies if and only if there is a $y \in S$ such that $xy \in E$ and $i_\tau(y) > x$.
2. *If $diam(S) = 2$:*
   (a) *If there is a $y \in S$ such that $xy \in E$ and $i_\tau(y) > x$, then x neighbour flies.*
   (b) *If there is no such $y \in S$, and x neighbour flies, then all neighbours of x that fly are universal to S.*

Proof.

1. $diam(S) \neq 2$:
   Thus y, an arbitrary neighbour of x in S is not universal to S and there is a vertex $v \in S$ such that $yv \notin E$. Furthermore by the P_3 Rule in τ, $xv \notin E$.
   For the "only if" part, suppose x neighbour flies with yz a flying edge where $y \in S$ and $xy \in E$. In σ^* there is a vertex z' in $cl_*(S)$ such that $z'z \notin E$. By the P_3 Rule, $z' <_\tau x$ and $zv \notin E$ (by the P_3 Rule in σ^*). Suppose that $z <_\tau x$. Now in τ we have unrelated vertices v and z with respect to x contradicting Theorem 3.23 (the paths are x, z', v and x, y, z). Recall that since x neighbour flies, it does not fly and thus $xz \notin E$. Thus $i_\tau(y) > x$.
   For the "if" part, suppose $i_\tau(y) > x$ and let z be a vertex after x in τ such that $yz \in E$. Then $xz \notin E$ since $i_\tau(x) = x$ and thus z is not universal to S. In σ^*, z must be after S by Claim 3.5.
2. $diam(S) = 2$:
   (a) As in "if" part above.
   (b) Suppose the claim is false and that there is y, a non-universal vertex of S that flies in S via edge yz. In τ, z is before x. Since y is not universal to S, there is a vertex v such that $yv \notin E$. Now using the same technique

as in “only if” part above, we see that x is not admissible in τ.

□

We now elaborate on some steps of the interval graph recognition algorithm and discuss the complexity of the algorithm.

Interval Graph Recognition Algorithm

Step 1. LBFS from an arbitrary vertex x and let y be the last vertex visited by this sweep, π'.

Step 2. Using π', LBFS+ (from y) and let z be the last vertex visited by this sweep, π.

Step 3. Using π, LBFS+ (from z) and let y be the last vertex visited by this sweep, σ.

Step 4. Using σ, LBFS+ (from y); by Claim 5.1, z is the last vertex visited by this sweep, σ^+.
In this sweep, for every vertex $v \in V$, calculate its index $i_+(v)$ and its A set where $A_v = \{z | zv \in E, z <_+ v \wedge i_+(z) > v\}$. Note that we link the elements of these sets so that a particular element may be removed from all A sets in O(the number of occurrences of the element). We also store the cardinalities of these sets and update these cardinalities every time a vertex is removed.

Step 5. Using σ^+, LBFS+ (from z, ending at y, by Claim 5.1), creating the sweep σ^{++}.
In this sweep, for every vertex $v \in V$, calculate its index $i_{++}(v)$ and its B set where $B_v = \{z | zv \in E, z <_{++} v \wedge i_{++}(z) > v\}$. Again the elements are linked to allow fast deletion and the cardinalities are stored and updated.

Step 6. Using σ^+ and σ^{++}, LBFS*, creating the sweep σ^*.
As each vertex is visited in this sweep, it is removed from all the A and B sets that contain it and the cardinalities of these sets are updated. The cost to do the removal of vertex v is bounded by $O(deg(v))$. We now indicate how ties are broken in this sweep, i.e. how the table in §5.1 is implemented. Assume that slice S has been identified with α the last vertex of $sc_+(S)$ and β the last vertex of $sc_{++}(S)$.

1. if $i_+(\alpha) > \alpha$ (i.e. α flies), then choose β.
2. else if $i_{++}(\beta) > \beta$ (i.e. β flies), then choose α.
3. else if $(|B_\beta| = 0) \vee (|A_\alpha| \neq 0)$, then choose β.
4. else let y be an arbitrary element of B_β; if $i_+(y) = \alpha$ then choose β else choose α.

Note that since the cardinalities are stored, all steps can be done in $O(1)$ time.

Step 6. If σ^* is umbrella-free then G is an interval graph; otherwise, G is not interval.

CLAIM 7.4. *LBFS* can be implemented to run in linear time.*

Proof. First we show that the above algorithm correctly determines whether α or β should be chosen. If S is a clique, then we do not care which vertex is chosen since there can be no umbrella u, v, w where u and v are in S. Thus we assume that S is not a clique and go through each part of Step 5 and show that we have made the correct choice for α and β. We also refer to the appropriate row and column of the algorithm table. Again, all flying and neighbour flying are assumed to be with respect to S.

(a) If $i_+(\alpha) > \alpha$ (i.e. α flies by Corollary 7.2), then choose β.
From the first row of the table, if α flies, then β must be chosen.

(b) If $i_{++}(\beta) > \beta$ (i.e. β flies by Corollary 7.2), then choose α.
Since α does not fly, and β does, entries (2,1) and (3,1) of the table show that α has to be chosen.
Note that at this point neither α nor β fly.

(c) If $(|B_\beta| = 0) \vee (|A_\alpha| \neq 0)$, then choose β.
Case 1: $|B_\beta| = 0$.
There are a few subcases to consider here. If $diam(S) \neq 2$ then we know by Case 1 of Claim 7.3 that β does not neighbour fly and thus is "OK". Regardless of α's condition we choose β. (See entries (2,3) and (3,3) of the table.) If $diam(S) = 2$, then we know by Case 2(b) of Claim 7.3 that if β neighbour flies, then all neighbours of β that fly are universal to S and thus are adjacent to α too (i.e. α neighbour flies). Thus either β is "OK" and we choose β (see entries (2,3) and (3,3)) or β neighbour flies, in which case so does α and we choose β (entry (2,2)).
Case 2: $|A_\alpha| \neq 0$.
Thus α neighbour flies and regardless of β's condition we choose β (entries (2,2) and (2,3)).

(d) If $i_+(y) = \alpha$, then choose β else choose α.
Now β neighbour flies but we're not sure whether α neighbour flies (in the case that $diam(S) = 2$). First we show that α neighbour flies if and only if for arbitrary y in B_β, $\alpha y \in E$. If $\alpha y \in E$, then clearly α neighbour flies.
Now assume α neighbour flies. Thus, by Claim 7.3, $diam(S) = 2$. If $\alpha\beta \in E$, then immediately $\alpha y \in E$ by the P_3 Rule on the σ^{++} sweep. Next we show that the only neighbours of β that fly are universal to S. To see this, note that since both α and β neighbour fly, the algorithm will choose β (entry (2,2)). If β has a neighbour that flies and this neighbour is not adjacent to α, then there will be an umbrella over α contradicting the correctness of the algorithm. (Note that all neighbours of β in S will be visited before non-neighbours.)
Our final point is that $ay \in E$ if and only if $i_+(y) = \alpha$. If $ay \in E$ and if $i_+(y) > \alpha$, then immediately we contradict $|A_\alpha| = 0$. If $i_+(y) = \alpha$, then by definition $ay \in E$.
Now we have shown that α neighbour flies if and only if $i_+(y) = \alpha$. If so, we choose β (entry (2,2)) and otherwise we choose α (entry (3,2)).

The description of the data structures used to choose α and β show that the algorithm can be easily implemented to run in linear time. □

THEOREM 7.5. *The six sweep interval graph recognition algorithm is implementable to run in $O(n + m)$ time.*

Proof. This follows immediately from Claim 7.4 and the observation that it is easy to determine if a given ordering of V, in our case the σ^* ordering, is an I-ordering. □

8. Concluding remarks and open problems. Apart from being umbrella free, the final LBFS of our interval graph recognition algorithm has other interesting properties that we now discuss.

First, directly from this final sweep we can construct an interval representation of G by representing vertex v by the interval $[\sigma^*(v), i_*(v)]$. (Recall that $i_*(v) = \sigma^*(v)$ if v is not adjacent to any vertices that follow it in σ^*.) The interval representation for the graph in Figure 12 is presented in Figure 15. Note that we have slightly extended the intervals to make the intersections more obvious.

Secondly, either from such an interval representation, or directly from the final

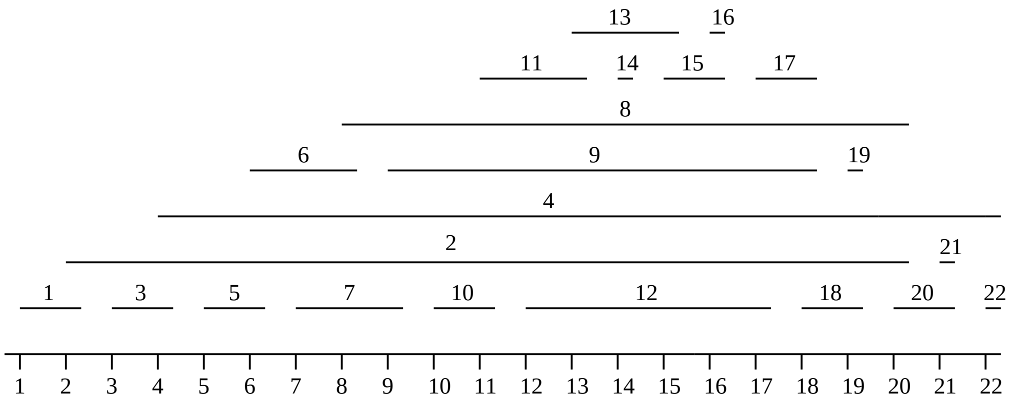

FIG. 15. *Interval representation for the graph in Figure 7*

sweep, we can easily construct a linear ordering of the maximal cliques of G (i.e. an ordering that satisfies Theorem 1.2). To construct such a set from the interval representation, sweep from left to right. For each position determine the intervals that overlap this position; this set clearly forms a clique, although it may not be maximal. To ensure that we only keep maximal cliques, the clique at position i is rejected if it is strictly included in the clique at position $i+1$. For our example the set of cliques in linear order are:

1, 2
2, 3, 4
2, 4, 5, 6
2, 4, 6, 7, 8
2, 4, 7, 8, 9
2, 4, 8, 9, 10, 11
2, 4, 8, 9, 11, 12, 13
2, 4, 8, 9, 12, 13, 14
2, 4, 8, 9, 12, 13, 15
2, 4, 8, 9, 12, 15, 16
2, 4, 8, 9, 12, 17
2, 4, 8, 9, 17, 18
2, 4, 8, 18, 19
2, 4, 8, 20
4, 20, 21
4, 22

We now turn our attention to two open problems. The first is whether the five sweep algorithm is correct (as we believe). The second is whether the techniques used in this paper can be used to recognize other families of graphs. (Note that unit interval graphs [5] and cographs [3] also have multisweep LBFS recognition algorithms.) The most important such family is *cocomparability graphs* (the complement has a transitive orientation). Cocomparability graphs are strictly between AT-free graphs and interval graphs and have a characterization very similar to that of interval graphs (i.e. Theorem 1.3) as observed by Kratsch and Stewart [23].

OBSERVATION 8.1. *[23] A graph $G = (V, E)$ is a cocomparability graph if and only if there exists a linear order $\prec$ on the set of its vertices such that for every choice* *of vertices* u, v, w, *with* $u \prec v$ *and* $v \prec w$, $uw \in E$ *implies* $uv \in E$ *or* $vw \in E$.

We call such an ordering a *CO-ordering*. By examining the LBFS+ of any CO-ordering, it is straightforward to see that this new sweep is both an LBFS and a CO-ordering [7]. This raises the fascinating question of whether there is a multi-sweep LBFS algorithm to find a CO-ordering of a cocomparability graph.

Finally we note that the LBFS structure on interval graphs and their superfamilies presented in §4 and §3 respectively should be of use in solving other algorithmic problems on interval graphs as well as to give insight into the various applications of interval graphs.

Acknowledgements. The authors thank Michel Habib, Pavol Hell and Christophe Paul for their many helpful suggestions in the preparation of the paper.

REFERENCES

[1] S. Benzer, *On the topology of the genetic fine structure*, Proc. Natl. Acad. Sci. USA, 45 (1959), pp. 1607–1620.

[2] K. S. Booth and G. S. Lueker, *Testing for the consecutive ones property, interval graphs and graph planarity using PQ-tree algorithms*, J. Comput. System Sci., 13 (1976), pp. 335–379.

[3] A. Bretscher, D. G. Corneil, M. Habib and C. Paul, *A simple linear time LexBFS cograph recognition algorithm*, SIAM J. Discrete Math., 22 (2008), pp. 1277–1296.

[4] V. Choi and M. Farach-Colton, *Barnacle: an assembly algorithm for clone-based sequences of whole genomes*, Gene, 320 (2003), pp. 165–176.

[5] D. G. Corneil, *A simple 3-sweep LBFS algorithm for the recognition of unit interval graphs*, Discrete Appl. Math., 138 (2004), pp. 371–379.

[6] D. G. Corneil, *Lexicographic breadth first search - a survey*, in 30th International Workshop on Graph Theoretic Concepts in Computer Science, WG2004, LNCS 3353 (Springer) (2004) pp. 1–19.

[7] D. G. Corneil, B. Dalton, and M. Habib, *LDFS based algorithms for Hamiltonicity problems on cocomparability graphs*, in preparation.

[8] D. G. Corneil, E. Köhler, and J-M. Lanlignel, *On end-vertices of Lexicographic Breadth First Searches*, Discrete Appl. Math., to appear.

[9] D. G. Corneil, S. Olariu, and L. Stewart, *Asteroidal triple-free graphs*, SIAM J. Discrete Math., 10 (1997), pp. 399–430.

[10] D. G. Corneil, S. Olariu, and L. Stewart, *The ultimate interval graph recognition algorithm? (Extended Abstract)*, in Ninth Annual ACM-SIAM Symposium on Discrete Algorithms SODA, 1998, pp. 175–180.

[11] D. G. Corneil, S. Olariu, and L. Stewart, *Linear time algorithms for dominating pairs in asteroidal triple-free graphs*, SIAM J. Comput., 28 (1999), pp. 1284–1297.

[12] G. Dirac, *On rigid circuit graphs*, Abh. Math. Sem. Univ. Hamburg, 25 (1961), pp. 71–76.

[13] F. F. Dragan, F. Nicolai, and A. Brandstädt, *LexBFS-orderings and powers of graphs*, in Graph Theoretic Concepts in Computer Science WG'96, Lecture Notes in Computer Science 1197, Springer-Verlag, Berlin, 1997, pp. 166–180.

[14] P. C. Gilmore and A. J. Hoffman, *A characterization of comparability graphs and interval graphs*, Canad. J. Math., 16 (1964), pp. 539–548.

[15] M. C. Golumbic, *Algorithmic Graph Theory and Perfect Graphs*, Academic Press, New York, 1980. 2nd Edition: Volume 57 of Annals of Discrete Math., Elsevier, 2004.

[16] M. Habib, R. M. McConnell, C. Paul and L. Viennot, *Lex-BFS and partition refinement, with applications to transitive orientation, interval graph recognition and consecutive ones testing*, Theoret. Comput. Sci., 234 (2000), pp. 59–84.

[17] W. L. Hsu, *A simple test for interval graphs*, in Graph Theoretic Concepts in Computer Science WG'92, Wiesbaden-Naurod, Germany, June 1992, Lecture Notes in Computer Science 657, E. W. Mayr, ed., Springer-Verlag, Berlin, 1993, pp. 11–16.

[18] W. L. Hsu and T. H. Ma, *Fast and simple algorithms for recognizing chordal comparability graphs and interval graphs*, SIAM J. Comput. 28 (1999), pp. 1004–1020.

[19] W. L. Hsu and R. M. McConnell, *PC trees and circular-ones arrangements*, Theoret. Comput. Sci., 296 (2003), pp. 59–74.

[20] B. Jamison and S. Olariu, *On the semi-perfect elimination*, Adv. in Appl. Math., 9 (1988), pp. 364–376.

[21] N. Korte and R. H. Möhring, *An incremental linear-time algorithm for recognizing interval graphs*, SIAM J. Comput., 18 (1989), pp. 68–81.

[22] D. Kratsch, R. M. McConnell, K. Mehlhorn, and J. P. Spinrad, *Certifying algorithms for recognizing interval graphs and permutation graphs*, SIAM J. Comput., 36 (2006), pp. 326–353.

[23] D. Kratsch and L. Stewart, *Domination on cocomparability graphs*, SIAM J. Discrete Math., 6 (1993), pp. 400–417.

[24] J-M. Lanlignel, Private communications, 1999.

[25] C. G. Lekkerkerker and J. C. Boland, *Representation of a finite graph by a set of intervals on the real line*, Fund. Math., 51 (1962), pp. 45–64.

[26] T. Ma, unpublished manuscript.

[27] S. Olariu, *An optimal greedy heuristic to color interval graphs*, Inform. Process. Lett., 37 (1991), pp. 65–80.

[28] G. Ramalingam and C. Pandu Rangan, *A uniform approach to domination problems on interval graphs*, Inform. Process. Lett., 27 (1988), pp. 271–274.

[29] A. Raychaudhuri, *On powers of interval and unit interval graphs*, Congr. Numer. 59 (1987), pp. 235–242.

[30] D. J. Rose, R. E. Tarjan, and G. S. Lueker, *Algorithmic aspects of vertex elimination on graphs*, SIAM J. Comput., 5 (1976), pp. 266–283.

[31] K. Simon, *A new simple linear algorithm to recognize interval graphs*, in International Workshop on Computational Geometry CG'91, Bern, Switzerland, March 1991, Lecture Notes in Computer Science 553, H. Bieri and H. Noltemeyer, eds., Springer-Verlag, Berlin, 1992, pp. 289–308.

Index.