

A survey on Algorithmic Aspects of Modular Decomposition¹

Michel Habib^a and Christophe Paul^b

^a*LIAFA, Univ. Paris Diderot - Paris VII, France*

^b*CNRS, LIRMM, Univ. Montpellier II, France*

Abstract

The modular decomposition is a technique that applies but is not restricted to graphs. The notion of module naturally appears in the proofs of many graph theoretical theorems. Computing the modular decomposition tree is an important preprocessing step to solve a larger number of combinatorial optimization problems. Since the first polynomial time algorithm in the early 70's, the algorithmic of the modular decomposition has known an important development. This paper survey the ideas and techniques that arose from this line of research.

Key words: Combinatorial algorithms, Graph theory, Modular decomposition

1 Introduction

Modular decomposition is a technique at the crossroads of several domains of combinatorics which applies to many discrete structures such as graphs, set systems, matroids among others. As a graph decomposition technique it has been introduced by Gallai [Gal67] to study the structure of comparability graphs (those graphs whose edge set can be transitively oriented). Roughly speaking a *module* in graph is a subset M of vertices which share the same neighbourhood outside M . Galai showed that the family of modules of an undirected graph can be represented by a tree, the *modular decomposition tree*. The notion of module appeared in the litterature as *closed sets* [Gal67],

Email addresses: habib@liafa.jussieu.fr (Michel Habib), paul@lirmm.fr (Christophe Paul).

¹ Research supported by the project ANR "Décomposition des graphes et algorithmes"

clan [EGMS94], *autonomous sets* [Möh85b], *clumps* [Bla78]. . . while the modular decomposition is also called *substitution decomposition* [Möh85a] or *X-join decomposition* [HM79]. See [MR84] for an early survey on this topic.

There is a large variety of combinatorial applications of modular decomposition. Modules can help proving results on graph structure like Galai did in its study of comparability graphs. More generally modular decomposition appears in (but is not limited to) the context of perfect graph theory. Indeed Lovász's proof of the perfect graph theorem [Lov72] involves cliques modules. Notice also that a number of perfect graph classes can be characterized by properties of their modular decomposition tree: cographs, P_4 -sparse graphs, permutation graphs, interval graphs. . . Refer to the books of Golubic [Gol80], Brandstädt *et al.* [BLS99] for graph classes. We should also mention that the modular decomposition tree is useful to solve optimization problems on graphs or other discrete structures (see [Möh85b]). An example of such use is given in the last section.

In the late 70's, the modular decomposition has been independently generalized to *partitive set families* [CHM81] and to a combinatorial decomposition theory [CE80] which applies to graphs, matroids and hypergraphs. More recently, the theory of partitive families and its variants had been the foundation of decomposition schemes for various discrete structures among which 2-structures [EHR99] and permutations [UY00, BCdMR08]. Beside, based on efficiently representable set families, different graph decompositions had been proposed. The *split decomposition* of [CE80] relies on a *bipartitive family* on the vertex set. Refer to [BX08] for a survey on the recent developments of these techniques.

A good feature of these decomposition scheme is that they be computed in polynomial time. Indeed, since the early 70's, there have been a number algorithms for computing the modular decomposition of a graph (or for some variants of this problem). The first polynomial algorithm is due to Cowan, James and Stanton [CJS72] and runs in $O(n^4)$. Successive improvements are due to Habib and Maurer [HM79] who proposed a cubic time algorithm, and to Müller and Spinrad who designed a quadratic time algorithm. The first two linear time algorithms appeared independently in 1994 [CH94, MS94]. Since then a series of simplified algorithms has been published, some of them running in linear time [MS99, TCHP08], some others in almost linear time [DGM01, MS00, HPV99]. The list is not exhaustive. In this line of research, a series of very interesting algorithmic techniques had been developed, which we believe, could be useful in other applications or topics of computer science. This paper surveys the algorithmic theory of modular decomposition.

The paper is organized as follows. The partitive family theory and its application to modular decomposition of graphs is presented in Section 2. As

an algorithmic appetizer, Section 3 addresses the special case of totally decomposable graphs, namely the *cographs*, for which a linear time algorithm is known since 1985 [CPS85]. *Partition refinement* is an algorithmic techniques that reveals to be really powerful for the modular decomposition problem, but also in for other graphs applications (see *e.g.* [PT87, HPV99]). Section 4 is devoted to partition refinement. Section 5 describes the principle of a series of modular decomposition algorithms developed in the mid 90's. Section 6 explains how the modular decomposition can be efficiently computed via the recent concept of *factoring permutation* [CHdM02]. Let us mention that we do not discuss the recent linear time algorithm of [TCHP08] as it mainly merge the ideas developed in Sections 5 and 6. The purpose of this paper is not to enter into the details of all the algorithm techniques but rather to present their main lines. Finally the last section presents three recent applications of the modular decomposition in three different domains of computer science, namely pattern matching, computational biology and parameterized complexity.

2 Partitive families

The *modular decomposition theory* has to be understood as a special case of the theory of *partitive family* whose study dates back to the early 80's [CE80, CHM81]. We briefly present the mains concepts and theorems of the partitive family theory. We then introduce the *modular decomposition* of graphs and discuss its elementary algorithmic aspects. This section ends with a discussion on two important class of graphs: indecomposable graphs (the *prime* graphs) and on totally decomposable graphs (known as the *cographs*)

2.1 Decomposition theorem of partitive families

The *symmetric difference* between two sets A and B is denoted by $A \Delta B = (A \setminus B) \cup (B \setminus A)$. Two subsets A and B of a set S *overlap* if $A \cap B \neq \emptyset$, $A \setminus B \neq \emptyset$ and $B \setminus A \neq \emptyset$, we write $A \perp B$.

Definition 1 A family $\mathcal{S} \subseteq 2^S$ of subsets of S is partitive if:

- (1) $S \in \mathcal{S}$, $\emptyset \notin \mathcal{S}$ and for all $x \in S$, $\{x\} \in \mathcal{S}$;
- (2) For any pair of subsets $A, B \in \mathcal{S}$ such that $A \perp B$:
 - (a) $A \cap B \in \mathcal{S}$;
 - (b) $A \setminus B \in \mathcal{S}$ and $B \setminus A \in \mathcal{S}$;
 - (c) $A \cup B \in \mathcal{S}$;
 - (d) $A \Delta B \in \mathcal{S}$.

A family is *weakly partitive* whenever condition (2.d) is not satisfied. Some definitions and results have to be adapted in the context of weakly partitive family [CE80]. Unless explicitly mentioned, we will consider partitive families.

Definition 2 An element $F \in \mathcal{S}$ is *strong* if it does not overlap any other element of \mathcal{S} . The set of strong elements of \mathcal{S} is denoted \mathcal{S}_F .

Obviously any trivial subset of \mathcal{S} , namely S or $\{x\}$ (for $x \in S$), is a strong element. Let us remark that \mathcal{S}_F is nested, *i.e.* the transitive reduction of the inclusion order of \mathcal{S}_F is a tree $T_{\mathcal{S}}$ (see Figure 1). It follows that $|\mathcal{S}_F| = O(|S|)$.

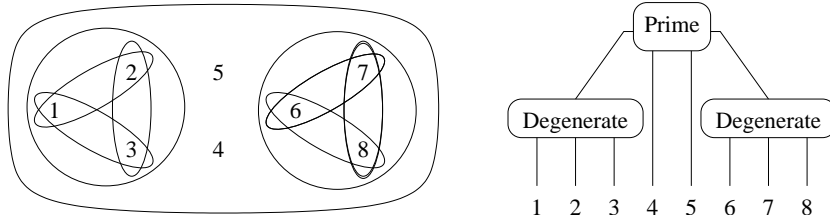


Fig. 1. The inclusion tree of the strong elements of the family $\mathcal{S} = \{\{1, 2, 3, 4, 5, 6, 7, 8\}, \{1, 2, 3\}, \{7, 8, 9\}, \{1, 2\}, \{2, 3\}, \{3, 4\}, \{4, 1\}, \{1, 3\}, \{2, 4\}, \{7, 8\}, \{8, 9\}, \{7, 9\}, \{1\}, \{2\}, \{3\}, \{4\}, \{5\}, \{6\}, \{7\}, \{8\}, \{9\}\}$

Definition 3 Let f_1^q, \dots, f_k^q be the children of a node q of $T_{\mathcal{S}}$. The node q is *degenerate* if for all non-empty subset $J \subset [1, k]$, $\cup_{j \in J} f_j^q \in \mathcal{S}$. A node is *prime* if for all non-empty subset $J \subset [1, k]$, $\cup_{j \in J} f_j^q \notin \mathcal{S}$.

It is not difficult to see that any strong element is either prime or degenerate. Moreover the following theorem tell us that the tree $T_{\mathcal{S}}$ is a representation of the family \mathcal{S} and the subfamily of strong elements \mathcal{S}_F of \mathcal{S} defines a "basis" of \mathcal{S} .

Theorem 1 [CHM81] Let \mathcal{S} be a partitive family on S . The subset $A \subseteq S$ belongs to \mathcal{S} if and only if A is strong or there exists a degenerate strong element A' (or a node of $T_{\mathcal{S}}$) such that A is the union of a strict subset of the children of A' in $T_{\mathcal{S}}$.

As a consequence, any partitive family, even of exponential size, admits a representation linear in the size of S . Such a representation property trivial is also known for other families of subsets of a set, such as laminar families, cross-free families [EG97]... as well as for some families of bipartitions of a set, such as splits [CE80]. Recently, a similar result has been shown for *union-difference* families of subsets of a set, *i.e.* families closed under the union and the difference of its overlapping elements [BXH08]. In this latter case, the size of the representation amounts to $O(|S|^2)$. For a detailed study of this aspects, the reader should refer to [BX08].

2.2 Factoring Permutations

Although the idea of *factoring permutation* implicitly appeared in some early papers (see *e.g.* [HM91, Hsu92, HHS95]), it has only been formalized in [CH97, Cap97]. This concept turns out to be central to recent modular decomposition algorithms.

Let σ be a permutation of a set S of size n . By $\sigma(x)$, we mean the rank i of x in σ and $\sigma^{-1}(x)$ stands for the i -th element of σ . A subset $I \subseteq S$ is a *factor* or an *interval* of a permutation σ if there exist $i \in [1, n]$ and $j \in [1, n]$ such that $I = \{x \mid x = \sigma^{-1}(k), i \leq k \leq j\}$. In other words, the elements of I occur consecutively in σ .

Definition 4 [Cap97] *Let \mathcal{S} be a (weakly) partitive family of a set S and let \mathcal{S}_F be the strong elements of \mathcal{S} . A permutation σ of S is factoring for \mathcal{S} if for any $F \in \mathcal{S}_F$, F is a factor of σ .*

For example, $\pi = 1\ 2\ 3\ 4\ 5\ 6\ 7\ 8\ 9$, $\pi_1 = 6\ 5\ 2\ 1\ 4\ 3\ 7\ 9\ 8$ and $\pi_2 = 9\ 8\ 7\ 5\ 1\ 3\ 2\ 4\ 6$ are three factoring permutations of the family \mathcal{S} depicted in Figure 1. One can check that, in each of these three permutations, the two non-trivial strong elements of \mathcal{S}_F , namely $\{1, 2, 3, 4\} \in \mathcal{S}_F$ and $\{7, 8, 9\} \in \mathcal{S}_F$, are factors.

Given a layout of the tree structure of a partitive family, a left-to-right enumeration of the leaves results in a factoring permutation. In many cases it is easier to compute a factoring permutation than the modular decomposition tree. We explain in Section 6.3 how to obtain the tree from a factoring permutation.

2.3 Modules of a graph

For the sake of the presentation we only consider undirected, simple and loopless graphs. We use the classical notations (*e.g.* see [BLS99]). The neighbourhood of a vertex x in a graph $G = (V, E)$ is denoted $N_G(x)$ and its non-neighbourhood $\overline{N}_G(x)$ (subscript G will be omitted when the context is clear). The complementary graph of a graph G is denoted by \overline{G} . Given a subset of vertices $X \subseteq V$, $G[X]$ is the subgraph induced by X (any edge in G between two vertices in X belongs to $G[X]$).

Let M be a set of vertices of a graph $G = (V, E)$ and x be a vertex of $V \setminus M$. Vertex x *splits* M (or is a *splitter* of M), if there exist $y \in M$ and $z \in M$ such that $xy \in E$ and $xz \notin E$. If x is not a splitter of M , then M is *uniform* or *homogeneous* with respect to x .

Definition 5 Let $G = (V, E)$ be a graph. A set $M \subseteq V$ of vertices is a module if M is homogeneous with respect to any $x \notin M$ (i.e. $M \subseteq N(x)$ or $M \cap N(x) = \emptyset$).

Aside the singletons and the whole vertex sets, the unions of any subset connected components (or of co-connected components) of a graph are simple examples of modules.

Lemma 1 [CHM81] The family \mathcal{M} of modules of a graph is *partitive*.

The notions of *trivial* and *strong* module and *degenerate* and *prime* graph are defined according to the terminology of Section 2.1. By Lemma 1, if M and M' are overlapping modules, then $M \setminus M'$, $M' \setminus M$, $M \cap M'$, $M \cup M'$ and $M \Delta M'$ are modules of G .

Let M and M' be disjoint sets. We say that M and M' are *adjacent* if any vertex of M is adjacent to all the vertices of M' and *non-adjacent* if any vertex of M is non-adjacent to all the vertices of M' .

Observation 1 Two disjoint modules are either adjacent or non-adjacent.

A module M is *maximal* with respect to a set S of vertices, if $M \subset S$ and there is no module M' such that $M \subset M' \subset S$. If the set S is not specified, we shall assume $S = V$.

Definition 6 Let $\mathcal{P} = \{M_1, \dots, M_k\}$ be a partition of the vertex set of a graph $G = (V, E)$. If for all i , $1 \leq i \leq k$, M_i is a module of G , then \mathcal{P} is a modular partition (or congruence partition) of G .

A modular partition $\mathcal{P} = \{M_1, \dots, M_k\}$ which only contains maximal strong modules is a *maximal modular partition*. Notice that each graph has a unique maximal modular partition. If G (resp. \overline{G}) is not connected then its (resp. co-connected) connected components are the elements of the maximal modular partition. From Observation 1, we can define a *quotient graph* whose vertices are the parts (or modules) belonging to the modular partition \mathcal{P} .

Definition 7 With a modular partition $\mathcal{P} = \{M_1, \dots, M_k\}$ of a graph $G = (V, E)$, we associated a quotient graph $G_{/\mathcal{P}}$, whose vertices are in one-to-one correspondence with the parts of \mathcal{P} . Two vertices v_i and v_j of $G_{/\mathcal{P}}$ are adjacent if and only if the corresponding modules M_i and M_j are adjacent in G .

Let us remark that the quotient graph $G_{/\mathcal{P}}$ with $\mathcal{P} = \{M_1, \dots, M_k\}$ is isomorphic to any subgraph induced by a set $V' \subseteq V$ such that $\forall i \in [1, k]$, $|M_i \cap V'| = 1$. The *representative graph* of a module M is the quotient graph $G[M]_{/\mathcal{P}}$ where \mathcal{P} is the maximal modular partition of $G[M]$: it is thereby the subgraph induced by a set containing a unique – *representative* – ver-

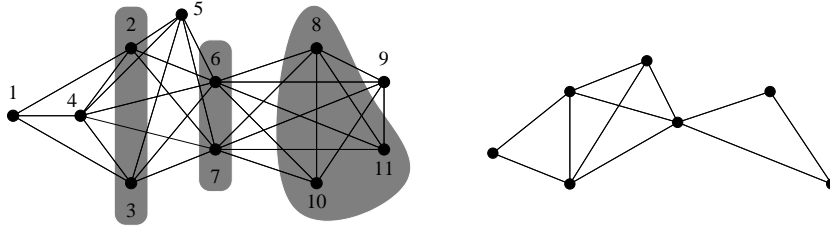


Fig. 2. On the left, the grey sets are modules of the graph G . $\mathcal{Q} = \{\{1\}, \{2, 3\}, \{4\}, \{5\}, \{6, 7\}, \{9\}, \{8, 10, 11\}\}$ is a modular partition of G . On the right, is represented the quotient graph G/\mathcal{Q} . The maximal modular partition of G is $\mathcal{P} = \{\{1\}, \{2, 3, 4\}, \{5\}, \{6, 7\}, \{8, 9, 10, 11\}\}$ and its quotient graph is represented in Figure 3 (aside the top node of the tree).

text per maximal strong module of $G[M]$. See Figure 2. By extension, for a module M , we denote by G/M the graph quotiented by the modular partition $\{M\} \cup \{\{x\} \mid x \notin M\}$.

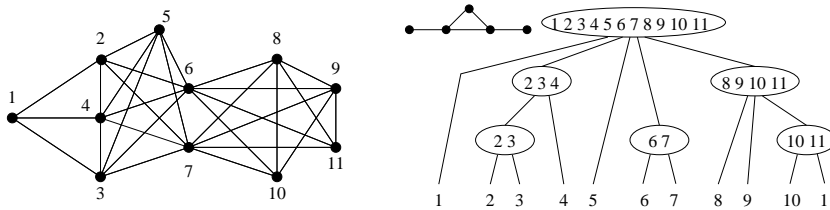


Fig. 3. The inclusion tree $MD(G)$ of the strong modules of G . The representative graph associated to the root is G/\mathcal{P} with $\mathcal{P} = \{\{1\}, \{2, 3, 4\}, \{5\}, \{6, 7\}, \{8, 9, 10, 11\}\}$.

The inclusion tree of the strong modules of G , denoted $MD(G)$, entirely represents the graph if the representative graph of each strong module is attached to each of its nodes (see Figure 3). Indeed any adjacency of G can be retrieved from $MD(G)$. Let x and y be two vertices of G and let G_N be the representative graph of node N , their least common ancestor. Then x and y are adjacent in G if and only if their representative vertices in G_N are adjacent.

Before we state the *modular decomposition theorem* (Theorem 2), let us present two more properties of modular partitions and quotient graphs which are central to efficient modular decomposition algorithms (see Section 5).

Lemma 2 [Möh85b] *Let \mathcal{P} be a modular partition of a graph $G = (V, E)$. Then $\mathcal{X} \subseteq \mathcal{P}$ is a module of G/\mathcal{P} iff $\bigcup_{M \in \mathcal{X}} M$ is a module of G .*

This lemma can be strengthened in order to observe the same correspondance between the strong modules of G and those of G/\mathcal{P} .

Lemma 3 *Let \mathcal{P} be a modular partition of a graph $G = (V, E)$. Then $\mathcal{X} \subset \mathcal{P}$ is a non-trivial strong module of G/\mathcal{P} iff $\bigcup_{M \in \mathcal{X}} M$ is a non-trivial strong module of G .*

Theorem 2 (Modular decomposition theorem) [Gal67, CHM81]

For any graph $G = (V, E)$, one of the following three conditions is satisfied:

- (1) G is not connected;
- (2) \overline{G} is not connected;
- (3) G and \overline{G} are connected and the quotient graph $G_{/\mathcal{P}}$, with \mathcal{P} the maximal modular partition of G , is a prime graph.

What does the modular decomposition theorem say is twofold. First, the quotient graphs associated with the node of the inclusion tree $MD(G)$ of the strong modules are of three types: an *independent set* if G is not connected (the node is labelled *parallel*); a *clique* (complete graph) if \overline{G} is not connected (the node is labelled *series*); a prime graph otherwise. It also follows that $MD(G)$ is unique and does not contain two consecutive series nodes nor two consecutive parallel nodes. Parallel and series nodes of $MD(G)$ are often called *degenerate* nodes.

The tree $MD(G)$ is called the *modular decomposition tree*. Theorem 2 yields a natural polynomial time recursive algorithm to compute $MD(G)$: 1) compute the maximal modular partition \mathcal{P} of G ; 2) label the root node according to the parallel, series or prime type of G ; 3) for each module M of \mathcal{P} , compute $MD(G[M])$ and attach it to the root node. Computing the maximal modular partition seems to be a central problem in the computation of $MD(G)$. It can be avoided if a non-trivial module M is identified. By Lemma 2 and Lemma 3, it suffices to recursively compute $MD(G[M])$ and $MD(G_{/M})$, and then to paste $MD(G[M])$ on the leaf of $MD(G_{/M})$ corresponding to the representative vertex of M .

Before we present some structural properties of prime and totally decomposable graphs, let us introduce some notations and briefly discuss the composition view of the theory of modules in graphs.

Notation 3 The strong modules corresponding to nodes p and q of $MD(G)$ will be respectively denoted $M(p)$ (or P) and $M(q)$ (or Q).

The minimal strong module containing two vertices x and y will be denoted $m(x, y)$, while $M(x, \bar{y})$ will design the maximal strong module containing x but not y .

The *substitution* operation is the reverse of the quotient operation. It consists of replacing a vertex x of G by a graph $H = (V', E')$ while preserving the neighbourhood. The resulting graph is:

$$G_{x \rightarrow H} = ((V \setminus \{x\}) \cup V', (E \setminus \{xy \in E\}) \cup E' \cup \{yz : xy \in E \text{ et } z \in V'\})$$

The *parallel composition* or *disjoint union* of k connected graphs G_1, \dots, G_k defines a graph whose connected components are the graphs G_1, \dots, G_k . This

composition operation is usually denoted $G_1 \oplus \dots \oplus G_k$.

The *series composition* of k co-connected graphs G_1, \dots, G_k defines a graph whose co-connected components are the graphs G_1, \dots, G_k (for any pair x, y of vertices belonging to different graphs G_i and G_j , the edge xy has been added). The series composition is generally denoted $G_1 \otimes \dots \otimes G_k$.

These three operations are classical graph operations that have been used in various contexts among which the clique-width theory [CER93].

2.4 Prime graphs

The structure of prime graphs has been extensively studied (e.g. see [CI98]). For example, it is easy to check that the smallest prime graph is the P_4 , the path on 4 vertices (see Figure 4). As witnessed by the following result, P_4 's play an important role in the structure of prime graphs.

Lemma 4 [CI98] *Let G be a prime graph. Then any vertex, but at most one, is contained in an induced P_4 . A vertex not contained in any P_4 is called the "nose of the bull" (see Figure 4).*




Fig. 4. The vertices a, b, c, d form a P_4 whose *extremities* are a and d , and midpoints b and c . The graph on the right is the *bull* the "nose" of which is vertex x .

Jamison and Olariu proposed an extension of Theorem 2 by considering the structure of prime graphs [JO95]. A subset C of vertices of a graph $G = (V, E)$ is *P-connected* if for any bipartition $\{A, B\}$ of C , there is a P_4 intersecting both A and B . For example the bull is not *P-connected* (consider the vertex partition $\{\{x\}, \{a, b, c, d\}\}$). A *P-connected* component is a maximal *P-connected* set of vertices. The set of *P-connected* components defines a partition of the vertices. A *P-connected* component H is *separable* if there is a bipartition (H_1, H_2) of H such that for any P_4 intersecting H_1 and H_2 , the extremities are in H_1 and the mid-vertices in H_2 .

Theorem 4 [JO95] *Let $G = (V, E)$ be a connected graph such that \overline{G} is connected. then G is either *P-connected* or there exists a unique *P-connected* component H which is separable in (H_1, H_2) such that any vertex $x \notin H$, $H_1 \subseteq N(x)$ and $H_2 \cap N(x) = \emptyset$.*

A hierarchy of graph families have been proposed based on the above Theorem by restricting the number of P_4 's in small subgraphs (or equivalently by

restricting the structure of prime graphs). For example, P_4 -sparse graphs are defined as the graphs for which there is at most one P_4 in any induced subgraph on 5 vertices [JO92a, JO92b]. Let us also mention the the P_4 -reducible graphs [JO95]. See [BLS99] for a complete presentation of these graph families.

2.5 Totally decomposable graphs

A graph is *totally decomposable* if any induced subgraph of size at least 4 has a non-trivial module. It follows from Theorem 1 that any node of modular decomposition tree $MD(G)$ of a totally decomposable graph G is degenerate.

The family \mathcal{F} of totally decomposable graphs is natural and arose in many different contexts (see [Sum73, CLSB81, CPS85] for references) even recently (see [BBCP04, BRV07]) as any graph of \mathcal{F} can be obtained by a sequence of disjoint and series compositions starting from single vertex graph belongs to \mathcal{F} . Let us remark that if G is totally decomposable then also is its complement. The family of totally decomposable is also known as the *cographs* for *complement reducible graphs* [CLSB81, Sum73]. From the following characterization theorem, the cograph family is *hereditary* (any induced subgraph of a cograph is a cograph).

Theorem 5 [Sum73] *The cographs are exactly the P_4 -free graphs.*

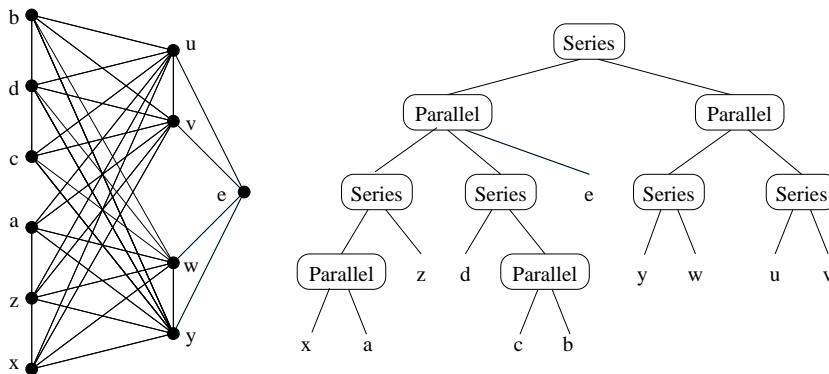


Fig. 5. A cograph and its modular decomposition tree (also called cotree).

The first linear time recognition algorithm was proposed in 1985 by Corneil, Perl and Stewart [CPS85]. It incrementally construct the modular decomposition tree, also called *cotree* for the cographs, as long as the graph induced by the processed vertices is a cograph. Even if alternative recognition algorithms have recently been proposed [Dah95, HP05, BCHP03], the seminal algorithm of [CPS85] is a corner stone in the algorithmic of the modular decomposition and turns out to have a large impact even for other decomposition technics (e.g. for the split decomposition [GP07]).

The following lemma states classical properties of cographs whose proofs (left to the reader) are good exercises to understand the structure of cographs.

Lemma 5 *Let x, y and v be vertices of a cograph $G = (V, E)$.*

- (1) *If $xv \in E$, $yv \notin E$ and $xy \in E$, then $m(v, y) \subseteq M(v, \bar{x})$*
- (2) *If $xv \in E$, $yv \in E$ and $xy \notin E$, then $M(v, \bar{x}) = M(v, \bar{y})$ and $m(v, x) = m(v, y)$*

2.6 Bibliographic notes

The seminal paper on modular decomposition of graphs is probably Gallai's one [Gal67] on transitive orientation. Up to our knowledge, the only survey paper is due to Möhring and Radermacher [MR84]. More recently, Ehrenfeucht, Harju and Rozenberg [EHR99] published a book on the decomposition of 2-structures (a generalization of graphs) which presents the modular decomposition in a more general framework. In its PhD thesis [BX08], Bui Xuan proposes a survey as well as original results on the representation of set families. Many graph families are well-structured with respect to the modular decomposition, *e.g.* comparability graphs, permutation graphs, cographs... For these aspects, the reader should refer to the books of Golubic [Gol80] and more recently [BLS99, Spi03]. The algorithmic aspect is particularly developed in [Gol80, Spi03].

3 Cographs algorithms as an appetizer

In 1985, Corneil, Pearl and Stewart [CPS85] proposed the first linear time algorithm for the cograph recognition problem. If the input graph is a cograph, this vertex incremental algorithm builds the cotree by adding the vertices one by one in an arbitrary order. In the late 80's, Müller and Spinrad generalized it to the first quadratic modular decomposition algorithm [MS89]. Their algorithm is also incremental. But for the sake of adjacency tests, the whole graph has to be known when the algorithm starts. Even though that most of linear time modular decomposition algorithms for the are based on different approaches (*e.g.* partition refinement, see Section 4), Corneil *et al*'s algorithm should be considered as one of the most important step towards the linear time computation of the modular decomposition tree of a graph. Also, as we shall discuss in Section 3.2, it has recently been generalized to fully dynamic algorithms for more general graph classes: the problem consists in updating the modular decomposition tree of a graph when a vertex or an edge is removed or added and the goal is to obtain a better complexity than computing the modular decomposition from scratch.

3.1 Cograph recognition algorithms

Consider the following subproblem: given a cograph $G = (V, E)$ together with its cotree $MD(G)$, a vertex x and a subset of vertices $S \subseteq V$, test whether the graph $G + (x, S) = (V \cup \{x\}, E \cup \{xy \mid y \in S\})$ is a cograph and if so output the cotree $MD(G + x)$. Corneil *et al*'s [CPS85] showed that whether $G + x$ is a cograph or not can be characterized by a labelling of the nodes of the cotree $MD(G)$. A node p receives the label: *empty*, if the corresponding module $M(p)$ does not intersect S ; *adjacent* if $M(p) \subseteq S$; and *mixed* otherwise. Remark that by definition any child of a node labelled adjacent (resp. empty) is also labelled adjacent (resp. empty).

Lemma 6 [CPS85] *Let G be a cograph, x a vertex of V and $S \subseteq V$. The graph $G + (x, S)$ is a cograph iff*

- (1) *either none of the nodes of the cotree $MD(G)$ is mixed;*
- (2) *or the set of mixed nodes induces a path π from the root of $MD(G)$ to some node p and*
 - (a) *the children of the series nodes of π different than p are all adjacent;*
 - (b) *the children of the parallel nodes of π different than p are all empty.*

The main idea expressed by the conditions of Lemma 6 is that the modifications of the cotree implied by the insertion of vertex x are localized in the subtree $T(p)$ root at node p . Indeed any module disjoint from $M(p)$ is not affected by x 's insertion (the corresponding nodes are labelled empty or adjacent). In a sense, node p should be considered as the insertion node. The cotree updates only depends on node p (*e.g.* whether it is mixed or adjacent). An example is depicted in Figure 6.

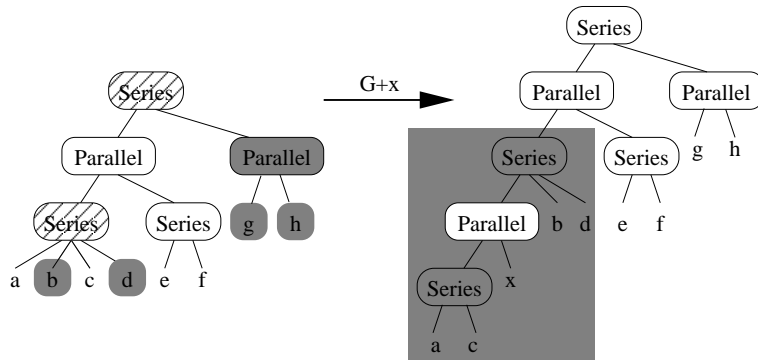


Fig. 6. Insertion of the vertex x adjacent to $S = \{b, d, g, h\}$. Grey nodes are the adjacent labelled node and dashed nodes are the mixed nodes.

The algorithm first labels the cotree in a bottom-up manner. The leaves corresponding to vertices of S are labelled adjacent. A node labelled adjacent forwards a partial mark to its father. When a node have received a mark from

each of its children, it is labelled adjacent. At the end of this process the empty node have never been searched, while the partially marked nodes corresponds, if $G + x$ is a cograph, to the parallel node of the path π . It is not difficult to see that the number of the marked node is linear in the size of S meaning that the labelling process runs in time $O(|S|)$. Testing the condition of the above lemma can be done within the same complexity as well.

Theorem 6 [CPS85] *The family of cographs can be recognized in linear time.*

Let us now turn to the edge modification problem which consists in updating the cotree of a cograph G under an edge insertion or deletion. Since the cotree of a cograph can be obtained from the cotree of its complement by flipping the parallel and the series nodes, deleting or inserting an edge in a cograph are equivalent problems.

Lemma 7 [SS04] *Let x and y be two non-adjacent vertices of a cograph $G = (V, E)$. Then $G + xy = (V, E \cup \{xy\})$ is a cograph iff x is a child of $m(x, y)$ and $M(y, \bar{x}) \subseteq N(x)$.*

Let us sketch the argument proof. As $xy \notin E$, the module $m(x, y)$ is represented by a parallel node. Assume the conditions of Lemma 7 do not hold. Then the path in the cotree from $m(x, y)$ to x (resp. y) contains a series nodes p_x (resp. p_y) which is the least common ancestor of x (resp. y) and some leaf u_x (resp. u_y). Then the vertices $\{u_x, x, y, u_y\}$ induces a P_4 in the graph $G + xy$.

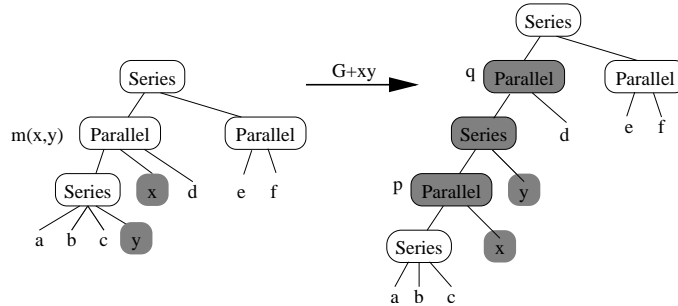


Fig. 7. Update of the cotree to insert the edge xy in a cograph. The node $m(x, y)$ is splitted into two parallel nodes, say p and q , one being the father of x , the another the father of the other children of $m(x, y)$. Then leaf y is extracted from the cotree and attached to a new series node inserted between nodes p and q .

It follows from Lemma 7 that as long as the modified graph remains a cograph, the modifications in the cotree are local and can be done in constant time. From results presented in this section, we obtain that:

Theorem 7 [SS04, CPS85] *There exists a fully-dynamic algorithm maintaining the modular decomposition tree of a cograph which runs in time $O(d)$ per modification (edge or vertex insertion and deletion), where d is the number involved in the modification.*

3.2 Bibliographic notes

Concerning the cograph recognition problem, new algorithms also appeared recently. Habib and Paul [HP05] proposed a partition refinement based algorithm (see Section 4) and Bretscher *et al* [BCHP08] discovered a simple Lexicographic Breadth First Search [RTL76] based algorithm.

Aside the two cograph algorithmic results presented above, fully-dynamic algorithms have recently been proposed to maintain a representation based on the modular decomposition tree under vertex and edge modifications for various graph classes: permutation graphs [CP06], interval graphs [Cre09, Iba09]. . . The fully-dynamic representation problem has also been solved for other families of graphs, *e.g.* proper interval graphs [HSS01], using other decomposition schemes.

Beside, Corneil *et al*'s algorithm has been generalized to the split decomposition [CE80] to obtain an optimal fully dynamic algorithm for the distance hereditary graphs recognition problem [GP07]. More recently by the same technique, Gioan *et al.* derived an almost linear time split decomposition algorithm [GPTC09a] and the first subquadratic circle graph recognition algorithm [GPTC09b].

4 Partition refinement

Partition refinement, as an algorithmic technique, has been used in a number of problems, the first of which is probably the deterministic automata minimization [Hop71]. Paige and Tarjan [PT87] wrote a synthesis paper on this technique. Since then, the number of problems solved by partition refinement keeps increasing: interval graph recognition [HPV99] and completion [RST08], transitive orientation [HMPV00] are example among others. As we will see, this technique turns out to be a powerful and simple algorithmic paradigm that plays an important role in the context of modular decomposition.

We first present the data-structure and the elementary operation, namely the *refine* operation, of the partition refinement technique. Then, we illustrate this technique with an algorithm that computes a modular partition of a graph. Let us mention that this algorithm really follows the lines of Hopcroft's deterministic automaton minimization algorithm [Hop71].

4.1 Data-structures and algorithmic scheme

Let \mathcal{P} and \mathcal{P}' be two partitions of the same set V . The partition \mathcal{P} is *smaller* than \mathcal{P}' , denoted $\mathcal{P} \triangleleft \mathcal{P}'$, if $\mathcal{P} \neq \mathcal{P}'$ and any part of \mathcal{P} is a subset of some part of \mathcal{P}' . The partition \mathcal{P} is *stable* with respect to a set S if none of the parts of \mathcal{P} overlaps S .

Partition refinement consists of repeating, as long as needed, the operation described in Algorithm 1. The initial partition and the sequence of *pivot sets* used in the successive refinement steps have a large impact on the whole complexity of the algorithm. Partitioning the vertex set of a graph with respect to the neighbourhood of some vertex is a common operation in graph algorithms. Indeed in our examples, the pivot sets correspond to some neighbourhoods.

Algorithm 1: $Refine(\mathcal{P}, S)$

Input: A partition \mathcal{P} of a set V and a subset $S \subseteq V$, called *pivot set*

Output: The coarsest partition refining \mathcal{P} and stable for S

begin

foreach part $\mathcal{X} \in \mathcal{P}$ **do**

if $\mathcal{X} \cap S \neq \emptyset$ and $\mathcal{X} \cap S \neq \mathcal{X}$ **then** replace \mathcal{X} by $\mathcal{X} \cap S$ and $\mathcal{X} \setminus S$;

end

Let us briefly describe the data-structure (see Figure 8). The elements of the set V to be partitioned are stored in a doubly linked list. Each element of V is assigned a pointer towards the part it belongs to. The elements of a part \mathcal{X} remains consecutive in the doubly linked list (they form an interval). So that each part maintain a pointer towards its first and its last element in the list.

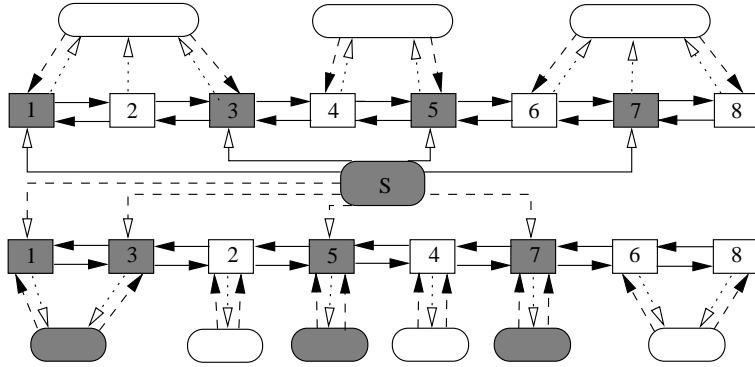


Fig. 8. $\mathcal{P}' = Refine(\mathcal{P}, S)$.

Notation 8 *The data-structure implicitly represent an ordered partition: the parts are totally ordered. Depending of the application, this aspect may or may not be important. In order to distinguish the two different cases, an ordered*

partition will be denoted $\mathcal{P} = [\mathcal{X}_1, \dots, \mathcal{X}_k]$ while a non-ordered partition will be denoted $\mathcal{P} = \{\mathcal{X}_1, \dots, \mathcal{X}_k\}$.

Lemma 8 *The time complexity of the operation $\text{Refine}(\mathcal{P}, S)$ is $O(|S|)$.*

We conclude this brief introduction by a few remarks. Refining a partition by a subset S or its complement $\bar{S} = V \setminus S$ are equivalent operations: $\text{Refine}(\mathcal{P}, S) = \text{Refine}(\mathcal{P}, V \setminus S)$. It is thereby possible to deal with the complement of the input graph without explicitly storing its edge set. Partition refinement is usually used either to compute a total ordering of the vertices (*e.g.* Lex-BFS) or the equivalence classes of some equivalence relations (*e.g.* maximal set of twin vertices). McConnell and Spinrad [MS00] showed how to augment the data-structure in order to extract within the same complexity, at each refinement step, the edges incident to vertices belonging to different parts. This operation is useful to efficiently compute the quotient graph associated to a modular partition. For a more detailed presentation of partition refinement refer to [PT87, HPV98, HPV99, HMPV00].

4.2 Hopcroft's rule and computation of a modular partition

Partition refinement is the right tool to compute a modular partition, an important subproblem towards efficient modular decomposition algorithms. In this section, we focus on the problem of computing the *coarsest modular partition* (see Definition 8) of a given vertex partition. The algorithm we present runs in time $O(n + m \log n)$ and is based on the *Hopcroft's rule* which is used in various simple quasi-linear time modular decomposition algorithms.

Definition 8 *Let \mathcal{P} be a partition of the vertices of a graph $G = (V, E)$. The coarsest modular partition of G with respect to \mathcal{P} is the largest modular partition \mathcal{Q} such that $\mathcal{Q} \triangleleft \mathcal{P}$.*

The main idea of the algorithm is the following: as long as there is a part \mathcal{X} which is not uniform for some vertex $x \notin \mathcal{X}$, the current partition \mathcal{P} is refined with the neighbourhood $N(x)$. When the algorithm ends, all the parts are modules. Finding, at each step, a vertex x whose neighbourhood strictly refines the partition \mathcal{P} , is the usual barrier to linear time complexity. However, using the so-called *Hopcroft's rule*, one get a fairly simple solution that uses the neighbourhood of each vertex at most $\log n$ times.

Lemma 9 *Let \mathcal{P} be a partition of the vertices of a graph $G = (V, E)$ and x be a vertex of some part \mathcal{X} . If \mathcal{P} is stable with respect to $N(y)$, $\forall y \notin \mathcal{X}$, then \mathcal{X} is a module of G and the partition $\mathcal{Q} = \text{Refine}(\mathcal{P}, N(x))$ is stable with respect to $N(x')$, $\forall x' \in \mathcal{X}$.*

The above lemma (which is a direct consequence of the definition of module) shows that using as pivots the vertices of all the parts of \mathcal{P} but one, say \mathcal{Z} , plus one vertex z of \mathcal{Z} is enough. For complexity issues, the avoided part \mathcal{Z} has to be chosen as the largest part of \mathcal{P} . Similarly, once a part \mathcal{X} has been split, the process continues recursively on the subgraph induced by \mathcal{X} and the resulting largest subpart can be avoided (meaning that only one of its vertices has to be used as pivot). This "avoid-the-largest" technique is known as the *Hopcroft's rule* and has been first proposed in the deterministic automata minimization algorithm [Hop71].

Algorithm 2: Modular Partition

Input: A partition \mathcal{P} of the vertex set V of a graph G

Output: The coarsest modular partition \mathcal{Q} smaller than \mathcal{P}

begin

 Let \mathcal{Z} be the largest part of \mathcal{P} ;

$\mathcal{Q} \leftarrow \mathcal{P}$; $K \leftarrow \{\mathcal{Z}\}$; $L \leftarrow \{\mathcal{X} \mid \mathcal{X} \neq \mathcal{Z}, \mathcal{X} \in \mathcal{P}\}$;

while $L \cup K \neq \emptyset$ **do**

if there exists $\mathcal{X} \in L$ **then** $S \leftarrow \mathcal{X}$ and $L \leftarrow L \setminus \{\mathcal{X}\}$;

else

1 Let \mathcal{X} be the first part K and $x \in \mathcal{X}$;

$S \leftarrow \{x\}$ and $K \leftarrow K \setminus \{\mathcal{X}\}$;

foreach vertex $x \in S$ **do**

foreach part $\mathcal{Y} \neq \mathcal{X}$ such that $N(x) \perp \mathcal{Y}$ **do**

2 Replace in \mathcal{Q} , \mathcal{Y} by $\mathcal{Y}_1 = \mathcal{Y} \cap N(x)$ and $\mathcal{Y}_2 = \mathcal{Y} \setminus N(x)$;

 Let \mathcal{Y}_{min} (resp. \mathcal{Y}_{max}) be the smallest part (resp. largest) among \mathcal{Y}_1 and \mathcal{Y}_2 ;

if $\mathcal{Y} \in L$ **then** $L \leftarrow L \cup \{\mathcal{Y}_{min}, \mathcal{Y}_{max}\} \setminus \{\mathcal{Y}\}$;

else

$L \leftarrow L \cup \{\mathcal{Y}_{min}\}$;

if $\mathcal{Y} \in K$ **then** Replace \mathcal{Y} by \mathcal{Y}_{max} in K ;

else Add \mathcal{Y}_{max} at the end of K ;

end

To implement this rule, the parts are stored in two disjoint lists K and L . The neighbourhoods of all the vertices of parts belonging to L will be used to refine the partition. For the parts belonging to K , only the neighbourhood of one selected vertex is used. The list K works in a FIFO manner, which guarantees that the first part of the list, when extracted, is a module.

Theorem 9 *Let \mathcal{P} be a partition of the vertices of a graph $G = (V, E)$. Algorithm 2 computes the coarsest modular partition for G and \mathcal{P} in time $O(n + m \log n)$.*

The correctness of the algorithm follows from the next three invariant properties. The first invariant shows that a module contains in some part of the given

partition cannot be split, while the third one guarantees that the algorithm outputs a modular partition.

- (1) *If M is a module of G contained in a part $\mathcal{X} \in \mathcal{P}$, then there exists a part \mathcal{Y} of the current partition containing M .*
- (2) *If $L = \emptyset$, then the first part \mathcal{Y} of K is a module.*
- (3) *If the current partition contains a part \mathcal{X} that is not a module, then there exists $\mathcal{Y} \in L \cup K$ different from \mathcal{X} and containing a splitter y for \mathcal{X} .*

4.3 Bibliographic notes

As already mentioned, the use of partition refinement technique dates to 1971 for the deterministic automata minimization problem [Hop71]. In 1987, Paige and Tarjan used again this technic to solve three different problems: functional partition, coarsest relational partition problems and doubly lexicographic ordering of a boolean matrix. In the late 90's, it has been used more systematically in the context of modular decomposition and transitive orientation yielding $O(n+m \log n)$ practical and simple algorithms (see e.g. [MS00, HMPV00]).

AUTRES EXEMPLES D'APPLICATIONS ???

5 Recursive computation of the modular decomposition tree

In 1994, Ehrenfeucht, Gabow, McConnell and Sullivan [EGMS94] proposed a quadratic algorithm for the modular decomposition². The principle of this algorithm, which we will call the *skeleton algorithm*, is the basis of a large number of the known subquadratic algorithms proposed in the late 90's (see e.g. [MS00, DGM01]), which could abusively be considered as a series of different implementations of the skeleton algorithm. The complexity of these implementations are respectively $O(n + m \cdot \alpha(n, m))$ or $O(n + m)$ [DGM01], and finally $O(n + m \log n)$ [MS00]. We describe the principle of the skeleton algorithm without considering the complexity issues. We then discuss the differences in the time complexity of the known algorithms.

5.1 The skeleton algorithm

Let us first mention that the skeleton algorithm computes a *non-reduced* form of the modular decomposition tree $MD(G)$: the resulting tree may contain

² This algorithm is designed for 2-structures, a classical generalization of graphs.

some series (or parallel) node child of a series (or parallel) node. All the algorithms we describe in this section will do so. It does not impact the complexity issues as a single search of the tree is enough reduce it in time $O(n)$. In the following, we will abusively denote $MD(G)$ the (non-reduced) decomposition tree returned by these algorithms.

The main idea developed by Ehrenfeucht et al. [EGMS94] is to first compute a "spine" of the modular decomposition tree $MD(G)$, then to recursively compute the modular decomposition trees of some induced subgraphs which are eventually padded to the spine. More formally:

Definition 9 Let v be an arbitrary vertex of a graph $G = (V, E)$. The v -modular partition is the following modular partition:

$$\mathcal{M}(G, v) = \{v\} \cup \{M \mid M \text{ is a maximal module not containing } v\}$$

We define $spine(G, v)$ as the modular decomposition tree $MD(G/\mathcal{M}(G, v))$.

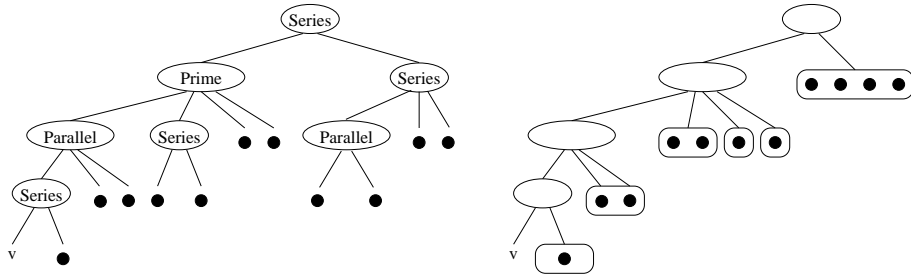


Fig. 9. On the left, a modular decomposition tree $MD(G)$ and on the right, the modular partition $\mathcal{M}(G, v)$ with the corresponding spine between v and the root of $MD(G)$.

Algorithm 3: Ehrenfeucht et al. [EGMS94]

Input: An arbitrary vertex v of $G = (V, E)$, $spine(G, v)$ and $\{MD(G[X]) \mid X \in \mathcal{M}(G, v)\}$

Output: The modular decomposition tree $MD(G) = (\mathcal{M}, F)$

```

begin
  foreach leaf  $X$  of  $T$  do
    Let  $T_X = MD(G[X])$  and  $p(X)$  be  $X$ 's father  $X$  in  $T$ ;
    Replace  $X$  by  $T_X$  in  $T$ ;
    if the root  $r(T_X)$  and  $p(X)$  are both parallel or series then
       $\lfloor$  Remove  $r(T_X)$  and connect the children of  $r(T_X)$  to  $p(X)$ 
    end
  end

```

Let us notice that any degenerate strong module (series or parallel) containing v will be represented in $spine(G, v)$ by a binary node. The purpose of test of Line 1 in Algorithm 3 is to correctly fix those binary nodes. The correctness of Algorithm 3 is a consequence of the following properties:

Lemma 10 [EGMS94] *Let v be a vertex of a graph $G = (V, E)$ and $\mathcal{M}(G, v)$ be the associated modular partition. Then:*

- (1) *Any non-trivial module of $G_{/\mathcal{M}(G,v)}$ contains v ;*
- (2) *A set $\mathcal{X} \subset \mathcal{M}(G, v)$ is a non-trivial strong module of $G_{/\mathcal{M}(G,v)}$ iff $\bigcup_{M \in \mathcal{X}} M$ is an ancestor of v in $MD(G)$;*
- (3) *Any module not containing v is a subset of a part $M \in \mathcal{M}(G, v)$.*

Using Algorithm 2, $\mathcal{M}(G, v)$ can be computed in time $O(n + m \log n)$. Computing $spine(G, v)$ is a more difficult and technical task. It is indeed the main bottleneck in the complexity of the skeleton algorithm. The solution we present hereafter has been proposed in [EGMS94] and yields quadratic running time. Later on, Dahlhaus et al. [DGM01] improved this step and obtained a subquadratic running time (see discussion of Section 5.3).

5.2 Computation of $spine(G, v)$.

Definition 10 *A graph $G = (V, E)$ is nested if there exists a vertex $v \in V$ which is contained in all the non-trivial modules of G . Such a vertex is called an inner vertex of G .*

As a direct consequence of Lemma 10, first property, the quotient graph $G_{/\mathcal{M}(G,v)}$ is a nested graph with inner vertex v .

In order to compute, the modules of $G_{/\mathcal{M}(G,v)}$ and $spine(G, v)$, Ehrenfeucht et al. [EGMS94] introduced an auxiliary *forcing digraph* the arc set of which guarantees the existence of a directed path from any vertex u to any vertex $w \in m(u, v)$, the smallest module containing u and v . As v belongs to all the modules of $G_{/\mathcal{M}(G,v)}$, a simple search on the forcing graph will suffice to compute $spine(G, v)$.

Definition 11³ *Let v be an arbitrary vertex of a graph $G = (V, E)$. The forcing graph $\mathcal{F}(G, v)$ is a directed graph whose vertex set is $V \setminus \{v\}$. The arc \overrightarrow{xy} exists if y is a splitter for $\{x, v\}$.*

In other words, if \overrightarrow{xy} exists then y belongs to any module containing v and x .

Lemma 11 [EGMS94] *If X is the set of vertices that can be reached from vertex x in the forcing graph $\mathcal{F}(G, v)$, then $\{v\} \cup X = m(v, x)$.*

In the following we will only consider the graph $G_{/\mathcal{M}(G,v)}$ and we will denote the associated forcing graph by $\mathcal{F}(\mathcal{M}(G, v))$ (whose vertices are $\mathcal{M}(G, v) \setminus \{v\}$).

³ The definition we propose slightly differs from the original one of [EGMS94]. This modification will ease the relationship with the results of [DGM01].

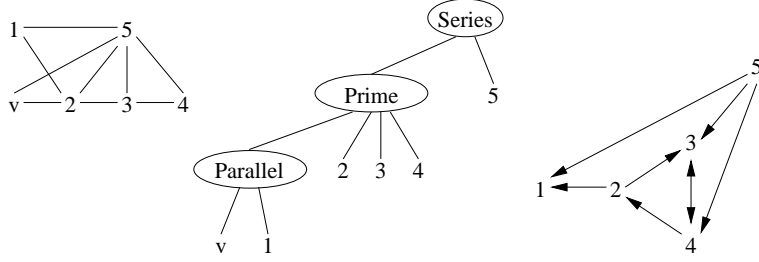


Fig. 10. A nested graph $G = (V, E)$ together with its modular decomposition tree $MD(G)$ and on its right the forcing graph $\mathcal{F}(G, v)$. The strongly connected components of $\mathcal{F}(G, v)$ are $\{1\}$, $\{2, 3, 4\}$, $\{5\}$. Any module of G containing 3 and V also contains $\{1, 2, 4\}$, the vertices that can be reached from vertex 3 in $\mathcal{F}(G, v)$.

As $G_{/\mathcal{M}(G,v)}$, applying Lemma 11 to $\mathcal{F}(\mathcal{M}(G, v))$, we obtain the following property.

Corollary 1 [EGMS94] *Let M_x be the module of $\mathcal{M}(G, v)$ containing the vertex x . If \mathcal{X} is the set of modules that can be reached from M_x in $\mathcal{F}(\mathcal{M}(G, v))$, then $\bigcup_{M \in \mathcal{X}} M = m(v, x)$.*

We now consider the *block graph* $\mathcal{B}(G, v)$ of $\mathcal{F}(\mathcal{M}(G, v))$ (see [CLR90]) whose vertices are the strongly connected components of $\mathcal{F}(\mathcal{M}(G, v))$, also called the *blocks* of (G, v) . An arc of $\mathcal{B}(G, v)$ between the block B and B' exists if the vertices of B' can be reached in $\mathcal{F}(\mathcal{M}(G, v))$ from the vertices of B .

Lemma 12 [EGMS94] *The transitive reduction of the block graph $\mathcal{B}(G, v)$ is a chain.*

A set of vertices of a digraph is a *sink* if it has no out-neighbour. By Lemma 12, any sink set of $\mathcal{F}(\mathcal{M}(G, v))$ is the union of consecutive blocks containing the last one in the transitive reduction of $\mathcal{B}(G, v)$. Each sink set corresponds to a module of $G_{/\mathcal{M}(G,v)}$.

Corollary 2 [EGMS94] *Let v be a vertex of a graph $G = (V, E)$. A set M of vertices containing v is a module of $G_{/\mathcal{M}(G,v)}$ iff M is the union of $\{v\}$ and the modules of $\mathcal{M}(G, v)$ belonging to a sink set X of $\mathcal{B}(G, v)$.*

Thereby the forcing graph $\mathcal{F}(\mathcal{M}(G, v))$ describes the modules of $G_{/\mathcal{M}(G,v)}$ and the block graph $\mathcal{B}(G, v)$ allows us to compute $spine(G, v)$. Finally, $MD(G)$ is obtained recursively by following the lines of Lemma 10.

5.3 Time complexity

Rather than detailing the complexity analysis, we point out the differences between the original skeleton algorithm presented in [EGMS94] and its later

versions improved in [DGM01]. The interested reader should access the original papers for details. As already mentioned, a quadratic time complexity analysis is proposed in [EGMS94]. The main bottlenecks are the computation of the partition $\mathcal{M}(G, v)$ and the construction of $MD(G/\mathcal{M}(G, v))$.

Two new versions of the skeleton algorithm proposed by Dahlhaus, Gustedt and McConnell [DGM01], respectively run in $O(n + m \cdot \alpha(n, m))$ time and in linear time. To improve the time complexity, the authors of [DGM01] borrowed from [Dah95] the idea to first recursively compute the modular decomposition trees of the subgraphs induced by $N(v)$ and by $\bar{N}(v)$. It follows from the next Lemma, that $\mathcal{M}(G, v)$ is easy to retrieve from those trees.

Lemma 13 *If \mathcal{X} is a module of $\mathcal{M}(G, v)$, then \mathcal{X} is either a module of $G[N(v)]$ or a module of $G[\bar{N}(v)]$.*

As in [EGMS94], the technique used to compute $spine(G, v)$ relies on a forcing digraph. Remind that the vertices of $\mathcal{F}(\mathcal{M}(G, v))$ are the modules of G (indeed the modules of $\mathcal{M}(G, v)$) which turns out to be a too strong condition for time complexity issues. In [DGM01], the forcing digraph is rather defined with the help of an equivalence relation. The idea is that each equivalence class gathers vertices of $N(v)$ or of $\bar{N}(v)$ which appear in a set of sibling modules of some ancestor node of v in $MD(G)$ (or $spine(G, v)$). The partition defined by the equivalence classes is a coarser partition than $\mathcal{M}(G, v)$.

The final trick is that given $MD(G[N(v)])$ and $MD(G[\bar{N}(v)])$, the computation of $\mathcal{M}(G, v)$, $spine(G, v)$ and finally $MD(G)$ has to be done in time linear in the number of *active edges*, *i.e.* the edges incident to v and the edges linking vertices of $N(v)$ and $\bar{N}(v)$. The $\alpha(n, m)$ factor in the first version of the skeleton algorithm presented in [DGM01] is due to the use of some *union-find* data-structures required to update the current tree. A clever time complexity analysis yields linear time if a careful pre-processing step is used to fix the recursion tree.

5.4 Bibliographic notes

Let us mention that the problem of finding a simple linear time algorithm for the modular decomposition is presented in [MS00] or [Spi03] as an open problem. In its book [Spi03], Spinrad wrote p.149:

"I hope and believe that in a number of years the linear algorithm can be simplified as well"

Based on partition refinement techniques, a simplified $O(n + m \log n)$ version of the skeleton algorithm has been developed in [MS00].

6 Factoring permutation algorithm

In its PhD Thesis, Capelle [Cap97] proved that computing the modular decomposition tree of a graph and computing a *factoring permutation* (see Definition 4 and Figures 3, 12) are two equivalent tasks, as one can be retrieved from each another in linear time [CHdM02]. It follows that computing the modular decomposition of a graph can be divided into two different steps: 1) computation of a factoring permutation; 2) computation of the modular decomposition tree given the factoring permutation. The main interest of such a strategy is to obtain an algorithm that avoids the auxiliary data-structures such as *union-find* and *least common ancestor*. Moreover, in some recent applications (*e.g.* comparative genomics [UY00, BHS02, HMS09]), the given data is not the graph nor the partitive family but rather a factoring permutation. This concept turns out to be of interest by itself.

As noticed by Capelle [Cap97], this strategy was already used in few cases such as the computation of the modular decomposition tree of chordal graph [HM91] and the block tree of inheritance graphs [HHS95]. In [HPV98, HPV99], a partition refinement algorithm is proposed to compute a factoring permutation of a graph in time $O(n + m \log n)$. Restricted to cographs, the complexity can be improved down to linear time [HP05].

We will first revisit Algorithm 1 of [HPV98] and show how it can be adapted to compute a factoring permutation in time $O(n + m \log n)$. This algorithm has to be compared to the McConnell and Spinrad's implementation [MS00] of Ehrenfeucht et al.'s algorithm. The main differences are that the modular decomposition tree is never built and the relative order between the different parts of the partition is important.

There exist several linear time algorithms that given a factoring permutation of a graph compute its modular decomposition tree. A recent one is proposed in [BCdMR05, BCdMR08]. We describe the principle first one due to Capelle, Habib and de Montgolfier [CHdM02].

6.1 Computing a factoring permutation

An ordered partition $\mathcal{P} = [\mathcal{X}_1, \dots, \mathcal{X}_k]$ of a set \mathcal{E} defines a partial order on \mathcal{E} , the set of anti-chains of which is the set of parts of \mathcal{P} . In other words, we have $x_i <_{\mathcal{P}} x_j$ iff $x_i \in \mathcal{X}_i$, $x_j \in \mathcal{X}_j$ and $i < j$. Thereby refining an ordered partition could be understood as computing an extension the corresponding partial order.

We will abusively write $x <_{\mathcal{P}} M$, for $x \in \mathcal{E}$ and $M \subset \mathcal{E}$, if $x <_{\mathcal{P}} y$ for all $y \in M$. To prove the correctness of the algorithm, we need to generalize the definition of interval of permutations to ordered partitions.

Definition 12 *Let \mathcal{P} be an ordered partition of a set \mathcal{E} . A subset $S \subseteq \mathcal{E}$ is an interval of \mathcal{P} iff there are two parts $\mathcal{L} \in \mathcal{P}$ and $\mathcal{R} \in \mathcal{P}$ (not necessarily distinct) intersecting S such that for any part \mathcal{X} :*

- if $\mathcal{L} <_{\mathcal{P}} \mathcal{X} <_{\mathcal{P}} \mathcal{R}$, then $\mathcal{X} \subset S$;
- if $\mathcal{X} <_{\mathcal{P}} \mathcal{L}$ or $\mathcal{R} <_{\mathcal{P}} \mathcal{X}$, then $\mathcal{X} \cap S = \emptyset$.

To compute a factoring permutation, the main steps of the algorithm we present are: 1) computation of an ordered partition that is a modular partition $\mathcal{M}(G, v)$ such that the strong modules containing a vertex v are intervals of $\mathcal{M}(G, v)$; and 2) recursive computation of a factoring permutation of each of the subgraphs induced by a module $M \in \mathcal{M}(G, v)$.

Algorithm 4: *Factoring-permutation(G, v)*

Input: A graph $G = (V, E)$ and a vertex $v \in V$

Output: A factoring permutation of G

begin

Let $\mathcal{P} = [\overline{N}(v), \{v\}, N(v)]$ be an ordered partition;
 Apply Algorithm 2 with the following refinement rule;
 Let x be the current pivot vertex and \mathcal{Y} a part such that $N(x) \perp \mathcal{Y}$;
if $x \leq_{\mathcal{P}} v \leq_{\mathcal{P}} \mathcal{Y}$ **or** $\mathcal{Y} \leq_{\mathcal{P}} v \leq_{\mathcal{P}} x$ **then**
 | Substitute \mathcal{Y} by $[\mathcal{Y} \cap \overline{N}(x), \mathcal{Y} \cap N(x)]$;
else
 | Substitute \mathcal{Y} by $[\mathcal{Y} \cap N(x), \mathcal{Y} \cap \overline{N}(x)]$;
foreach part $\mathcal{X} \in \mathcal{M}(G, v)$, such that $|\mathcal{X}| > 1$ **do**
 | Let x be the last vertex of \mathcal{X} used as pivot;
 | $\mathcal{P}_{\mathcal{X}} \leftarrow$ Factoring-permutation($G[\mathcal{X}], x$);
 | Substitute \mathcal{X} by $\mathcal{P}_{\mathcal{X}}$;

end

Theorem 10 *Algorithm 4 compute in time $O(n + m \log n)$ a factoring permutation of a graph $G = (V, E)$.*

Proof: We already know that Algorithm 2 applied to $[\overline{N}(v), \{v\}, N(v)]$ computes $\mathcal{M}(G, v)$. By Lemma 10, any module not containing v is a subset of some module of $\mathcal{M}(G, v)$. It thereby suffices to prove that the following invariant is satisfied by Algorithm 4 (see Figure 11):

$\Pi =$ any strong module containing v is an interval of the current partition

The property Π is obviously satisfied by the initial partition $[\overline{N}(v), \{v\}, N(v)]$. Assume by induction Π holds before the current partition \mathcal{P} is refined by $N(x)$

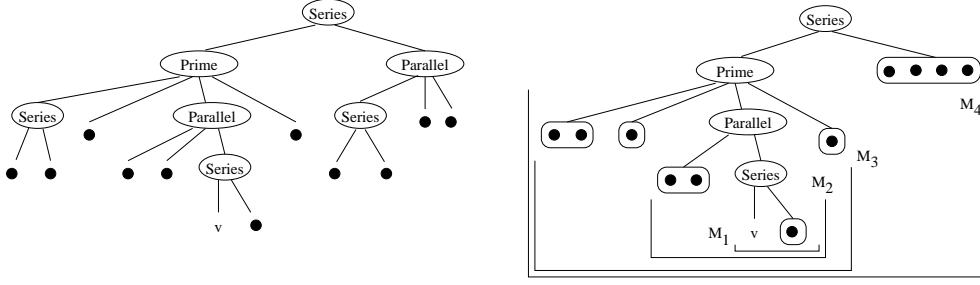


Fig. 11. Layout of the modular decomposition tree $MD(G)$ such that the neighbours of v are placed on the right of v and the non-neighbours on the left. The right tree enlightens the modules of $\mathcal{M}(G, v)$ and the strong modules M_1, M_2, M_3 and M_4 containing v . Algorithm 4 first compute the partition $\mathcal{M}(G, v)$ and then recursively solve the problem on each module of $\mathcal{M}(G, v)$

for some vertex x . Let M be a module containing v and \mathcal{X} be a part of \mathcal{P} such that $\mathcal{X} \perp N(x)$. There are two distinct cases:

- $x \notin M$: no vertex y of $\mathcal{X} \cap N(x)$ belongs M , otherwise x would be a splitter for v and y ;
- $x \in M$: if $\mathcal{X} \subset N(v)$, then any vertex $y \in \mathcal{X} \cap \overline{N}(x)$ belong to M , otherwise y would be a splitter for x and v . Similarly if $\mathcal{X} \subset \overline{N}(v)$, then any vertex $y \in \mathcal{X} \cap N(x)$ belongs to M .

It follows that $\mathcal{P}' = \text{Refine}(\mathcal{P}, N(x))$ also satisfies the invariant II. The complexity analysis is similar to the analysis of Algorithm 2. \square

6.2 The case of cographs

The natural question is how to get rid of the $\log n$ factor in the complexity of Algorithm 4. Restricting the problem to cographs (or totally decomposable graphs - see Section 2.5) gives some ideas. The reader should keep in mind that the $\log n$ factor corresponds to the number of times the neighbourhood of a vertex can be used to refine the partition. So, a linear time algorithm should use each vertex as a pivot a constant number of times.

The linear time cograph recognition algorithm proposed in [HP05] computes a factoring permutation as a preliminary step. It roughly proceeds as follows. It uses at most one vertex per partition part to refine the ordered partition $[\overline{N}(v), \{v\}, N(v)]$. Assuming the input graph is a cograph, when none of the parts of the current partition is free of pivot, it can be proved that one of the two non-singleton parts closest to v in the current partition, say \mathcal{X} , can be refined into $[\overline{N}(x) \cap \mathcal{X}, \{x\}, N(x) \cap \mathcal{X}]$ (x being the used pivot of \mathcal{X}). This step creates at least one new part free of pivot and thereby relaunches the refining process.

6.3 From factoring permutation to modular decomposition tree

As already notice, a natural idea to compute the modular decomposition tree is to compute for each pair x, y of vertices the set of *splitter* $S(x, y)$. Unfortunately a linear time algorithm could not afford the computation of all these $O(n^2)$ sets. But if one has in hand a factoring permutation σ , it is then sufficient to consider the pairs of consecutive vertices in σ . Indeed, Capelle et al.'s algorithm [CHdM02] only computes for each pair of vertices $x = \sigma(i)$ and $y = \sigma(i + 1)$ ($i \in [1, n - 1]$) the leftmost and the rightmost (in σ) splitter of x and y . These two splitters define two intervals of σ , which are both contained in $m(x, y)$, the smallest module containing both x and y :

- the *left fracture* $Fg(x, y) = [z, x]$ if z is the leftmost splitter of $\{x, y\}$ in $[\sigma(1), y]$ (if any);
- the *right fracture* $Fd(x, y) = [y, z]$ if z is the rightmost splitter of $\{x, y\}$ in $[x, \sigma(n)]$ (if any).

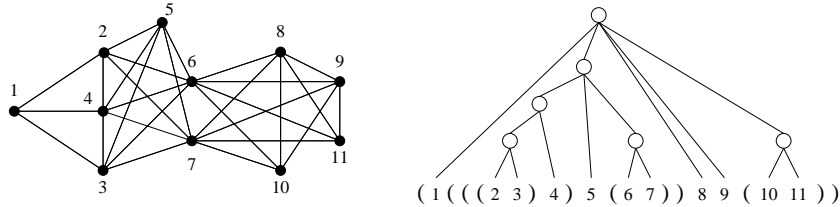


Fig. 12. A graph $G = (V, E)$ for which $\sigma = \mathbf{1\ 2\ 3\ 4\ 5\ 6\ 7\ 8\ 9\ 10\ 11}$ is a factoring permutation (see Definition 4). The right fracture of $(3, 4)$ does not exist but $Fg(3, 4) = [2, 3]$. We also have $Fd(1, 2) = [2, 7] = Fg(7, 8)$.

The set of fractures (left and right) defines a parenthesis system. Forgetting the initial pairing of the parenthesis, this system naturally yields a tree, called the *fracture tree* and denoted $FT(G)$ (see Figure 12). The fracture tree is actually a good estimation of the $MD(G)$ (see Lemma 14) which can be computed in linear time by two traversals of σ : the first traversal computes the fractures, the second builds the tree.

Lemma 14 [CHdM02] *Let σ be a factoring permutation of a graph G and M be a strong module of G . If M is a prime node of $MD(G)$ or if the father of M is a degenerate, then there exists a node N of the fracture tree $FT(G)$ such that M is the set of leaves of the subtree of $FT(G)$ rooted at N*

For example, in Figure 12, any strong module but $M = \{8, 9, 10, 11\}$ is represented by some node of $FT(G)$. Let us notice that the above lemma does not imply that the strong module $\{2, 3, 4\}$ has a corresponding node in $FT(G)$.

Henceforth to compute $MD(G)$, the fracture tree $FT(G)$ has to be cleaned. To that aim, Capelle et al. [CHdM02] use four extra traversals of the factoring permutation. The first one identifies the strong modules represented by some

of $FT(G)$; the second finds the dummy nodes of $FT(G)$; the third search for strong modules that are merged in a single node of $FT(G)$; and the last one remove the nodes of $FT(G)$ that does not represent strong modules. The complexity of each of these four traversal is linear in the size of G , $O(n + m)$.

6.4 Bibliographic notes

An attempt to generalize to arbitrary the linear time algorithm which computes a factoring permutation of a cograph has been proposed in [HdMP04]. Unfortunately the algorithm of [HdMP04] contains a flaw. The recent linear time modular decomposition algorithm presented in [TCHP08] mixes the ideas from the factoring permutation algorithms and the skeleton algorithm. It generalizes the ordered partition refining technique to tree partition and avoids union-find or least-common ancestor data-structures. In that sense this new algorithm may be considered as the positive answer to Spinrad comment (see Section 5.4).

7 Three novel applications of the modular decomposition

As mentioned in the introduction modular decomposition is used in a number of algorithmic graph theory applications and more generally applies to various discrete structures. We end this survey with the presentation of three novel applications which are good witnesses of the use of modular decomposition. The first one is a pattern matching problem which is closely related to the concept of factoring permutations. The second one provides an example of dynamic programming on the modular decomposition tree in the context of comparative genomics. Finally, we list a series of parameterized problems for which module based data-reduction rules leads to polynomial size kernel.

7.1 Pattern matching - common intervals of two permutations

Motivated by a series of genetic algorithms for sequencing problems, *e.g.* the TSP, Uno and Yagiura [UY00] formalized the concept of *common interval* of two permutations. As we will see in the next subsection, in the context of comparative genomics, common intervals reveal conserved structures in chromosomal material.

Definition 13 *A set S of elements is a common interval of a set of permutations Σ if in each permutation $\sigma \in \Sigma$, the elements of S form an interval of σ (see Section 2.2 for the definition of an interval).*

It is fairly easy to check that the family \mathcal{I} of common intervals of two permutations is a weakly partitive family (see Definition 1). It follows that all the theory presented in Section 2.1 applies and thus the set of strong common intervals are organized into a tree.

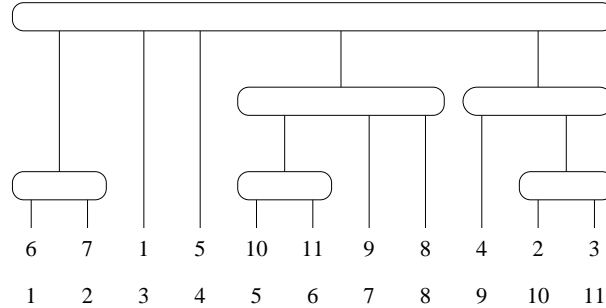


Fig. 13. The strong interval tree of two permutations. Remark $\{9, 10, 11\}$ and $\{8, 9\}$ are also a common interval, but they are not strong as the overlap.

In their paper, Uno and Yagiura proposed the first linear time algorithm to enumerate the common intervals of two permutations. More precisely, it runs in $O(n + K)$ time, where K is the number of those common intervals (which is possibly quadratic). Alternative algorithms have been recently proposed [HMS09, BCdMR08]. We sketch Uno and Yagiura’s algorithm and discuss how it can be adapted to compute the modules of a graph when a factoring permutation is given.

Without loss of generality, we will consider the problem of computing the common intervals of a permutation σ and the identity permutation \mathbb{I}_n . To identify the common intervals of a permutation σ and \mathbb{I}_n , the algorithm traverses σ only once. We denote by $[i, j]$ the interval of σ composed by the elements whose indexes are between i and j in σ : i.e. $[i, j] = \{x \mid i \leq \sigma(x) \leq j\}$. An element $x \notin [i, j]$ is a *splitter* of the interval $[i, j]$ if there exist $y \in [i, j]$ and $z \in [i, j]$ such that $y < x < z$. By $s([i, j])$ we denote the number $s([i, j])$ of *splitters* of the interval $[i, j]$. The algorithm uses a list `Potentiel` to filter and extract σ the common intervals of σ and \mathbb{I}_n . An element i belongs to the list `Potentiel` as long as it may be the right boundary of a common interval. The step i consists in removing those elements which we know they cannot be the left boundary of a common containing. This filtering can be done efficiently by computing $s([i, j])$ (see Lemmas 15 and 16).

The following properties are fundamental in the correctness of the algorithm:

Lemma 15 [UY00] *An interval $[i, j]$ of σ is a common interval of σ and \mathbb{I}_n iff $s([i, j]) = 0$.*

Lemma 16 [UY00, BXHP05] *If $s([i, j]) > s([i, j + 1])$, then it does not exist $r < i$ such that $[r, j]$ is a common interval of σ and \mathbb{I}_n .*

Algorithm 5: Uno and Yagiura’s algorithm [UY00]

Input: A permutation σ **Output:** The set of intervals common to σ and the identity permutation \mathbb{I}_n **begin** Let **Potentiel** be an empty list; **for** $i = n$ *downto* 1 **do** (Filter) Remove from **Potentiel** the boundaries r such $\forall j \leq i, [j, r]$
 is not a **common interval** of σ and \mathbb{I}_n ; (Addition) Add i to **Potentiel**; (Extraction) Search **Potentiel** to find the boundaries r such $[i, r]$ is
 a **common interval** of σ and \mathbb{I}_n and output those intervals $[i, r]$;**end**

The second lemma above means that if $s([i, j]) > s([i, j + 1])$ then the vertex $\sigma^{-1}(j + 1)$ is a splitter of $[i, j]$. Thereby any common interval containing $[i, j]$ as a subset has to extend up to $\sigma^{-1}(j + 1)$.

Application to factoring permutations of a graph. The most striking link between common intervals and modules of graphs is observed on permutation graphs (see Lemma 17). *Permutation graphs* are defined as the intersection graphs of a set of segments between two parallel lines (see [Gol80, BLS99] for example). It follows that the vertices of a permutation graph $G = (V, E)$ can be numbered from 1 to n such that there exists a permutation σ of $[1, n]$ such that vertex numbered i is adjacent to vertex numbered j iff $i < j$ and $\sigma(j) < \sigma(i)$. The permutations σ and \mathbb{I}_n form the *realizer* of G .

Lemma 17 [dM03] *Let $G = (V, E)$ be a permutation graph and (\mathbb{I}_n, σ) be its realizer. A set of vertices M is a strong module iff M is a strong common interval of \mathbb{I}_n and σ .*

The permutation graph corresponding to the permutations depicted in Figure 14 is the graph G of Figure 3. Notice that the strong interval tree of these two permutations is isomorphic the modular decomposition tree of G .

It follows from Lemma 17 that applied to the realizer of a permutation graph, Algorithm 5 computes the strong modules. Though some extra work is required to obtain the modular decomposition tree, the complexity remains linear time. Moreover, as shown in [BXHP05], Uno and Yagiura’s algorithm can directly be adapted to compute, given a factoring permutation, the strong modules of a graph. The number $s([i, j])$ becomes the number of splitters (in the sense of the modular decomposition, see Section 2.3) of the vertices contained in the interval $[i, j]$ of the factoring permutation. Now notice that Algorithm 5 does not only output the strong common intervals. In order to restrict the enumeration to strong modules, a slight modification is required.

A first traversal computes the strong right modules (i.e. the modules that are intervals of σ and which are not overlapped on their right boundary by any other module). Then a second traversal can detect those modules which are overlapped on the left boundary.

7.2 Comparative genomic - perfect sorting by reversals

A *reversal* in a permutation σ consists in reversing the order of the elements of an interval of σ . When dealing with *signed* permutations (whose elements are positive or negative), a reversal also flips the sign of the element of the reserved interval. Given two (signed) permutations σ and τ , the problem of *sorting by reversals* asks for a series of reversals (a *scenario*) to transform σ into τ .

Sorting by reversals is used in comparative genomic to measure the evolutionary distance between the genomes of two chromosomes, modeled as signed permutations [BHS02]. When comparing two genomic sequences, it can be assumed that the intervals having the same gene content are likely to have been present in their common ancestor and may witness to some functionally interacting proteins. Such a conserved genomic structure in the signed permutation model corresponds to common intervals. So to guess an evolution scenario between two genomic sequences represented by signed permutations σ and τ , one could ask for the smallest *perfect scenario*, which is a series of reversals that preserves any common interval of σ and τ . For further details on this topics, the reader could refer to [BHS02, BBCP04].

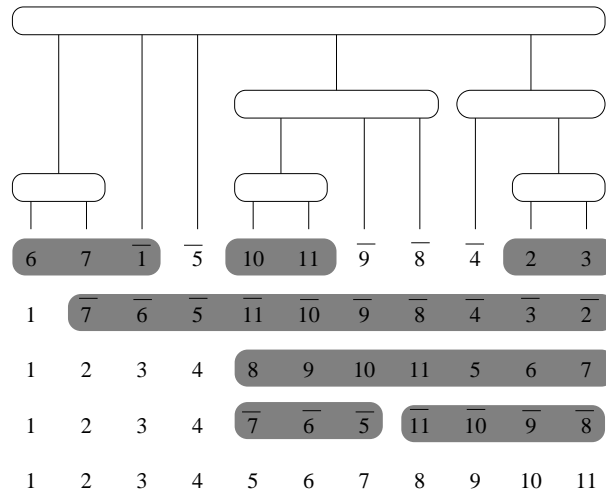


Fig. 14. A perfect scenario of length 7.

As mentioned in the previous subsection, the set of common intervals of two permutations (signed or not) defines a weakly partitive family. It follows that one can distinguish prime from degenerate strong common intervals. As shown

by the following lemma, we can be read on the strong interval tree which are the perfect scenarios.

Lemma 18 [BBCP04] *A reversal scenario for two signed permutations σ and τ is perfect iff any reversed interval is either a prime common interval of σ and τ , or the union of strong common intervals which form a subset of the children of a prime common interval.*

It follows from the previous lemma that the strong interval tree is useful to compute minimum perfect scenarios. Indeed with some extra technical properties to deal with the signs it can be shown that a simple dynamic programming algorithm on the strong interval tree solves the problem in time $O(2^k \times n\sqrt{n \log n})$, where k is the maximum number of prime nodes which are children of the same prime node. In practice, the parameter k keeps very small [BCP08]: e.g. when comparing the chromosome X of the mouse and the rat, we have $k = 0$ [BBCP07].

7.3 Parameterized complexity and kernel reductions - cluster editing

Parameterized algorithms is, among others, one of the modern techniques to cope with NP-hard problems. A problem Π is *fixed parameter tractable* (FPT) with respect to parameter k if it can be solved in time $f(k).n^{O(1)}$ where n is the input size. The idea behind parameterized algorithms is to find a parameter k , as small as possible, which controls the combinatorial explosion. Many algorithm techniques have been developed in the context of fixed parameter complexity, among which *kernelization*. A parameterized problem (Π, k) admits a *polynomial kernel* if there is a polynomial time algorithm (a set of *reduction rules*) that reduces the input instance to an instance whose size is bounded by a polynomial $p(k)$ depending only in k , while preserving the output. The classical example of parameterized problem having a polynomial kernel is the problem VERTEX COVER parameterized by k the solution size, which has a $2k$ size kernel. For textbooks on this topics, the reader should refer to [DF99, Nie06, FG06].

Recently, the modular decomposition appeared in kernelization algorithms for a series of parameterized problems among which: CLUSTER EDITING [Nie06], BICLUSTER EDITING [PdSS07], FAST (feedback arc set in tournament) [DGH⁺06], CLOSEST 3-LEAF POWER [BPP09], FLIP CONSENSUS TREE [BBT08]. We discuss the CLUSTER EDITING problem. Concerning the others, the reader should refer to the original papers.

The parameterized CLUSTER EDITING problem asks whether the edge set of an input graph G can be modified by at most k modifications (deletions or insertions) such that the resulting graph H is the disjoint union of cliques

(e.g. clusters). This problem is NP-complete but can be solved in time $O^*(3^k)$ by a simple bounded search tree algorithm [Cai96], which iteratively branches on at most k P_3 's. Recent papers [Guo07, FLRS07] showed the existence of a linear kernel (best bound is $4k$). The reduction rules used for these linear kernels are crown rules involving modules. For the sake of simplicity we only present the two basic reduction rules which leads to a quadratic kernel.

Lemma 19 *Let $G = (V, E)$ be a graph. A quadratic kernel for the CLUSTER EDITING problem is obtained by the following reduction rules:*

- (1) *Remove from G the connected components which are cliques.*
- (2) *If G contains a clique module C of size at least $k + 1$, then remove from $|C| - k - 1$ vertex from C .*

It is clear that these rules can be applied in linear time using modular decomposition algorithms. The proof idea works as follows. The first rule is obviously safe. For the second rule, simply observe that to disconnect a clique module of size $k + 1$ from the rest of the graph, at least $k + 1$ edge deletions are required. Now assuming G is a positive instance, each cluster of the resulting graph H can be bipartitioned into the vertices non-incident to a modified edge and the other vertices (the *affected vertices*). Finally k edge modifications can create at most $2k$ clusters and the total number of affected vertices is bounded by $2k$. This shows that the number of vertices in the reduced graph H is at most $2k^2 + 4k$.

The BICLUSTER EDITING problem edits the edge set of a graph to obtain a disjoint union of complete bipartite graphs. Instead of considering clique modules, we need to consider independent set modules [PdSS07]. The proof is then slightly more complicated and relies a careful analysis of the modification of the modular decomposition under edge insertion or deletion. In the case of FAST, similar rules involving to transitive modules also yields a quadratic kernel bound. Note that for these two problems, linear kernels can be obtained with more sophisticated reduction rules [GHKZ08, BFG+09].

References

- [BBCP04] A. Bergeron, S. Bérard, C. Chauve, and C. Paul. Sorting by reversal is not always difficult. In *Workshop on Algorithm for Bio-Informatics (WABI)*, Lecture Notes in Computer Science, 2004.
- [BBCP07] S. Bérard, A. Bergeron, C. Chauve, and C. Paul. Perfect sorting by reversals is not always difficult. *IEEE/ACM Trans. Comput. Biology Bioinform.*, 4(1):4–16, 2007.
- [BBT08] S. Böcker, Q.B. Anh Bui, and A. Truß. An improved fixed-

- parameter algorithm for minimum-flip consensus trees. In *International Workshop on Parameterized and Exact Computation (IWPEC)*, number 5018 in Lecture Notes in Computer Science, pages 43–54, 2008.
- [BCdMR05] A. Bergeron, C. Chauve, F. de Montgolfier, and M. Raffinot. Computing common intervals of k permutations, with applications to modular decomposition of graphs. In *European Symposium on Algorithms (ESA)*, number 3669 in Lecture Notes in Computer Science, pages 779–790, 2005.
- [BCdMR08] Anne Bergeron, Cedric Chauve, Fabien de Montgolfier, and Mathieu Raffinot. Computing common intervals of k permutations, with applications to modular decomposition of graphs. *SIAM Journal on Discrete Mathematics*, 22(3):1022–1039, 2008.
- [BCHP03] A. Bretscher, D.G. Corneil, M. Habib, and C. Paul. A simple linear time LexBFS cograph recognition algorithm. In *29th International Workshop on Graph-Theoretic Concepts in Computer Science (WG)*, volume 2880 of *Lecture Notes in Computer Science*, pages 119–130, 2003.
- [BCHP08] A. Bretscher, D. Corneil, M. Habib, and C. Paul. A simple linear time lexbfs cograph recognition algorithm. *SIAM Journal on Discrete Mathematics*, 22(4):1277–1296, 2008.
- [BCP08] S. Bérard, C. Chauve, and C. Paul. A more efficient algorithm for perfect sorting by reversals. *Information Processing Letters*, 106:90–95, 2008.
- [BFG⁺09] Stéphane Bessy, Fedor V. Fomin, Serge Gaspers, Christophe Paul, Anthony Perez, Saket Saurabh, and Stéphan Thomassé. Kernels for feedback arc set in tournaments. Technical Report CoRR abs/0907.2165, arxiv, 2009.
- [BHS02] A. Bergeron, S. Herber, and J. Stoye. Common intervals and sorting by reversals: a marriage of necessity. In *European Conference on Computational Biology*, Bioinformatics, pages 54–63, 2002.
- [Bla78] A. Blass. Graphs with unique maximal clumpings. *Journal of Graph Theory*, 2:19–24, 1978.
- [BLS99] A. Brandstädt, VB. Le, and J. Spinrad. *Graph classes: a survey*. SIAM Monographs on Discrete Mathematics and Applications. Society for Industrial and Applied Mathematics, 1999.
- [BPP09] S. Bessy, C. Paul, and A. Perez. Polynomial kernels for 3-leaf power graph modification problems. In *International Workshop on Combinatorial Algorithm (IWOCA)*, Lecture Notes in Computer Science, 2009.
- [BRV07] M. Bouvel, D. Rossin, and S. Vialette. Longest common separable pattern among permutations. In *Annual Symposium on Combinatorial Pattern Matching (CPM)*, number 4580 in Lecture Notes in Computer Science, pages 316–327, 2007.

- [BX08] B.M. Bui-Xuan. *Tree-representation of set families in graph decompositions and efficient algorithms*. PhD thesis, Univ. de Montpellier II, 2008.
- [BXH08] B.-M. Bui-Xuan and M. Habib. A representation theorem for union-difference families and application. In *Latin American Symposium on Theoretical Informatics (LATIN)*, volume 4957 of *Lecture Notes in Computer Science*, pages 492–503, 2008.
- [BXHP05] B.-M. Bui-Xuan, M. Habib, and C. Paul. Revisiting uno and yagiura’s algorithm. In *16th International Symposium on Algorithms and Computation (ISAAC)*, volume 3827 of *Lecture Notes in Computer Science*, pages 146–155, 2005.
- [Cai96] L. Cai. Fixed-parameter tractability of graph modification problems for hereditary properties. *Information Processing Letters*, 58(4):171–176, 1996.
- [Cap97] C. Capelle. *Décomposition de graphes et permutations factorisantes*. PhD thesis, Univ. de Montpellier II, 1997.
- [CE80] W.H. Cunningham and J. Edmonds. A combinatorial decomposition theory. *Canadian Journal of Mathematics*, 32(3):734–765, 1980.
- [CER93] B. Courcelle, J. Engelfriet, and G. Rozenberg. Handle rewriting graph grammars. *Journal of Computer and System Science*, 46:218–270, 1993.
- [CH94] A. Cournier and M. Habib. A new linear algorithm of modular decomposition. In *Trees in algebra and programming (CAAP)*, volume 787 of *Lecture Notes in Computer Science*, pages 68–84, 1994.
- [CH97] C. Capelle and M. Habib. Graph decompositions and factorizing permutations. In *Fifth Israel Symposium on Theory of Computing and Systems ISTCS’97, IEEE Conf. PR08037, Ramat-Gan, Israel*, pages 132–143, 1997.
- [CHdM02] C. Capelle, M. Habib, and F. de Montgolfier. Graph decompositions and factorizing permutations. *Discrete Mathematics and Theoretical Computer Science*, 5:55–70, 2002.
- [CHM81] M. Chein, M. Habib, and M.-C. Maurer. Partitive hypergraphs. *Discrete Mathematics*, 37:35–50, 1981.
- [CI98] A. Cournier and P. Ille. Minimal indecomposable graphs. *Discrete Mathematics*, 183:61–80, 1998.
- [CJS72] D.D. Cowan, L.O. James, and R.G. Stanton. Graph decomposition for undirected graphs. In *3rd S-E Conference on Combinatorics, Graph Theory and Computing, Utilitas Math*, pages 281–290, 1972.
- [CLR90] T.H. Cormen, C.E. Leiserson, and R.L. Rivest. *Algorithms*. MIT Press, 1990.
- [CLSB81] D.G. Corneil, H. Lerchs, and L.K. Stewart-Burlingham. Complement reducible graphs. *Discrete Applied Mathematics*, 3(1):163–

- 174, 1981.
- [CP06] C. Crespelle and C. Paul. Fully-dynamic recognition algorithm and certificate for directed cographs. *Discrete Applied Mathematics*, 2006. A paraître.
- [CPS85] D.G. Corneil, Y. Perl, and L.K. Stewart. A linear time recognition algorithm for cographs. *SIAM Journal on Computing*, 14(4):926–934, 1985.
- [Cre09] C. Crespelle. Fully dynamic representation of interval graphs. In *International Workshop on Graph Theoretical Concepts in Computer Science (WG)*, 2009.
- [Dah95] E. Dahlhaus. Efficient parallel algorithms for cographs and distance hereditary graphs. *Discrete Applied Mathematics*, 57:29–54, 1995.
- [DF99] R.G. Downey and M.R. Fellows. *Parameterized complexity*. Springer, 1999.
- [DGH⁺06] M. Dom, J. Guo, F. Hüffner, R. Niedermeier, and A. Truss. Fixed-parameter tractability results for feedback set problems in tournaments. In *Italian Conference on Algorithms and Complexity (CIAC)*, volume 3998 of *Lecture Notes in Computer Science*, pages 320–331, 2006.
- [DGM01] E. Dahlhaus, J. Gustedt, and R.M. McConnell. Efficient and practical algorithm for sequential modular decomposition algorithm. *Journal of Algorithms*, 41(2):360–387, 2001.
- [dM03] F. de Mongolfier. *Décomposition modulaire des graphes - Théorie, extensions et algorithmes*. PhD thesis, Univ. de Montpellier II, 2003.
- [EG97] J. Edmonds and R. Giles. A min-max relation for submodular functions on graphs. *Annals of Discrete Mathematics*, 1:185–204, 1997.
- [EGMS94] A. Ehrenfeucht, H.N. Gabow, R.M. McConnell, and S.L. Sullivan. An $o(n^2)$ divide-and-conquer algorithm for the prime tree decomposition of two-structures and modular decomposition of graphs. *Journal of Algorithms*, 16:283–294, 1994.
- [EHR99] A. Ehrenfeucht, T. Harju, and G. Rozenberg. *The theory of 2-structures*. World Scientific, 1999.
- [FG06] J. Flum and M. Grohe. *Parameterized complexity theory*. Texts in Theoretical Computer Science. Springer, 2006.
- [FLRS07] M.R. Fellows, M. Langston, F. Rosamond, and P. Shaw. Efficient parameterized preprocessing for cluster editing. In *International Symposium on Fundamentals of Computation Theory (FCT)*, number 4639 in *Lecture Notes in Computer Science*, pages 312–321, 2007.
- [Gal67] T. Gallai. Transitiv orientierbare graphen. *Acta Mathematica Acad. Sci. Hungar.*, 18:25–66, 1967.
- [GHKZ08] J. Guo, F. Hüffner, C. Komusiewicz, and Y. Zhang. Improved

- algorithms for bicluster editing. In *Annual Conference on Theory and Applications of Models of Computation (TAMC)*, volume 4978 of *Lecture Notes in Computer Science*, pages 445–456, 2008.
- [Gol80] M.C. Golumbic. *Algorithmic graph theory and perfect graphs*. Academic Press, 1980.
- [GP07] E. Gioan and C. Paul. Dynamic distance hereditary graphs using split decomposition. In *International Symposium on Algorithms and Computation (ISAAC)*, volume 4884 of *Lecture Notes in Computer Science*, 2007.
- [GPTC09a] E. Gioan, C. Paul, M. Tedder, and D. Corneil. Practical split-decomposition via graph-labelled trees. Submitted, 2009.
- [GPTC09b] E. Gioan, C. Paul, M. Tedder, and D. Corneil. Quasi-linear circle graph recognition. Submitted, 2009.
- [Guo07] J. Guo. A more effective linear kernelization for cluster editing. In *International Symposium Combinatorics, Algorithms, Probabilistic and Experimental Methodologies (ESCAPE)*, volume 4614 of *Lecture Notes in Computer Science*, pages 36–47, 2007.
- [HdMP04] M. Habib, F. de Montgolfier, and C. Paul. A simple linear-time modular decomposition algorithm. In *9th Scandinavian Workshop on Algorithm Theory (SWAT)*, volume 3111 of *Lecture Notes in Computer Science*, pages 187–198, 2004.
- [HHS95] M. Habib, M. Huchard, and J.S. Spinrad. A linear algorithm to decompose inheritance graphs into modules. *Algorithmica*, 13:573–591, 1995.
- [HM79] M. Habib and M.-C. Maurer. On the x -join decomposition of undirected graphs. *Discrete Applied Mathematics*, 1:201–207, 1979.
- [HM91] W.-L. Hsu and T.-H. Ma. Substitution decomposition on chordal graphs and applications. In *2nd International Symposium on Algorithms (ISA)*, volume 557 of *Lecture Notes in Computer Science*, pages 52–60, 1991.
- [HMPV00] 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. *Theoretical Computer Science*, 234:59–84, 2000.
- [HMS09] S. Herber, R. Mayr, and J. Stoye. Common intervals of multiple permutations. *Algorithmica*, 2009.
- [Hop71] J. Hopcroft. An $n \log n$ algorithm for minimizing states in a finite automaton. *Theory of machines and computations*, pages 189–196, 1971.
- [HP05] M. Habib and C. Paul. A simple linear time algorithm for co-graph recognition. *Discrete Applied Mathematics*, 145(2):183–187, 2005.
- [HPV98] M. Habib, C. Paul, and L. Viennot. A synthesis on partition refinement: a useful routine for strings, graphs, boolean matri-

- ces and automata. In *15th Symposium on Theoretical Aspect of Computer Science (STACS)*, volume 1373 of *Lecture Notes in Computer Science*, pages 25–38, 1998.
- [HPV99] M. Habib, C. Paul, and L. Viennot. Partition refinement : an interesting algorithmic tool kit. *International Journal of Foundation of Computer Science*, 10(2):147–170, 1999.
- [HSS01] P. Hell, R. Shamir, and R. Sharan. A fully dynamic algorithm for recognizing and representing proper interval graphs. *SIAM Journal on Discrete Mathematics*, 31(1):289–305, 2001.
- [Hsu92] W.-L. Hsu. A simple test for interval graphs. In *3rd International Symposium on Algorithms and Computation (ISAAC)*, volume 557 of *Lecture Notes in Computer Science*, pages 459–468, 1992.
- [Iba09] L. Ibarra. A fully dynamic graph algorithm for recognizing interval graphs. *Algorithmica*, 2009.
- [JO92a] B. Jamison and S. Olariu. Recognizing p_4 -sparse graphs in linear time. *SIAM Journal on Discrete Mathematics*, 21:381–406, 1992.
- [JO92b] B. Jamison and S. Olariu. A tree representation of p_4 -sparse graphs. *Discrete Applied Mathematics*, 35:115–129, 1992.
- [JO95] B. Jamison and S. Olariu. P-components and the homogeneous decomposition of graphs. *SIAM Journal on Discrete Mathematics*, 8(3):448–463, 1995.
- [Lov72] L. Lovász. A characterization of perfect graphs. *Journal of Combinatorial Theory Series B*, 13:95–98, 1972.
- [Möh85a] R.H. Möhring. Algorithmic aspect of comparability graphs and interval graphs. *Graphs and orders*, pages 42–101, 1985.
- [Möh85b] R.H. Möhring. Algorithmic aspect of the substitution decomposition in optimization over relations, set systems and boolean functions. *Annals of Operations Research*, 4:195–225, 1985.
- [MR84] R.H. Möhring and F.J. Rademacher. Substitution decomposition for discrete structures and connections with combinatorial optimization. *Annals of Discrete Mathematics*, 19:257–356, 1984.
- [MS89] J.H. Muller and J.P. Spinrad. Incremental modular decomposition. *Journal of the ACM*, 36(1):1–19, 1989.
- [MS94] R.M. McConnell and J.P. Spinrad. Linear-time modular decomposition and efficient transitive orientation of comparability graphs. In *5th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 536–545, 1994.
- [MS99] R.M. McConnell and J.P. Spinrad. Modular decomposition and transitive orientation. *Discrete Mathematics*, 201:189–241, 1999.
- [MS00] R.M. McConnell and J.P. Spinrad. Ordered vertex partitioning. *Discrete Mathematics and Theoretical Computer Science*, 4:45–60, 2000.
- [Nie06] R. Niedermeier. *Invitation to fixed parameter algorithms*, volume 31 of *Oxford Lectures Series in Mathematics and its Applications*. Oxford University Press, 2006.

- [PdSS07] F. Protti, M. Dantas da Silva, and J.L. Szwarcfiter. Applying modular decomposition to parameterized cluster editing problems. *Theory of Computing Systems*, 2007.
- [PT87] R. Paige and R.E. Tarjan. Three partition refinement algorithms. *SIAM Journal on Computing*, 16(6):973–989, 1987.
- [RST08] I. Rapaport, K. Suchan, and I. Todinca. Minimal proper interval completions. *Information Processing Letters*, 106(5):195–202, 2008.
- [RTL76] D.J. Rose, R.E. Tarjan, and G.S. Lueker. Algorithmic aspects of vertex elimination on graphs. *SIAM Journal on Computing*, 5(2):266–283, 1976.
- [Spi03] J.P. Spinrad. *Efficient graph representation*, volume 19 of *Fields Institute Monographs*. American Mathematical Society, 2003.
- [SS04] R. Shamir and R. Sharan. A fully dynamic algorithm for modular decomposition and recognition of cographs. *Discrete Applied Mathematics*, 136(2-3):329–340, 2004.
- [Sum73] D.P. Sumner. Graphs indecomposable with respect to the x -join. *Discrete Mathematics*, 6:281–298, 1973.
- [TCHP08] M. Tedder, D. Corneil, M. Habib, and C. Paul. Simpler linear-time modular decomposition via recursive factorizing permutations. In *International Colloquium on Automata, Languages and Programming (ICALP)*, volume 5125 of *Lecture Notes in Computer Science*, pages 634–645, 2008.
- [UY00] T. Uno and M. Yagiura. Fast algorithms to enumerate all common intervals of two permutations. *Algorithmica*, 26(2):290–309, 2000.