Available online at www.sciencedirect.com

ScienceDirect

Procedia Computer Science 195 (2021) 239–247

Procedia Computer Science

www.elsevier.com/locate/procedia

XI Latin and American Algorithms, Graphs and Optimization Symposium

Bounds on the Diameter of Graph Associahedra

Jean Cardinal[a,1,4], Lionel Pournin[b,2,4], Mario Valencia-Pabon[b,3,4]

[a] *Université libre de Bruxelles (ULB)*
[b] *LIPN, Université Sorbonne Paris Nord*

Abstract

Graph associahedra are generalized permutohedra arising as special cases of nestohedra and hypergraphic polytopes. The graph associahedron of a graph G encodes the combinatorics of search trees on G, defined recursively by a root r together with search trees on each of the connected components of $G - r$. In particular, the skeleton of the graph associahedron is the rotation graph of those search trees. We investigate the diameter of graph associahedra as a function of some graph parameters. It is known that the diameter of the associahedra of paths of length n, the classical associahedra, is $2n - 6$ for a large enough n. We give a tight bound of $\Theta(m)$ on the diameter of trivially perfect graph associahedra on m edges. We consider the maximum diameter of associahedra of graphs on n vertices and of given tree-depth, treewidth, or pathwidth, and give lower and upper bounds as a function of these parameters. Finally, we prove that the maximum diameter of associahedra of graphs of pathwidth two is $\Theta(n \log n)$.

Keywords: generalized permutohedra; graph associahedra; tree-depth; treewidth; pathwidth

1. Introduction

The vertices and edges of a polyhedron form a graph whose diameter (often referred to as the diameter of the polyhedron for short) is related to a number of computational problems. For instance, the question of how large the diameter of a polyhedron arises naturally from the study of linear programming and the simplex algorithm (see, for instance [27] and references therein). The case of associahedra [18, 31, 32]—whose diameter is known exactly [25]—is particularly interesting. Indeed, the diameter of these polytopes is related to the worst-case complexity of rebalancing binary search trees [29]. Here, we consider the same question on *graph associahedra* [6], a large family of *generalized permutohedra* in the sense of Postnikov [22] that can be built from an underlying graph. The question has already been

[1] jcardin@ulb.ac.be
[2] lionel.pournin@univ-paris13.fr
[3] valencia@lipn.univ-paris13.fr
[4] Supported by the grant “Propriétés combinatoires, algorithmiques et métriques des graphes de flips”, from the Partenariat Hubert Curien (PHC Tournesol grant number 42703TD).

1877-0509 © 2021 The Authors. Published by Elsevier B.V.

10.1016/j.procs.2021.11.030

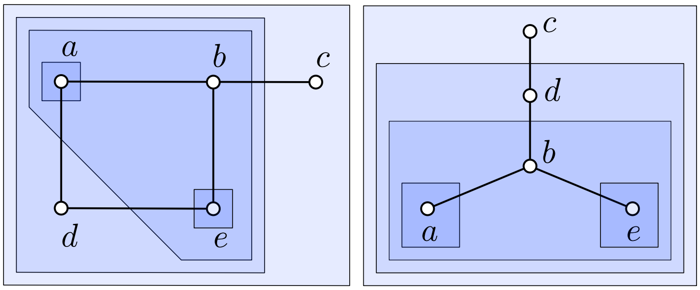

Fig. 1. An inclusionwise maximal tubing on a graph (left) and the search tree corresponding to it (right). The tubing is the set $\{\{a, b, c, d, e\}, \{a, b, d, e\}, \{a, b, e\}, \{a\}, \{e\}\}$. The tubes $\{a, b, e\}$ and $\{a\}$ are nested; the tubes $\{a\}$ and $\{e\}$ are non-adjacent.

studied by Manneville and Pilaud [19], by Pournin [26] in the special case of cyclohedra, and by Cardinal, Langerman and Pérez-Lantero [5] in the special case of tree associahedra. Here, we aim at giving tighter bounds on the diameter of graph associahedra, in terms of some structural invariants of the underlying graphs.

Graph associahedra have been defined by several authors including Davis, Januszkiewicz, and Scott [7], Carr and Devadoss [6], and Postnikov [22]. We give their definition in terms of so-called *tubings* on graphs, following Carr and Devadoss [6]. Thereafter, we let $G = (V, E)$ be a simple, connected graph with $|V| = n$. A *tube* S in G is a subset $S \subseteq V$ such that the induced subgraph $G[S]$ is connected. We say that a pair (S, S') of tubes is *nested* if either $S \subset S'$ or $S' \subset S$, and that it is *non-adjacent* if $G[S \cup S']$ is disconnected. A *tubing* $\mathcal{S}$ on G is a collection of tubes, every pair of which is either nested or non-adjacent. An example of (inclusionwise) maximal tubing on a graph is shown on the left of Figure 1.

The *graph associahedron* $\mathcal{A}(G)$ of G is a convex polytope whose face lattice is isomorphic to the inclusion order of tubings on G. In particular, the vertices are the maximal tubings on G. It is known that this polytope is always realizable, either by truncation of the permutohedron [10], or as a Minkowski sums of simplices [22]. The latter is in fact a generalization of Loday's realization of the associahedron [18].

When G is the complete graph on n vertices, no pair of tubes can be non-adjacent. Therefore, maximal tubings and vertices of the graph associahedron are in one-to-one correspondence with permutations of n elements. In that case, the graph associahedron $\mathcal{A}(G)$ is simply the $(n-1)$-dimensional permutohedron. Another special case of interest is when the graph G is a path on n vertices. In that case, maximal tubings are a Catalan family, and the graph associahedron of G is the $(n-1)$-dimensional classical associahedron. Similarly, the graph associahedra of cycles are the cyclohedra [26], and the graph associahedra of stars are stellohedra [6, 14, 23]. The number of maximal tubings on a graph G is known as the G-Catalan number [22].

Just as in the classical associahedron, whose edges correspond to flips in the triangulations of a convex polygon, the edges of a graph associahedron can be interpreted as *flips* in tubings: pairs of maximal tubings whose symmetric difference has size two. An example of such a flip is shown on Figure 2.

1.1. Search trees

We are interested in the structure of the skeleton of the graph associahedron $\mathcal{A}(G)$. For our purpose, it is useful to consider a representation of the vertices of $\mathcal{A}(G)$ alternative to the inclusionwise maximal tubings. A *search tree* T on G is a rooted tree with vertex set V defined recursively as follows: The root of T is a vertex $r \in V$, and r is connected to the root of search trees on each connected component of $G - r$. We will use the standard terminology related to rooted trees, in particular the *parent*, *child*, *ancestor*, and *descendant* relations. A vertex v together with its descendants in a search tree T form a *subtree* of T *rooted at* v. Search trees are in one-to-one correspondence with maximal tubings. Indeed, the tubes are exactly the subsets of the vertices of T contained in a subtree of T, as illustrated on the right of Figure 1. Such trees have been defined in many different contexts before. They are called $\mathcal{B}$-*trees* in Postnikov, Reiner, and Williams [23], and *spines* in Manneville and Pilaud [19]. In the context of polymatroids, they are special cases of the partial orders studied by Bixby, Cunningham, and Topkis [2]. In combinatorial optimization and graph theory, they can be defined in terms of *vertex rankings* [28, 3], *ordered colorings* [17], or as *elimination trees* [24].

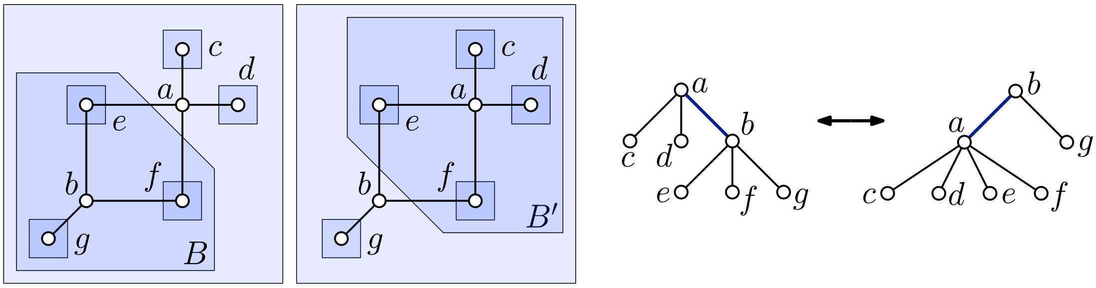

Fig. 2. A flip in a tubing, and the corresponding rotation in the associated search tree.

1.2. Flips in tubings and rotations in search trees

We will interpret the edges of $\mathcal{A}(G)$ as *rotations* in the search trees on G. A rotation in a search tree T involves a pair a, b of vertices, where a is the parent of b in T. The set A of vertices in the subtree rooted at a induces a connected subgraph of G, and forms a tube in the maximal tubing corresponding to T. The set B corresponding to the subtree rooted at b is a strict subset of A, and (A, B) is a nested pair of tubes. Note that B is one of the connected components of $G - a$. The rotation consists of picking b instead of a as the root of the subtree for the graph $G[A]$. The vertex a then becomes the root of the subtree on the connected component B' of $G[A] - b$ that contains a. After the rotation, each subtree rooted at a child of b is reattached to either a or b, depending on whether the child belongs to the same connected component of $G[A] - b$ as a or not. In terms of maximal tubings, it simply amounts to flipping the tubes B and B'.

The correpondence between flips in tubings and rotations in search trees is illustrated in Figure 2.

1.3. Applications to data structures

In a recent paper, Bose, Cardinal, Iacono, Koumoutsos, and Langerman [4] consider the design of competitive algorithms for the problem of searching in trees. They define a computation model involving search trees on trees, in which pointer moves and rotations all have unit cost. This can be seen as a generalization of the standard online binary search tree problem [30, 33, 9, 8], which has fostered developments in combinatorics, including the exact asymptotic estimate on the diameter of associahedra [29]. In the context of search trees on trees, the lower bound of Cardinal, Langerman, and Pérez-Lantero [5] on the diameter of tree associahedra ruled out some of the techniques that were known for binary search trees, and motivated the definition of *Steiner-closed* search trees. The importance of this notion is emphasized in a recent preprint by Berendsohn and Kozma [1]. Our results could give insights on potential generalizations to online search trees on graphs.

1.4. The diameter of graph associahedra

We will denote the diameter of the graph associahedron of G by $\delta(\mathcal{A}(G))$. Manneville and Pilaud proved the following tight bounds on that quantity as a function of the number of vertices and edges of G.

Theorem 1.1 (Manneville-Pilaud [19])**.** *For any connected graph G on n vertices and m edges, the diameter of the graph associahedron of G satisfies*

$$\max(m, 2n - 20) \le \delta(\mathcal{A}(G)) \le \binom{n}{2}.$$

In order to distinguish between these two extreme cases (linear versus quadratic diameter), we aim at bounds expressed as a function of some graph invariants.

The following result, also from Manneville and Pilaud, will be useful as well .

Theorem 1.2 (Manneville-Pilaud [19])**.** *The diameter $\delta(\mathcal{A}(G))$ is non-decreasing: $\delta(\mathcal{A}(G)) \le \delta(\mathcal{A}(G'))$ for any two graphs G, G' such that G is a subgraph of G'.*

1.5. Pathwidth, treewidth, and tree-depth

We consider three classical numerical invariants of a graph G. We refer the reader to the texts of Diestel [11] and Nešetřil and Ossona de Mendez [20] for details and alternative definitions of those parameters.

The *pathwidth* of a graph G is $\omega - 1$, where ω is the smallest clique number of an interval supergraph of G, that is, an interval graph that can be obtained from G by adding edges. Similarly, the *treewidth* of a graph G is exactly one less than the smallest clique number of a chordal supergraph of G. Pathwidth and treewidth can also be defined in terms of path and tree decompositions, respectively. Paths have pathwidth one, trees have treewidth one, and on an intuitive level, those two parameters quantify how close the graph is to a path or a tree. They also play a key role in the theory of graph minors and in graph algorithms.

The *tree-depth* of a graph G is the smallest height of a search tree on G, where a tree composed of a single vertex has height one. The tree-depth is definitely a natural invariant to consider, as it is a function of the exact same objects that form the vertices of the graph associahedron. Surprisingly, this connection does not seem to have been exploited in previous works.

1.6. Our Results

We first prove that the lower bound of m from Manneville and Pilaud on the diameter of the associahedra of a graph on m edges is essentially tight for all *trivially perfect* graphs. Those graphs appear naturally here, as they are maximal for a fixed tree-depth. In Section 2, we properly define trivially perfect graphs, and prove the following result.

Theorem 1.3. *Let G be a connected trivially perfect graph with m edges. Then the diameter of the graph associahedra of G is $\delta(\mathcal{A}(G)) = \Theta(m)$.*

In Section 3, we refine the bounds on the diameter of graph associahedra whose underlying graphs have bounded tree-depth or treewidth. Given a family $\mathcal{G}$ of graphs, we consider the worst-case diameter

$$\delta_{\mathcal{G}}(n) = \max_{G \in \mathcal{G} : |V(G)| = n} \delta(\mathcal{A}(G))$$

of their graph associahedra. The tree-depth is an example of parameter that precisely controls the behavior of the diameter of the associahedra, in the following worst-case sense.

Theorem 1.4. *Let $\mathcal{G}$ be the family of graphs on n vertices and of tree-depth at most $td(n)$. Then $\delta_{\mathcal{G}}(n) = \Theta(td(n) \cdot n)$.*

We obtain the following lower and upper bounds as a function of the treewidth of the graph.

Theorem 1.5. *Let $\mathcal{G}$ be the family of graphs on n vertices and of treewidth at most $tw(n)$. Then $\Omega(tw(n) \cdot n) \leq \delta_{\mathcal{G}}(n) \leq O(tw(n) \cdot n \log n)$.*

The same bounds hold as a function of the pathwidth of the graph.

Theorem 1.6. *Let $\mathcal{G}$ be the family of graphs on n vertices and of pathwidth at most $pw(n)$. Then $\Omega(pw(n) \cdot n) \leq \delta_{\mathcal{G}}(n) \leq O(pw(n) \cdot n \log n)$.*

The case of graphs of pathwidth at most *two* is intriguing, as one could have suspected that the diameter is close to that of the associahedra (the *path* associahedra). In fact, the diameter jumps from linear to linearithmic, as we show in Section 4.

Theorem 1.7. *Let $\mathcal{G}$ be the family of graphs of pathwidth two. Then $\delta_{\mathcal{G}}(n) = \Theta(n \log n)$.*

2. The diameter of trivially perfect graph associahedra

A graph is *trivially perfect* if it is both a *cograph* and an *interval graph*. Wolk [34] called these graphs *comparability graphs of trees* and gave characterizations of them. Golumbic [15] called them *trivially perfect graphs* because it is trivial to show that such a graph is a *perfect* graph. These graphs have the property that in each of their induced subgraphs the size of the maximum independent set is also the number of maximal cliques.

A *universal vertex* in a graph G is a vertex which is adjacent to all the other vertices in G. A *maximal universal clique* in a graph G is a maximal clique C in G such that each vertex in C is a universal vertex in G. Yan et al. [35] give some equivalent definitions of trivially perfect graphs. In particular, they define trivially perfect graph recursively as follows: (1) an isolated vertex (i.e. K_1) is a trivially perfect graph, (2) adding a new universal vertex to a trivially perfect graph results in a trivially perfect graph, and (3) the disjoint union of two trivially perfect graphs is a trivially perfect graph.

Drange et al. [12] give the following decomposition of a trivially perfect graph. Let T be a rooted tree and t be a node of T. We denote by T_t the maximal subtree of T rooted in t.

Definition 2.1. (Definition 2.3 in [12]) *A universal clique decomposition of a connected trivially perfect graph $G = (V, E)$ is a pair $(T = (V_T, E_T), \mathcal{B} = \{B_t\}_{t\in V_T})$, where T is a rooted tree and $\mathcal{B}$ is a partition of the vertex set V into disjoint nonempty subsets, such that*

- *if $vw \in E$, $v \in B_t$, and $w \in B_s$, then s and t are on a path from a leaf to the root (and, possibly $s = t$), and*
- *for every node $t \in V_T$, B_t is the maximal universal clique in the subgraph of G induced by $\bigcup_{s\in V(T_t)} B_s$.*

The vertices of T are called *nodes* and the sets in $\mathcal{B}$ *bags* of the universal clique decomposition $(T, \mathcal{B})$. Notice that in a universal clique decomposition, every nonleaf node t has at least two children, since otherwise the universal clique contained in the bag corresponding to t would not be maximal. Drange et al. [12] have shown that a connected graph G admits a universal clique decomposition if and only if it is trivially perfect. Moreover, such a decomposition is unique up to isomorphism. Figure 3 shows a trivially perfect graph (left) and its universal clique decomposition (center). It is well known that the tree-depth of a graph G is the minimum size of the largest clique in a trivially perfect supergraph of G (see [21]).

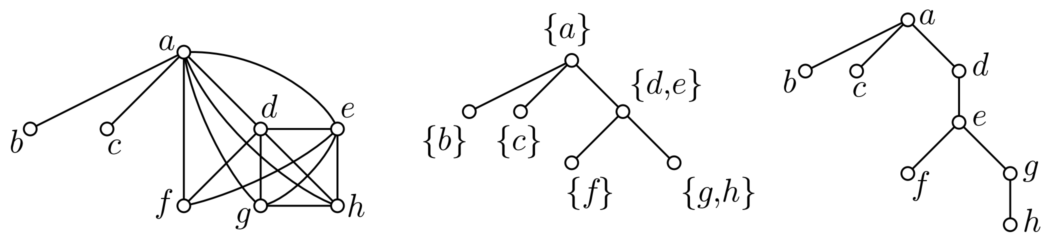

Fig. 3. A trivially perfect graph G (left), the universal clique decomposition of G (center), and a minimum height search tree on G constructed from its universal clique decomposition.

Consider a connected trivially perfect graph $G = (V, E)$ with clique number ω. One can use the universal clique decomposition $(T, \mathcal{B})$ of G in order to construct a search tree T' of G whose height is equal to the tree-depth of G and, therefore to ω: let r be the root of T and denote by r_1 to r_p its child nodes. Form a path P of length $|B_r|$ whose vertices are the elements of B_r (arranged in any order). One of the ends of this path will be the root of T'. Now, denote by C_i the connected component of $G[V\backslash B_r]$ that admits r_i as a subset of its vertices. The pair $(T_{r_i}, \{B_j : j \in V(T_{r_i})\})$ turns out to be the universal clique decomposition of C_i. Therefore, one can use the procedure recursively in order to build a search tree for each C_i. Connecting the root of these search trees by an edge to one of the ends of P results in the announced search tree T' of G with height ω. This procedure is illustrated in Figure 3, where the minimum height search tree is shown on the right.

Theorem 2.2. *Let G be a connected trivially perfect graph on m edges. Then $\delta(\mathcal{A}(G)) \leq 2m$.*

Proof. Let $td(G) = \omega(G) = k$. We prove by induction on k that any search tree on G can be transformed into a search tree of height k in at most m rotations. When $k = 2$, the graph G has a single edge and the result is immediate. Let S be a search tree on G of height k with root r and denote by r_1 to r_p the child vertices of r in S. Such a search tree can be obtained from the universal clique decomposition tree of G as described above. Let T another search tree on G. Then, there exists a sequence of at most $n-1$ rotations that transform T into a tree P where r has been lifted at the root. Clearly, r belongs to a maximal universal clique in G and thus, $G\backslash r$ has p connected components, say C_1 to C_p, each inducing a trivially perfect subgraph of G with clique number at most $k-1$. By the definition of the universal clique decomposition, each subtree S_{r_i} is a search tree of height at most $k-1$ on $G[C_i]$. Denote by s_1 to s_p the child vertices of r in P. By induction, there exists a sequence of at most $|E(G[C_i])|$ rotations that transform the tree P_{s_i}, rooted at s_i, into the subtree T_{r_i}. Therefore, T can be transformed into S in at most $n-1+\sum_{1\leq i\leq p}|E(G[C_i])| = m$ rotations. Since any search tree on G can be transformed into S in at most m rotations, the diameter of $\mathcal{A}(G)$ is at most $2m$. □

Theorem 1.3 is a consequence of Theorems 2.2 and 1.1.

3. Diameter, Tree-Depth, and Treewidth

In this section, we establish tight bounds on the diameter of graph associahedra in terms of the tree-depth of the underlying graph. We will make use of the following.

Lemma 3.1. *Let H be a trivially perfect graph on n vertices and m edges, with tree-depth td. Then $m < td \cdot n$.*

Proof. This is easily proved by induction. If $td = 1$, then $n = 1$ and $m = 0$. Assume that the statement holds for all graphs H with less than n vertices and tree-depth less than td. If r is the root of a search tree T on H of height td, then $m = n-1+|E(H-r)|$. By induction, $|E(H-r)| < (td-1)(n-1)$. As a consequence, $m \leq td\cdot(n-1) < td\cdot n$, as desired. □

Theorem 3.2. *Let G be a graph on n vertices, of tree-depth at most td. Then $\delta(\mathcal{A}(G)) \leq 2\cdot td\cdot n$.*

Proof. By definition, there exists a trivially perfect supergraph H of G with clique number td. By Lemma 3.1, H has at most $td\cdot n$ edges. Hence, Theorems 1.2 and 2.2, yield $\delta(\mathcal{A}(G)) \leq \delta(\mathcal{A}(H)) \leq 2\cdot td\cdot n$. □

Let us prove that this bound is tight for a wide range of values of $td(G)$, up to a constant factor.

Theorem 3.3. *For any two positive integers k and n such that k divides n, there exists a trivially perfect graph G on $n+1$ vertices such that $td(G) = n/k+1$ and $\delta(\mathcal{A}(G)) \geq td(G)\cdot n/2$.*

Proof. The graph G is composed of k cliques $C_1, C_2, \ldots, C_k$, each of size $n/k+1$, and a designated vertex v such that $C_i \cap C_j = \{v\}$ when i and j are distinct. Clearly, G is trivially perfect, $td(G) = n/k+1$, and the number of edges of G is $k\cdot\binom{n/k+1}{2} = n^2/(2k)+n/2 = td(G)\cdot n/2$. The desired bound on the diameter of $\mathcal{A}(G)$ therefore follows from the lower bound stated by Theorem 1.1. □

From this construction, we obtain families of polyhedra parameterized by a function $td(n)$ that interpolate between stellohedra (a star has tree-depth two) and permutohedra (a complete graph has tree-depth n). Theorems 3.2 and 3.3 together prove Theorem 1.4.

Corollary 3.4. *Let G be a graph on n vertices, of treewidth at most tw. Then $\delta(\mathcal{A}(G)) \leq c\cdot tw\cdot n\log n$ for some constant c.*

Proof. This follows from Theorem 3.2 and the known fact that $td(G) = O(tw(G)\cdot\log n)$ [20]. □

4. Diameter and Pathwidth

We first prove a lower bound on the diameter of graph associahedra in terms of pathwidth.

Theorem 4.1. *For any $k \geq 2$ and any n multiple of $k-1$, there exists an interval graph on $n+1$ vertices and of pathwidth $k-1$ such that $\delta(\mathcal{A}(G)) \geq nk/2$.*

Proof. We consider a graph G induced by a sequence of cliques $C_1, C_2, \ldots, C_{n/(k-1)}$, each of size k, such that any two consecutive cliques C_i and C_{i+1} have a single vertex in common, and no other pairs of cliques have a vertex in common. This graph is clearly an interval graph of pathwidth $k-1$, and its number of edges is $\binom{k}{2}n/(k-1) \geq nk/2$. The conclusion follows from the lower bound in Theorem 1.1. □

Theorem 4.1 and Corollary 3.4 together prove Theorems 1.5 and 1.6.

Note that connected graphs of pathwidth one are exactly paths, and their graph associahedra have linear diameter. Interestingly, as we shall see, the diameter jumps to $\Omega(n \log n)$ for graphs of pathwidth two. Our proof uses a construction similar to the one from Cardinal, Langerman, and Pérez-Lantero [5] for tree associahedra. We need some preliminaries involving chordal graphs and *projections* of rotation sequences.

4.1. Chordal graphs

A graph G is *chordal* if it does not contain induced cycles of length 4 or more. In other words, every cycle in G of length 4 or more has a *chord*. We denote by $N(v)$ the set of the neighbors of a vertex v in G. A vertex v is said to be *simplicial* if $G[N(v)]$ is a clique. It is known that a graph is chordal if and only if it has a *perfect elimination ordering*: an ordering of its vertices such that the set of neighbors of a vertex v that come after v in the ordering induce a clique. Hence a perfect elimination ordering is obtained by iteratively removing a simplicial vertex in the remaining subgraph. The set of the vertices remaining after removing the vertices in a prefix of a perfect elimination ordering is called *monophonically convex*, or *m-convex*; see Farber and Jamison [13] and references therein for details. In what follows, we will simply use the term *convex*.

4.2. Projections

We now introduce a tool that will turn out useful for performing induction on chordal graphs.

Observation 1. *Consider a chordal graph G on at least two vertices, a simplicial vertex v of G, and a search tree T on G. Then v has at most one child in T. Further consider the tree T' obtained as follows:*

1. *If v is a leaf of T, just remove v from the tree.*
2. *If v is the root of T, then remove v from T and designate its child as the new root.*
3. *If v has both a parent and a child, then remove v from T and replace the two edges between v and its parent and child by a single edge between its parent and its child.*

Then T' is a search tree on $G - v$.

Given an initial search tree on G, we can construct a search tree on $G[S]$ by induction, for any convex subset $S \subseteq V$. Note that the obtained search tree does not depend on the order in which the simplicial vertices have been removed. Letting $T_{|S}$ be the tree thus obtained, we call it the *projection* of T on S.

Observation 2. *Let T be a search tree on $G = (V, E)$, and let T' be obtained from T by performing a single rotation. Then the projections of T and T' on $V\backslash\{v\}$ are either identical or related by a single rotation. They are identical if and only if the rotation between T and T' involves v.*

The following lemma is a consequence of the above observations and generalizes Lemma 3 in [5].

Lemma 4.2 (Projection Lemma)**.** *Let $G = (V, E)$ be a chordal graph, and π a sequence of rotations transforming a search tree T on G into another one, say T'. Let S be a convex subset of V. The* projection $\pi_{|S}$ *of π on S is the sequence of rotations obtained by removing from π all rotations involving two vertices at least one of whose does not belong to S. Then $\pi_{|S}$ is a rotation sequence that transforms $T_{|S}$ into $T'_{|S}$.*

4.3. An $\Omega(n \log n)$ lower bound for graphs of pathwidth two

Let us first recall how *bit-reversal* permutations work [16, 33, 8]. We denote a permutation on n elements by a sequence composed of one occurence of each integer in $[n]$. The bit-reversal permutation of length one is $\sigma_1 = 1$. The kth bit-reversal permutation σ_k has length $n = 2^{k-1}$ and is obtained by concatenating $2\sigma_{k-1} - 1$ with $2\sigma_{k-1}$. Note that the permutation σ_k of length n alternates between entries $\leq n/2$ and entries $> n/2$. Here are the first five bit-reversal permutations:

$$\sigma_1 = 1$$
$$\sigma_2 = 1, 2$$
$$\sigma_3 = 1, 3, 2, 4$$
$$\sigma_4 = 1, 5, 3, 7, 2, 6, 4, 8$$
$$\sigma_5 = 1, 9, 5, 13, 3, 11, 7, 15, 2, 10, 6, 14, 4, 12, 8, 16.$$

Theorem 4.3. *Let n be a power of two. If $n \geq 4$, then there exists a graph G of pathwidth two on n vertices such that $\delta(\mathcal{A}(G)) \geq \Omega(n \log n)$.*

Proof. Let $n = 2^{k-1}$ for some k. We construct a graph G_n of pathwidth two composed of $2n$ vertices denoted by $a_1, a_2, \ldots, a_n, b_1, b_2, \ldots, b_n$. We let a_i be adjacent to a_{i+1} and b_{i+1} for all $i \in [n-1]$, b_i be adjacent to b_{i+1} for all $i \in [n-1]$, and a_i be adjacent to b_i for all $i \in [n]$. It can be checked that G_n is an interval graph with clique number three, hence of pathwidth two. The graph is shown on Figure 4.

We define the two subsets of vertices $L = \bigcup_{i \leq n/2}\{a_i, b_i\}$, and $R = \bigcup_{i > n/2}\{a_i, b_i\}$. Both subsets are convex. Further note that $G[L]$ and $G[R]$ are both isomorphic to $G_{n/2}$. We now consider the rotation distance between two search trees T and T' on G_n. These trees are paths of the following form, rooted at the first vertex:

$$\begin{array}{lrl} T & a_1, b_1, a_2, b_2, \ldots, a_{n/2}, b_{n/2}, & a_{n/2+1}, b_{n/2+1}, \ldots, a_n, b_n \\ T' & a_{\sigma_k(1)}, a_{\sigma_k(2)}, a_{\sigma_k(3)}, a_{\sigma_k(4)}, \ldots, a_{\sigma_k(n-1)}, a_{\sigma_k(n)}, & b_1, b_2, \ldots, b_{n/2}, b_{n/2+1}, \ldots, b_n \end{array}$$

The first path, T, is indeed a search tree because it corresponds to a perfect elimination ordering. The second path, T', is also a search tree, because the subgraphs corresponding to subtrees rooted at each of the a_i are connected, thanks to the b_i. By the recursive definition of bit-reversal permutations, the search trees $T_{|L}$ and $T'_{|L}$ on $G[L]$ are obtained in the exact same way, by permuting the elements of L according to σ_{k-1}. The same holds for $T_{|R}$ and $T'_{|R}$ on $G[R]$, up to a shift of the indices.

The sets L and R induce a two-coloring of any search tree on G_n. An edge of a search tree will be called *monochromatic* if both endpoints belong to the same set L or R, and *bichromatic* otherwise. Similarly, we distinguish monochromatic rotations, involving pairs of vertices from the same set L or R, from bichromatic rotations involving one vertex of each set.

Let π be a sequence of rotations transforming T into T', of minimum length $\ell(n)$. From our previous observations and Lemma 4.2, the number of monochromatic rotations in π is $|\pi_{|L}| + |\pi_{|R}| \geq 2\ell(n/2)$. We now give a lower bound on the number of bichromatic rotations in π. Given a search tree on G_n, we define its *alternation number* as the maximum, over all paths from the root to a leaf, of the number of bichromatic edges on the path. Note that, from the property of bit-reversal permutations, the alternation number of T' is $n+1$. On the other hand, the alternation number of T is 1. We make two observations. First, a monochromatic rotation cannot increase the alternation number of a tree. Indeed, consider a rotation involving vertices a, b, with b the child of a (refer to Figure 2). The only edges of the tree whose endpoints are changed by the rotation are the edge from the parent of a, and edges connecting b to the root of a subtree in the initial tree. If a and b have the same color, none of these edges can become bichromatic.

Second, a bichromatic rotation can only increase the alternation number of a search tree by two. Indeed again, on a path from the root to a leaf, only two edges can become bichromatic by a rotation involving a and b: the one from the parent of a and one from b to the root of a subtree.

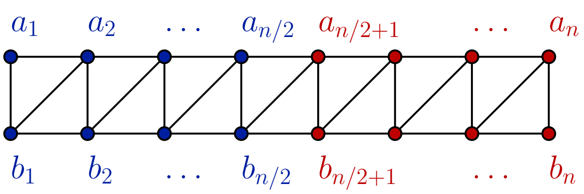

Fig. 4. The graph G_n used in the proof of Theorem 4.3.

We conclude that there must be at least $n/2$ bichromatic rotations in π. Summing the number of monochromatic and bichromatic rotations, we obtain $\ell(n) \geq 2\ell(n/2) + n/2 = \Omega(n \log n)$. □

References

[1] Berendsohn, B.A., Kozma, L., 2020. Splay trees on trees. arXiv:2010.00931 .
[2] Bixby, R.E., Cunningham, W.H., Topkis, D.M., 1985. The partial order of a polymatroid extreme point. Mathematics of Operations Research 10, 367–378.
[3] Bodlaender, H.L., Deogun, J.S., Jansen, K., Kloks, T., Kratsch, D., Müller, H., Tuza, Z., 1998. Rankings of graphs. SIAM J. Discrete Math. 11, 168–181.
[4] Bose, P., Cardinal, J., Iacono, J., Koumoutsos, G., Langerman, S., 2020. Competitive online search trees on trees, in: Proceedings of the Annual ACM-SIAM Symposium on Discrete Algorithms (SODA20), pp. 1878–1891.
[5] Cardinal, J., Langerman, S., Pérez-Lantero, P., 2018. On the diameter of tree associahedra. Electr. J. Comb. 25, P4.18.
[6] Carr, M., Devadoss, S.L., 2006. Coxeter complexes and graph-associahedra. Topology and its Applications 153, 2155–2168.
[7] Davis, M.W., Januszkiewicz, T., Scott, R.A., 2003. Fundamental groups of blow-ups. Advances in Mathematics 177, 115–179.
[8] Demaine, E.D., Harmon, D., Iacono, J., Kane, D.M., Patrascu, M., 2009. The geometry of binary search trees, in: Proceedings of the Twentieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2009, New York, NY, USA, January 4-6, 2009, pp. 496–505.
[9] Demaine, E.D., Harmon, D., Iacono, J., Patrascu, M., 2007. Dynamic optimality - almost. SIAM J. Comput. 37, 240–251.
[10] Devadoss, S.L., 2009. A realization of graph associahedra. Discrete Mathematics 309, 271–276.
[11] Diestel, R., 2005. Graph Theory. Graduate Texts in Mathematics, Springer.
[12] Drange, P.G., Fomin, F.V., Pilipczuk, M., Villanger, Y., 2015. Exploring the subexponential complexity of completion problems. ACM Transactions on Computation Theory 7, 14:1–14:38.
[13] Farber, M., Jamison, R.E., 1986. Convexity in graphs and hypergraphs. SIAM Journal on Algebraic Discrete Methods 7, 433–444.
[14] Forcey, S., Lauve, A., Sottile, F., 2009. New Hopf Structures on Binary Trees. Discrete Mathematics & Theoretical Computer Science .
[15] Golumbic, M.C., 1978. Trivially perfect graphs. Discrete Mathematics 24, 105–107.
[16] Karp, A.H., 1996. Bit reversal on uniprocessors. SIAM Review 38, 1–26.
[17] Katchalski, M., McCuaig, W., Seager, S., 1995. Ordered colourings. Discrete Mathematics 142, 141 – 154.
[18] Loday, J.L., 2004. Realization of the stasheff polytope. Archiv der Mathematik 83, 267–278.
[19] Manneville, T., Pilaud, V., 2015. Graph properties of graph associahedra. Séminaire Lotharingien de Combinatoire B73d.
[20] Nešetřil, J., Ossona de Mendez, P., 2012. Sparsity: Graphs, Structures, and Algorithms. Springer. chapter 6. pp. 115–144.
[21] Nešetřil, J., de Mendez, P.O., 2015. On low tree-depth decompositions. Graphs Comb. 31, 1941–1963.
[22] Postnikov, A., 2009. Permutohedra, associahedra, and beyond. International Mathematics Research Notices 2009, 1026–1106.
[23] Postnikov, A., Reiner, V., Williams, L.K., 2008. Faces of generalized permutohedra. Documenta Mathematica 13, 207–273.
[24] Pothen, A., 1988. The complexity of optimal elimination trees. Tech. Report CS-88-13, Pennsylvania State University.
[25] Pournin, L., 2014. The diameter of associahedra. Advances in Mathematics 259, 13–42.
[26] Pournin, L., 2017. The asymptotic diameter of cyclohedra. Israel Journal of Mathematics 219, 609–635.
[27] Santos, F., 2012. A counterexample to the Hirsch conjecture. Annals of Mathematics 176, 383–412.
[28] Schäffer, A.A., 1989. Optimal node ranking of trees in linear time. Information Processing Letters 33, 91–96.
[29] Sleator, D., Tarjan, R., Thurston, W., 1988. Rotation distance, triangulations, and hyperbolic geometry. Journal of the American Mathematical Society 1, 647–681.
[30] Sleator, D.D., Tarjan, R.E., 1985. Self-adjusting binary search trees. J. ACM 32, 652–686.
[31] Stasheff, J.D., 1963. Homotopy associativity of H-spaces. I. Transactions of the American Mathematical Society 108, 275–292.
[32] Tamari, D., 1951. Monoïdes préordonnés et chaînes de Malcev. Thèse de Mathématiques, Paris.
[33] Wilber, R.E., 1989. Lower bounds for accessing binary search trees with rotations. SIAM J. Comput. 18, 56–67.
[34] Wolk, E.S., 1962. The comparability graph of a tree. Proceedings of the American Mathematical Society 13, 789–795.
[35] Yan, J., Chen, J., Chang, G.J., 1996. Quasi-threshold graphs. Discrete Applied Mathematics 69, 247–255.