

GRAPH PROPERTIES OF GRAPH ASSOCIAHEDRA

THIBAUT MANNEVILLE AND VINCENT PILAUD

ABSTRACT. A graph associahedron is a simple polytope whose face lattice encodes the nested structure of the connected subgraphs of a given graph. In this paper, we study certain graph properties of the 1-skeleta of graph associahedra, such as their diameter and their Hamiltonicity. Our results extend known results for the classical associahedra (path associahedra) and permutahedra (complete graph associahedra). We also discuss partial extensions to the family of nestohedra.

KEYWORDS. Graph associahedron, nestohedron, graph diameter, Hamiltonian cycle.

1. INTRODUCTION

Associahedra are classical polytopes whose combinatorial structure was first investigated by J. Stasheff [Sta63] and later geometrically realized by several methods [Lee89, GKZ08, Lod04, HL07, PS12, CSZ15]. They appear in different contexts in mathematics, in particular in algebraic combinatorics (in homotopy theory [Sta63], for construction of Hopf algebras [LR98], in cluster algebras [CFZ02, HLT11]) and discrete geometry (as instances of secondary or fiber polytopes [GKZ08, BFS90] or brick polytopes [PS12, PS15]). The combinatorial structure of the n -dimensional associahedron encodes the dissections of a convex $(n + 3)$ -gon: its vertices correspond to the triangulations of the $(n + 3)$ -gon, its edges correspond to flips between them, etc. See Figure 1. Various combinatorial properties of these polytopes have been studied, in particular in connection with the symmetric group and the *permutahedron*. The combinatorial structure of the n -dimensional permutahedron encodes ordered partitions of $[n + 1]$: its vertices are the permutations of $[n + 1]$, its edges correspond to transpositions of adjacent letters, etc.

In this paper, we are interested in graph properties, namely in the diameter and Hamiltonicity, of the 1-skeleta of certain generalizations of the permutahedra and the associahedra. For the n -dimensional permutahedron, the diameter of the *transposition graph* is the number $\binom{n+1}{2}$ of inversions of the longest permutation of $[n + 1]$. Moreover, H. Steinhaus [Ste64], S. M. Johnson [Joh63], and H. F. Trotter [Tro62] independently designed an algorithm to construct a Hamiltonian cycle of this graph. For the associahedron, the diameter of the *flip graph* motivated intensive research and relevant approaches, involving volumetric arguments in hyperbolic geometry [STT88] and combinatorial properties of Thompson's groups [Deh10]. Recently, L. Pournin finally gave a purely combinatorial proof that the diameter of the n -dimensional associahedron is precisely $2n - 4$ as soon as $n \geq 9$ [Pou14]. On the other hand, J. Lucas [Luc87] proved that the flip graph is Hamiltonian. Later, F. Hurtado and M. Noy [HN99] obtained a simpler proof of this result, using a hierarchy of triangulations which organizes all triangulations of convex polygons into an infinite generating tree.

Generalizing the classical associahedron, M. Carr and S. Devadoss [CD06, Dev09] defined and constructed *graph associahedra*. For a finite graph G , a G -*associahedron* $\text{Asso}(G)$ is a simple convex polytope whose combinatorial structure encodes the connected subgraphs of G and their nested structure. To be more precise, the face lattice of the polar of a G -associahedron is isomorphic to the *nested complex* on G , defined as the simplicial complex of all collections of tubes (vertex subsets inducing connected subgraphs) of G which are pairwise either nested, or disjoint and non-adjacent. See Figures 1 and 3 for 3-dimensional examples. The graph associahedra of certain special families of graphs happen to coincide with well-known families of polytopes (see Figure 3):

VP was partially supported by grant MTM2011-22792 of the Spanish MICINN and by the French ANR grant EGOS (12 JS02 002 01).




FIGURE 1. The 3-dimensional associahedron and the graph associahedron of the tripod.

classical associahedra are path associahedra, cyclohedra are cycle associahedra, and permutahedra are complete graph associahedra. Graph associahedra have been geometrically realized in different ways: by successive truncations of faces of the standard simplex [CD06], as Minkowski sums of faces of the standard simplex [Pos09, FS05], or from their normal fans by exhibiting explicit inequality descriptions [Zel06]. However, we do not consider these geometric realizations as we focus on the combinatorial properties of the nested complex.

Given a finite simple graph G , we denote by $\mathcal{F}(G)$ the 1-skeleton of the graph associahedron $\text{Asso}(G)$. In other words, $\mathcal{F}(G)$ is the facet-ridge graph of the nested complex on G . Its vertices are maximal tubings on G and its edges connect tubings which differ only by two tubes. See Section 2 for precise definitions and examples. In this paper, we study graph properties of $\mathcal{F}(G)$. In Section 3, we focus on the diameter $\delta(\mathcal{F}(G))$ of the flip graph $\mathcal{F}(G)$. We obtain the following structural results.

Theorem 1. *The diameter $\delta(\mathcal{F}(G))$ of the flip graph $\mathcal{F}(G)$ is non-decreasing: $\delta(\mathcal{F}(G)) \leq \delta(\mathcal{F}(G'))$ for any two graphs G, G' such that $G \subseteq G'$.*

Related to this diameter, we investigate the *non-leaving-face property*: do all geodesics between two vertices of a face F of $\text{Asso}(G)$ stay in F ? This property was proved for the classical associahedron in [STT88] but the name was coined in [CP16]. Although not all faces of the graph associahedron $\text{Asso}(G)$ fulfill this property, we prove in the following statement that some of them do.

Proposition 2. *Any tubing on a geodesic between two tubings \mathbb{T} and \mathbb{T}' in the flip graph $\mathcal{F}(G)$ contains any common upper set to the inclusion posets of \mathbb{T} and \mathbb{T}' .*

In fact, we extend Theorem 1 and Proposition 2 to all *nestohedra* [Pos09, FS05], see Section 3. Finally, using Theorem 1 and Proposition 2, the lower bound on the diameter of the associahedron [Pou14], the usual construction of graph associahedra [CD06, Pos09] and the diameter of graphical zonotopes, we obtain the following inequalities on the diameter $\delta(\mathcal{F}(G))$ of $\mathcal{F}(G)$.

Theorem 3. *For any connected graph G with $n + 1$ vertices and e edges, the diameter $\delta(\mathcal{F}(G))$ of the flip graph $\mathcal{F}(G)$ is bounded by*

$$\max(e, 2n - 18) \leq \delta(\mathcal{F}(G)) \leq \binom{n + 1}{2}.$$

In Section 4, we study the Hamiltonicity of $\mathcal{F}(G)$. Based on an inductive decomposition of graph associahedra, we show the following statement.

Theorem 4. *For any graph G with at least two edges, the flip graph $\mathcal{F}(G)$ is Hamiltonian.*

2. PRELIMINARIES

2.1. Tubings, nested complex, and graph associahedron. Let V be an $(n + 1)$ -elements ground set, and let G be a simple graph on V with $\pi_0(G)$ connected components. We denote by $G[U]$ the subgraph of G induced by a subset U of V .

A *tube* of G is a non-empty subset t of V that induces a connected subgraph of G . A tube is *proper* if it does not induce a connected component of G . The set of all tubes of G is called the *graphical building set* of G and denoted by $B(G)$. We moreover denote by $B(G)_{\max}$ the set of inclusion maximal tubes of $B(G)$, *i.e.* the vertex sets of connected components of G .

Two tubes t and t' are *compatible* if they are

- nested, *i.e.* $t \subseteq t'$ or $t' \subseteq t$, or
- disjoint and non-adjacent, *i.e.* $t \cup t'$ is not a tube of G .

A *tubing* on G is a set of pairwise compatible tubes of G . A tubing is *proper* if it contains only proper tubes and *loaded* if it contains $B(G)_{\max}$. Since inclusion maximal tubes are compatible with all tubes, we can transform any tubing T into a proper tubing $T \setminus B(G)_{\max}$ or into a loaded tubing $T \cup B(G)_{\max}$, and we switch along the paper to whichever version suits better the current purpose. Observe by the way that maximal tubings are automatically loaded. Figure 2 illustrates these notions on a graph with 9 vertices.




FIGURE 2. A proper tube, a tubing, a maximal proper tubing, and a maximal (loaded) tubing.

The *nested complex* on G is the simplicial complex $\mathcal{N}(G)$ of all proper tubings on G . This complex is known to be the boundary complex of the *graph associahedron* $\text{Asso}(G)$, which is an $(n+1-\pi_0(G))$ -dimensional simple polytope. This polytope was first constructed in [CD06, Dev09]¹ and later in the more general context of nestohedra in [Pos09, FS05, Zel06]. In this paper, we do not need these geometric realizations since we only consider combinatorial properties of the nested complex $\mathcal{N}(G)$. In fact, we focus on the *flip graph* $\mathcal{F}(G)$ whose vertices are maximal proper tubings on G and whose edges connect adjacent maximal proper tubings, *i.e.* which only differ by two tubes. We refer to Figure 4 for an example, and to Section 2.3 for a description of flips. To avoid confusion, we always use the term *edge* for the edges of the graph G , and the term *flip* for the edges of the flip graph $\mathcal{F}(G)$. To simplify the presentation, it is sometimes more convenient to consider the *loaded flip graph*, obtained from $\mathcal{F}(G)$ by loading all its vertices with $B(G)_{\max}$, and still denoted by $\mathcal{F}(G)$. Note that only proper tubes can be flipped in each maximal tubing on the loaded flip graph.

Observe that if G is disconnected with connected components G_i , for $i \in [\pi_0(G)]$, then the nested complex $\mathcal{N}(G)$ is the join of the nested complexes $\mathcal{N}(G_i)$, the graph associahedron $\text{Asso}(G)$ is the Cartesian product of the graph associahedra $\text{Asso}(G_i)$, and the flip graph $\mathcal{F}(G)$ is the Cartesian

¹The definition used in [CD06, Dev09] slightly differs from ours for disconnected graphs, but our results still hold in their framework.

product of the flip graphs $\mathcal{F}(G_i)$. In many places, this allows us to restrict our arguments to connected graphs.

Example 5 (Classical polytopes). For certain families of graphs, the graph associahedra turn out to coincide (combinatorially) with classical polytopes (see Figure 3):

- (i) the path associahedron $\text{Asso}(P_{n+1})$ coincides with the n -dimensional associahedron,
- (ii) the cycle associahedron $\text{Asso}(O_{n+1})$ coincides with the n -dimensional cyclohedron,
- (iii) the complete graph associahedron $\text{Asso}(K_{n+1})$ coincides with the n -dimensional permutahedron $\text{Perm}(n) := \text{conv}\{(\sigma(1), \dots, \sigma(n+1)) \mid \sigma \in \mathfrak{S}_{n+1}\}$.




FIGURE 3. The associahedron, the cyclohedron, and the permutahedron are graph associahedra.

2.2. Spines. Spines provide convenient representations of the tubings on G . Given a tubing T on G , the corresponding *spine* S is the Hasse diagram of the inclusion poset on $T \cup B(G)_{\max}$, where the node corresponding to a tube $t \in T \cup B(G)_{\max}$ is labeled by $\lambda(t, T) := t \setminus \bigcup \{t' \in T \mid t' \subsetneq t\}$. See Figure 4.

The compatibility condition on the tubes of T implies that the spine S is a rooted forest, where roots correspond to elements of $B(G)_{\max}$. Spines are in fact called $B(G)$ -forests in [Pos09]. The labels of S define a partition of the vertex set of G . The tubes of $T \cup B(G)_{\max}$ are the descendant sets $\text{desc}(s, S)$ of the nodes s of the forest S , where $\text{desc}(s, S)$ denotes the union of the labels of the descendants of s in S , including s itself. The tubing $T \cup B(G)_{\max}$ is maximal if and only if all labels are singletons, and we then identify nodes with their labels, see Figure 4.

Let T, \bar{T} be tubings on G with corresponding spines S, \bar{S} . Then $\bar{T} \subseteq T$ if and only if \bar{S} is obtained from S by edge contractions. We say that S *refines* \bar{S} , that \bar{S} *coarsens* S , and we write $\bar{S} \prec S$. Given any node s of S , we denote by S_s the *subspine* of S induced by all descendants of s in S , including s itself.

2.3. Flips. As already mentioned, the nested complex $\mathcal{N}(G)$ is a simplicial sphere. It follows that there is a natural flip operation on maximal proper tubings on G . Namely, for any maximal proper tubing T on G and any tube $t \in T$, there exists a unique proper tube $t' \notin T$ of G such that $T' := T \Delta \{t, t'\}$ is again a proper tubing on G (where Δ denotes the symmetric difference operator). We denote this flip by $T \leftrightarrow T'$. This flip operation can be explicitly described both in terms of tubings and spines as follows.

Consider a tube t in a maximal proper tubing T , with $\lambda(t, T) = \{v\}$. Let \bar{t} denote the smallest element of $T \cup B(G)_{\max}$ strictly containing t , and denote its label by $\lambda(\bar{t}, T) = \{v'\}$. Then the unique tube t' such that $T' := T \Delta \{t, t'\}$ is again a proper tubing on G is the connected component of the induced subgraph $G[\bar{t} \setminus \{v\}]$ containing v' . See Figure 4.

This description translates to spines as follows. The flip between the tubings T and T' corresponds to a rotation between the corresponding spines S and S' . This operation is local: it only perturbs the nodes v and v' and their children. More precisely, v is a child of v' in S , and becomes the parent of v' in S' . Moreover, the children of v in S contained in t' become children of v' in S' . All other nodes keep their parents. See Figure 4.

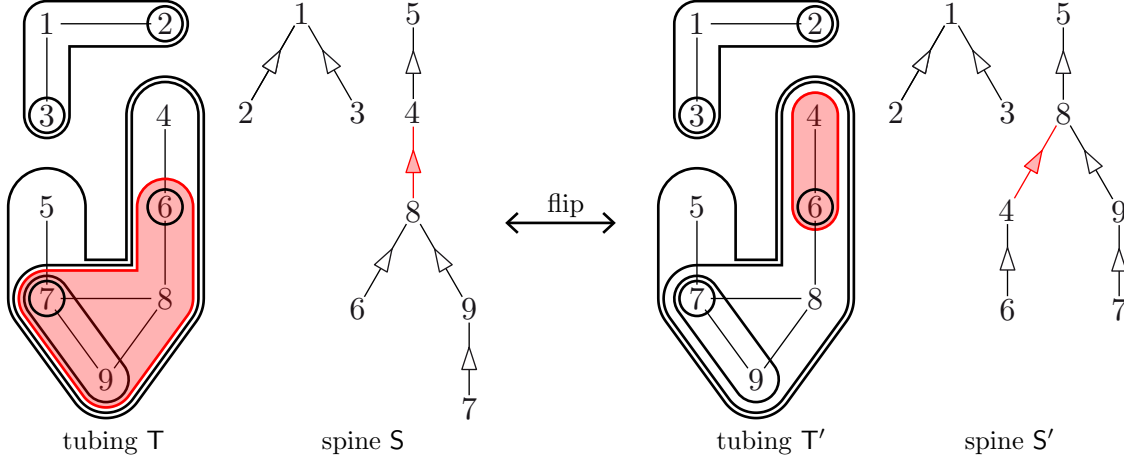


FIGURE 4. The flip of a proper tube (shaded, red) in a maximal tubing seen both on the tubings and on the corresponding spines.

3. DIAMETER

Let $\delta(\mathcal{F}(G))$ denote the diameter of the flip graph $\mathcal{F}(G)$. For example, for the complete graph K_{n+1} , the diameter of the n -dimensional permutahedron is $\delta(\mathcal{F}(K_{n+1})) = \binom{n+1}{2}$, while for the path P_{n+1} , the diameter of the classical n -dimensional associahedron is $\delta(\mathcal{F}(P_{n+1})) = 2n - 4$ for $n > 9$, by results of [STT88, Pou14]. We discuss in this section properties of the diameter $\delta(\mathcal{F}(G))$ and of the geodesics in the flip graph $\mathcal{F}(G)$. The results of Section 3.1 are extended to nestohedra in Section 3.2. We prefer to present the ideas first on graph associahedra as they prepare the intuition for the more technical proofs on nestohedra.

3.1. Non-decreasing diameters. Our first goal is to show that $\delta(\mathcal{F}(\cdot))$ is non-decreasing.

Theorem 6. $\delta(\mathcal{F}(\bar{G})) \leq \delta(\mathcal{F}(G))$ for any two graphs G, \bar{G} such that $\bar{G} \subseteq G$.

Remark 7. We could prove this statement by a geometric argument, using the construction of the graph associahedron of M. Carr and S. Devadoss [CD06]. Indeed, it follows from [CD06] that the graph associahedron $\text{Asso}(G)$ can be obtained from the graph associahedron $\text{Asso}(\bar{G})$ by successive face truncations. Geometrically, this operation replaces the truncated face F by its Cartesian product with a simplex of codimension $\dim(F) + 1$. Therefore, a path in the graph of $\text{Asso}(G)$ naturally projects to a shorter path in the graph of $\text{Asso}(\bar{G})$. Our proof is a purely combinatorial translation of this geometric intuition. It has the advantage not to rely on the results of [CD06] and to help formalizing the argument.

Observe first that deleting an isolated vertex in G does not change the nested complex $\mathcal{N}(G)$. We can thus assume that the graphs G and \bar{G} have the same vertex set and that $\bar{G} = G \setminus \{(u, v)\}$ is obtained by deleting a single edge (u, v) from G . We define below a map Ω from tubings on G to tubings on \bar{G} which induces a surjection from the flip graph $\mathcal{F}(G)$ onto the flip graph $\mathcal{F}(\bar{G})$. For consistency, we use t and T for tubes and tubings of G and \bar{t} and \bar{T} for tubes and tubings of \bar{G} .

Given a tube t of G (proper or not), define $\Omega(t)$ to be the coarsest partition of t into tubes of \bar{G} . In other words, $\Omega(t) = \{t\}$ if (u, v) is not an isthmus of $G[t]$, and otherwise $\Omega(t) = \{\bar{t}_u, \bar{t}_v\}$ where \bar{t}_u and \bar{t}_v are the vertex sets of the connected components of $\bar{G}[t]$ containing u and v respectively. For a set of tubes T of G , define $\Omega(T) := \bigcup_{t \in T} \Omega(t)$. See Figure 5 for an illustration.

Lemma 8. For any tubing T on G , the set $\Omega(T)$ is a tubing on \bar{G} and $|T| \leq |\Omega(T)|$.

Proof. It is immediate to see that Ω sends tubings on G to tubings on \bar{G} . We prove by induction on $|T|$ that $|T| \leq |\Omega(T)|$. Consider a non-empty tubing T , and let t be an inclusion maximal tube of T . By induction hypothesis, $|T \setminus \{t\}| \leq |\Omega(T \setminus \{t\})|$. We now distinguish two cases:




FIGURE 5. Two maximal tubings (left and middle) with the same image by the map Ω (right). The middle tubing is the preimage of the rightmost tubing obtained by the process described in the proof of Corollary 9 with $u = 6$ and $v = 8$.

- (i) If (u, v) is an isthmus of $G[t]$, then $\Omega(t) = \{\bar{t}_u, \bar{t}_v\} \not\subseteq \Omega(\mathbb{T} \setminus \{t\})$. Indeed, since \bar{t}_u and \bar{t}_v are adjacent in G , two tubes of \mathbb{T} whose images by Ω produce \bar{t}_u and \bar{t}_v must be nested. Therefore, one of them contains both \bar{t}_u and \bar{t}_v , and thus equals $t = \bar{t}_u \cup \bar{t}_v$ by maximality of t in \mathbb{T} .
- (ii) If (u, v) is not an isthmus of $G[t]$, then $\Omega(t) = \{t\} \not\subseteq \Omega(\mathbb{T} \setminus \{t\})$. Indeed, if $t' \in \mathbb{T}$ is such that $t \in \Omega(t')$, then $t \subseteq t'$ and thus $t = t'$ by maximality of t in \mathbb{T} .

We conclude that $|\Omega(\mathbb{T})| \geq |\Omega(\mathbb{T} \setminus \{t\})| + 1 \geq |\mathbb{T} \setminus \{t\}| + 1 = |\mathbb{T}|$. \square

Corollary 9. *The map Ω induces a graph surjection from the loaded flip graph $\mathcal{F}(G)$ onto the loaded flip graph $\mathcal{F}(\bar{G})$, i.e. a surjective map from maximal tubings on G to maximal tubings on \bar{G} such that adjacent tubings on G are sent to identical or adjacent tubings on \bar{G} .*

Proof. Let $\bar{\mathbb{T}}$ be a tubing on \bar{G} . If all tubes of $\bar{\mathbb{T}}$ containing u also contain v (or the opposite), then $\bar{\mathbb{T}}$ is a tubing on G and $\Omega(\bar{\mathbb{T}}) = \bar{\mathbb{T}}$. Otherwise, let $\bar{\mathbb{T}}_u$ denote the set of tubes of $\bar{\mathbb{T}}$ containing u but not v and \bar{t}_v denote the maximal tube containing v but not u . Then $(\bar{\mathbb{T}} \setminus \bar{\mathbb{T}}_u) \cup \{\bar{t}_u \cup \bar{t}_v \mid \bar{t}_u \in \bar{\mathbb{T}}_u\}$ is a tubing on G whose image by Ω is $\bar{\mathbb{T}}$. See Figure 5 for an illustration. The map Ω is thus surjective from tubings on G to tubings on \bar{G} . Moreover, any preimage \mathbb{T}_o of a maximal tubing $\bar{\mathbb{T}}$ can be completed into a maximal tubing \mathbb{T} with $\Omega(\mathbb{T}) \supseteq \Omega(\mathbb{T}_o) = \bar{\mathbb{T}}$, and thus satisfying $\Omega(\mathbb{T}) = \bar{\mathbb{T}}$ by maximality of \mathbb{T} .

Remember that two distinct maximal tubings on G are adjacent if and only if they share precisely $|V| - 1$ common tubes. Consider two adjacent maximal tubings \mathbb{T}, \mathbb{T}' on G , so that $|\mathbb{T} \cap \mathbb{T}'| = |V| - 1$. Since $\Omega(\mathbb{T} \cap \mathbb{T}') \subseteq \Omega(\mathbb{T}) \cap \Omega(\mathbb{T}')$ and $|\Omega(\mathbb{T} \cap \mathbb{T}')| \geq |\mathbb{T} \cap \mathbb{T}'|$ by Lemma 8, we have $|\Omega(\mathbb{T}) \cap \Omega(\mathbb{T}')| \geq |\mathbb{T} \cap \mathbb{T}'| = |V| - 1$. Therefore, the tubings $\Omega(\mathbb{T}), \Omega(\mathbb{T}')$ are adjacent if $|\Omega(\mathbb{T}) \cap \Omega(\mathbb{T}')| = |\mathbb{T} \cap \mathbb{T}'|$ and identical if $|\Omega(\mathbb{T}) \cap \Omega(\mathbb{T}')| > |\mathbb{T} \cap \mathbb{T}'|$. \square

Remark 10. We can in fact precisely describe the preimage $\Omega^{-1}(\bar{\mathbb{T}})$ of a maximal tubing $\bar{\mathbb{T}}$ on \bar{G} as follows. As in the previous proof, let $\bar{\mathbb{T}}_u$ denote the chain of tubes of $\bar{\mathbb{T}}$ containing u but not v and similarly $\bar{\mathbb{T}}_v$ denote the chain of tubes of $\bar{\mathbb{T}}$ containing v but not u . Any linear extension L of these two chains defines a preimage of $\bar{\mathbb{T}}$ where the tubes of $\bar{\mathbb{T}}_u \cup \bar{\mathbb{T}}_v$ are replaced by the tubes $\bigcup \{t' \in L \mid t' \leq_L t\}$ for $t \in L$. In terms of spines, this translates to shuffling the two chains corresponding to $\bar{\mathbb{T}}_u$ and $\bar{\mathbb{T}}_v$. Details are left to the reader.

Proof of Theorem 6. Consider two maximal tubings $\bar{\mathbb{T}}, \bar{\mathbb{T}}'$ on \bar{G} . Let \mathbb{T}, \mathbb{T}' be maximal loaded tubings on G such that $\Omega(\mathbb{T}) = \bar{\mathbb{T}}$ and $\Omega(\mathbb{T}') = \bar{\mathbb{T}}'$ (surjectivity of Ω), and $\mathbb{T} = \mathbb{T}_0, \dots, \mathbb{T}_\ell = \mathbb{T}'$ be a geodesic between them ($\ell \leq \delta(\mathcal{F}(G))$). Deleting repetitions in the sequence $\bar{\mathbb{T}} = \Omega(\mathbb{T}_0), \dots, \Omega(\mathbb{T}_\ell) = \bar{\mathbb{T}}'$ yields a path from $\bar{\mathbb{T}}$ to $\bar{\mathbb{T}}'$ (Corollary 9) of length at most $\ell \leq \delta(\mathcal{F}(G))$. So $\delta(\mathcal{F}(G)) \geq \delta(\mathcal{F}(\bar{G}))$. \square

3.2. Extension to nestohedra. The results of the previous section can be extended to the nested complex on an arbitrary building set. Although the proofs are more abstract and technical, the ideas behind are essentially the same. We recall the definitions of building set and nested complex needed here and refer to [CD06, Pos09, FS05, Zel06] for more details and motivation.

A *building set* on a ground set V is a collection \mathcal{B} of non-empty subsets of V such that

- (B1) if $\mathbf{b}, \mathbf{b}' \in \mathcal{B}$ and $\mathbf{b} \cap \mathbf{b}' \neq \emptyset$, then $\mathbf{b} \cup \mathbf{b}' \in \mathcal{B}$, and
- (B2) \mathcal{B} contains all singletons $\{v\}$ for $v \in V$.

We denote by \mathcal{B}_{\max} the set of inclusion maximal elements of \mathcal{B} and call *proper* the elements of $\mathcal{B} \setminus \mathcal{B}_{\max}$. The building set is *connected* if $\mathcal{B}_{\max} = \{V\}$. Graphical building sets are particular examples, and connected graphical building sets correspond to connected graphs.

A *B-nested set* on \mathcal{B} is a subset \mathcal{N} of \mathcal{B} such that

- (N1) for any $n, n' \in \mathcal{N}$, either $n \subseteq n'$ or $n' \subseteq n$ or $n \cap n' = \emptyset$, and
- (N2) for any $k \geq 2$ pairwise disjoint sets $n_1, \dots, n_k \in \mathcal{N}$, the union $n_1 \cup \dots \cup n_k$ is not in \mathcal{B} .

As before, a \mathcal{B} -nested set \mathcal{N} is *proper* if $\mathcal{N} \cap \mathcal{B}_{\max} = \emptyset$ and *loaded* if $\mathcal{B}_{\max} \subseteq \mathcal{N}$. The *B-nested complex* is the $(|V| - |\mathcal{B}_{\max}|)$ -dimensional simplicial complex $\mathcal{N}(\mathcal{B})$ of all proper nested sets on \mathcal{B} . As in the graphical case, the \mathcal{B} -nested complex can be realized geometrically as the boundary complex of the polar of the *nestohedron* $\text{Nest}(\mathcal{B})$, constructed *e.g.* in [Pos09, FS05, Zel06]. We denote by $\delta(\mathcal{F}(\mathcal{B}))$ the diameter of the graph $\mathcal{F}(\mathcal{B})$ of $\text{Nest}(\mathcal{B})$. As in the previous section, it is more convenient to regard the vertices of $\mathcal{F}(\mathcal{B})$ as maximal loaded nested sets.

The *spine* of a nested set \mathcal{N} is the Hasse diagram of the inclusion poset of $\mathcal{N} \cup \mathcal{B}_{\max}$. Spines are called \mathcal{B} -forests in [Pos09]. The definitions and properties of Section 2.2 extend to general building sets, see [Pos09] for details.

We shall now prove the following generalization of Theorem 6.

Theorem 11. $\delta(\mathcal{F}(\bar{\mathcal{B}})) \leq \delta(\mathcal{F}(\mathcal{B}))$ for any two building sets $\mathcal{B}, \bar{\mathcal{B}}$ on V such that $\bar{\mathcal{B}} \subseteq \mathcal{B}$.

The proof follows the same line as that of Theorem 6. We first define a map Ω which transforms elements of \mathcal{B} to subsets of $\bar{\mathcal{B}}$ as follows: for $\mathbf{b} \in \mathcal{B}$ (proper or not), define $\Omega(\mathbf{b})$ as the coarsest partition of \mathbf{b} into elements of $\bar{\mathcal{B}}$. Observe that $\Omega(\mathbf{b})$ is well-defined since $\bar{\mathcal{B}}$ is a building set, and that the elements of $\Omega(\mathbf{b})$ are precisely the inclusion maximal elements of $\bar{\mathcal{B}}$ contained in \mathbf{b} . For a nested set \mathcal{N} on \mathcal{B} , we define $\Omega(\mathcal{N}) := \bigcup_{n \in \mathcal{N}} \Omega(n)$. The following statement is similar to Lemma 8.

Lemma 12. For any nested set \mathcal{N} on \mathcal{B} , the image $\Omega(\mathcal{N})$ is a nested set on $\bar{\mathcal{B}}$ and $|\mathcal{N}| \leq |\Omega(\mathcal{N})|$.

Proof. Consider a nested set \mathcal{N} on \mathcal{B} . To prove that $\Omega(\mathcal{N})$ is a nested set on $\bar{\mathcal{B}}$, we start with condition (N1). Let $\bar{n}, \bar{n}' \in \Omega(\mathcal{N})$ and let $n, n' \in \mathcal{N}$ such that $\bar{n} \in \Omega(n)$ and $\bar{n}' \in \Omega(n')$. Since \mathcal{N} is nested, we can distinguish two cases:

- Assume that n and n' are disjoint. Then $\bar{n} \cap \bar{n}' = \emptyset$ since $\bar{n} \subseteq n$ and $\bar{n}' \subseteq n'$.
- Assume that n and n' are nested, *e.g.* $n \subseteq n'$. If $\bar{n} \cap \bar{n}' \neq \emptyset$, then $\bar{n} \cup \bar{n}'$ is in $\bar{\mathcal{B}}$ and is a subset of n' . By maximality of \bar{n}' in n' , we obtain $\bar{n} \cup \bar{n}' = \bar{n}'$, and thus $\bar{n} \subseteq \bar{n}'$.

To prove Condition (N2), consider pairwise disjoint elements $\bar{n}_1, \dots, \bar{n}_k \in \Omega(\mathcal{N})$ and $n_1, \dots, n_k \in \mathcal{N}$ such that $\bar{n}_i \in \Omega(n_i)$. We assume by contradiction that $\bar{n} := \bar{n}_1 \cup \dots \cup \bar{n}_k \in \bar{\mathcal{B}}$ and we prove that $n := n_1 \cup \dots \cup n_k \in \mathcal{B}$. Indeed, \bar{n}, n_1, \dots, n_k all belong to \mathcal{B} and $\bar{n} \cap n_i \neq \emptyset$ (it contains \bar{n}_i) so that $\bar{n} \cup n$ also belongs to \mathcal{B} by multiple applications of Property (B1) of building sets. Moreover, $\bar{n} \subseteq n$ so that $n = \bar{n} \cup n \in \mathcal{B}$. Finally, we conclude distinguishing two cases:

- If there is $i \in [k]$ such that n_i contains all n_j , then \bar{n}_i contains all \bar{n}_j and thus \bar{n} . This contradicts the maximality of \bar{n}_i in n_i since $\bar{n}_i \subsetneq \bar{n} \in \bar{\mathcal{B}}$.
- Otherwise, merging intersecting elements allows us to assume that n_1, \dots, n_k are pairwise disjoint and $n := n_1 \cup \dots \cup n_k \in \mathcal{B}$ contradicts Condition (N2) for \mathcal{N} .

This concludes the proof that $\Omega(\mathcal{N})$ is a nested set on $\bar{\mathcal{B}}$.

We now prove that $|\mathcal{N}| \leq |\Omega(\mathcal{N})|$ by induction on $|\mathcal{N}|$. Consider a non-empty nested set \mathcal{N} and let n_{\circ} be an inclusion maximal element of \mathcal{N} . By induction hypothesis, $|\mathcal{N} \setminus \{n_{\circ}\}| \leq |\Omega(\mathcal{N} \setminus \{n_{\circ}\})|$. Let $\Omega(n_{\circ}) = \{\bar{n}_1, \dots, \bar{n}_k\}$. Consider $n_1, \dots, n_k \in \mathcal{N}$ such that $\bar{n}_i \in \Omega(n_i)$, and let $n := n_1 \cup \dots \cup n_k$. Since $n_{\circ}, n_1, \dots, n_k$ all belong to \mathcal{B} and $n_{\circ} \cap n_i \neq \emptyset$ (it contains \bar{n}_i), we have $n_{\circ} \cup n \in \mathcal{B}$ by multiple

applications of Property (B1) of building sets. Moreover, $\mathfrak{n}_o \subseteq \mathfrak{n}$ so that $\mathfrak{n} = \mathfrak{n}_o \cup \mathfrak{n} \in \mathcal{B}$. It follows by Condition (N2) on \mathcal{N} that there is $i \in [k]$ such that \mathfrak{n}_i contains all \mathfrak{n}_j , and thus $\mathfrak{n}_o \subseteq \mathfrak{n}_i$. We obtain that $\mathfrak{n}_o = \mathfrak{n}_i$ by maximality of \mathfrak{n}_o . We conclude that \mathfrak{n}_o is the only element of \mathcal{N} such that $\bar{\mathfrak{n}}_i \in \Omega(\mathfrak{n}_o)$, so that $|\Omega(\mathcal{N})| \geq |\Omega(\mathcal{N} \setminus \{\mathfrak{n}_o\})| + 1 \geq |\mathcal{N} \setminus \{\mathfrak{n}_o\}| + 1 = |\mathcal{N}|$. \square

Corollary 13. *The map Ω induces a graph surjection from the loaded flip graph $\mathcal{F}(\mathcal{B})$ onto the loaded flip graph $\mathcal{F}(\bar{\mathcal{B}})$, i.e. a surjective map from maximal nested sets on \mathcal{B} to maximal nested sets on $\bar{\mathcal{B}}$ such that adjacent nested sets on \mathcal{B} are sent to identical or adjacent nested sets on $\bar{\mathcal{B}}$.*

Proof. To prove the surjectivity, consider a nested set $\bar{\mathcal{N}}$ on $\bar{\mathcal{B}}$. The elements of $\bar{\mathcal{N}}$ all belong to \mathcal{B} and satisfy Condition (N1) for nested sets. It remains to transform the elements in $\bar{\mathcal{N}}$ which violate Condition (N2). If there is no such violation, then $\bar{\mathcal{N}}$ is a nested set on \mathcal{B} and $\Omega(\bar{\mathcal{N}}) = \bar{\mathcal{N}}$. Otherwise, consider pairwise disjoint elements $\bar{\mathfrak{n}}_1, \dots, \bar{\mathfrak{n}}_k$ of $\bar{\mathcal{N}}$ such that $\mathfrak{n} := \bar{\mathfrak{n}}_1 \cup \dots \cup \bar{\mathfrak{n}}_k$ is in \mathcal{B} and is maximal for this property. Consider the subset $\bar{\mathcal{N}}' := (\bar{\mathcal{N}} \setminus \{\bar{\mathfrak{n}}_1\}) \cup \{\mathfrak{n}\}$ of \mathcal{B} . Observe that:

- $\bar{\mathcal{N}}'$ still satisfies Condition (N1). Indeed, if $\bar{\mathfrak{n}} \in \bar{\mathcal{N}}$ is such that $\mathfrak{n} \cap \bar{\mathfrak{n}} \neq \emptyset$, then $\bar{\mathfrak{n}}$ intersects at least one element $\bar{\mathfrak{n}}_i$. Since $\bar{\mathcal{N}}$ is nested, $\bar{\mathfrak{n}} \subseteq \bar{\mathfrak{n}}_i$ or $\bar{\mathfrak{n}}_i \subseteq \bar{\mathfrak{n}}$. In the former case, $\bar{\mathfrak{n}} \subseteq \mathfrak{n}$ and we are done. In the latter case, $\bar{\mathfrak{n}}$ and the elements $\bar{\mathfrak{n}}_j$ disjoint from $\bar{\mathfrak{n}}$ would contradict the maximality of \mathfrak{n} .
- $\bar{\mathcal{N}}'$ still satisfies $\Omega(\bar{\mathcal{N}}') = \bar{\mathcal{N}}$. Indeed, $\bar{\mathfrak{n}}_1 \in \Omega(\mathfrak{n})$ since $\Omega(\mathfrak{n}) = \{\bar{\mathfrak{n}}_1, \dots, \bar{\mathfrak{n}}_k\}$. For the latter equality, observe that $\{\bar{\mathfrak{n}}_1, \dots, \bar{\mathfrak{n}}_k\}$ is a partition of \mathfrak{n} into elements of $\bar{\mathcal{B}}$ and that a coarser partition would contradict Condition (N2) on $\bar{\mathcal{N}}$.
- \mathfrak{n} cannot be partitioned into two or more elements of $\bar{\mathcal{N}}'$. Such a partition would refine the partition $\Omega(\mathfrak{n})$, and would thus contradict again Condition (N2) on $\bar{\mathcal{N}}$. Therefore, $\bar{\mathcal{N}}'$ has strictly less violations of Condition (N2) than $\bar{\mathcal{N}}$.
- All violations of Condition (N2) in $\bar{\mathcal{N}}'$ only involve elements of $\bar{\mathcal{B}}$. Indeed, pairwise disjoint elements $\bar{\mathfrak{n}}'_1, \dots, \bar{\mathfrak{n}}'_\ell \in \bar{\mathcal{N}}'$ disjoint from \mathfrak{n} and such that $\mathfrak{n} \cup \bar{\mathfrak{n}}'_1 \cup \dots \cup \bar{\mathfrak{n}}'_\ell \in \mathcal{B}$ would contradict the maximality of \mathfrak{n} .

These four points enable us to decrease the number of violations of Condition (N2) until we reach a nested set \mathcal{N} on \mathcal{B} which still satisfies $\Omega(\mathcal{N}) = \bar{\mathcal{N}}$.

The second part of the proof is identical to that of Corollary 9. \square

From Corollary 13, the proof of Theorem 11 is identical to that of Theorem 6.

3.3. Geodesic properties. In this section, we focus on properties of the geodesics in the graphs of nestohedra. We consider three properties for a face F of a polytope P :

NLFP: F has the *non-leaving-face property* in P if F contains all geodesics connecting two vertices of F in the graph of P .

SNLFP: F has the *strong non-leaving-face property* in P if any path connecting two vertices v, w of F in the graph of P and leaving the face F has at least two more steps than a geodesic between v and w .

EFP: F has the *entering-face property* in P if for any vertices u, v, w of P such that $u \notin F$, $v, w \in F$, and u and v are neighbors in the graph of P , there exists a geodesic connecting u and w whose first edge is the edge from u to v .

For a face F of a polytope P , we have $\text{EFP} \iff \text{SNLFP} \implies \text{NLFP}$. However, the reverse of the last implication is wrong: all faces of a simplex have the NLFP (all vertices are at distance 1), but not the SNLFP. Alternative counter-examples with no simplicial face already exist in dimension 3. Among classical polytopes the n -dimensional cube, permutahedron, associahedron, and cyclohedron all satisfy the EFP. The NLFP is further discussed in [CP16].

Contrarily to the classical associahedron, not all faces of a graph associahedron have the NLFP. A counter-example is given by the star with n branches: Figure 6 shows a path of length $2n$ between two maximal tubings \mathbb{T}, \mathbb{T}' , while the minimal face containing \mathbb{T} and \mathbb{T}' is an $(n-1)$ -dimensional permutahedron (see the face description in [CD06, Theorem 2.9]) and the graph distance between \mathbb{T} and \mathbb{T}' in this face is $\binom{n}{2}$. It turns out however that the following faces of the graph associahedra, and more generally of nestohedra, always have the SNLFP.

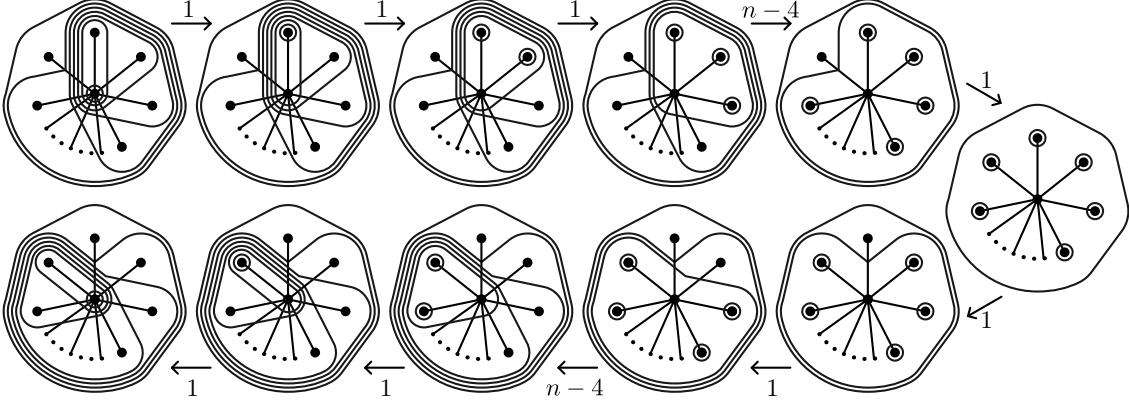


FIGURE 6. A geodesic (of length $2n$) between two maximal tubings of the star that flips their common tube (the central vertex).

Lemma 14. We call **upper ideal face** of the nestohedron $\text{Nest}(\mathbb{B})$ a face corresponding to a loaded nested set \mathbb{N}^\uparrow that satisfies the following equivalent properties:

- (i) any element of \mathbb{B} not in \mathbb{N}^\uparrow but compatible with \mathbb{N}^\uparrow is contained in an inclusion minimal element of \mathbb{N}^\uparrow ,
- (ii) the set $\lambda(n, \mathbb{N}^\uparrow) := n \setminus \bigcup \{n' \in \mathbb{N}^\uparrow \mid n' \subsetneq n\}$ is a singleton for any inclusion non-minimal element n of \mathbb{N}^\uparrow ,
- (iii) the forest obtained by deleting all leaves of the spine S^\uparrow of \mathbb{N}^\uparrow forms an upper ideal of any spine refining S^\uparrow .

Proof. We first prove that (i) \implies (ii). Assume that $n \in \mathbb{N}^\uparrow$ is not inclusion minimal and that $\lambda(n, \mathbb{N}^\uparrow)$ contains two distinct elements $v, w \in V$. One can then check that the maximal element of \mathbb{B} contained in n and containing v but not w is compatible with \mathbb{N}^\uparrow , but not contained in an inclusion minimal element of \mathbb{N}^\uparrow . This proves that (i) \implies (ii).

Conversely, assume (ii) and consider $b \in \mathbb{B}$ not in \mathbb{N}^\uparrow but compatible with \mathbb{N}^\uparrow . Since \mathbb{N}^\uparrow is loaded, there exists $n \in \mathbb{N}^\uparrow$ strictly containing b and minimal for this property. Since b is compatible with \mathbb{N}^\uparrow , we obtain that $\lambda(n, \mathbb{N}^\uparrow)$ contains at least one element from b and one from $n \setminus b$, and is thus not a singleton. It follows by (ii) that n is an inclusion minimal element of \mathbb{N}^\uparrow , and it contains b .

The equivalence (ii) \iff (iii) follows directly from the definition of the spines and their labelings, and the fact that a non-singleton node in a spine can be split in a refining spine. \square

Proposition 15. Any upper ideal face of the nestohedron $\text{Nest}(\mathbb{B})$ satisfies SNLFP.

Proof. Consider an upper ideal face F of $\text{Nest}(\mathbb{B})$ corresponding to the loaded nested set \mathbb{N}^\uparrow . We consider the building set $\bar{\mathbb{B}} \subseteq \mathbb{B}$ on V consisting of all elements of \mathbb{B} (weakly) contained in an inclusion minimal element of \mathbb{N}^\uparrow together with all singletons $\{v\}$ for elements $v \in V$ not contained in any inclusion minimal element of \mathbb{N}^\uparrow . The reader is invited to check that $\bar{\mathbb{B}}$ is indeed a building set on V . It follows from Lemma 14 that

- $\lambda(n, \mathbb{N}^\uparrow) = n$ if n is an inclusion minimal element of \mathbb{N}^\uparrow ,
- $\lambda(n, \mathbb{N}^\uparrow) = \{v\}$ for some v not contained in any inclusion minimal element of \mathbb{N}^\uparrow otherwise,

and thus that the map $\lambda(\cdot, \mathbb{N}^\uparrow)$ is a bijection from \mathbb{N}^\uparrow to $\bar{\mathbb{B}}_{\max}$.

Consider the surjection Ω from the maximal nested sets on \mathbb{B} to the maximal nested sets on $\bar{\mathbb{B}}$ as defined in the previous section: $\Omega(\mathbb{N}) = \bigcup_{n \in \mathbb{N}} \Omega(n)$ where $\Omega(n)$ is the coarsest partition of n into elements of $\bar{\mathbb{B}}$. Following [STT88, CP16], we consider the *normalization* Ω^* on maximal nested sets on \mathbb{B} defined by $\Omega^*(\mathbb{N}) := (\Omega(\mathbb{N}) \setminus \bar{\mathbb{B}}_{\max}) \cup \mathbb{N}^\uparrow$. We claim that $\Omega^*(\mathbb{N})$ is a maximal nested set on \mathbb{B} :

- it is nested since both $\Omega(\mathbb{N}) \setminus \bar{\mathbb{B}}_{\max}$ and \mathbb{N}^\uparrow are themselves nested, and all elements of $\Omega(\mathbb{N}) \setminus \bar{\mathbb{B}}_{\max}$ are contained in a minimal element of \mathbb{N}^\uparrow .

- it is maximal since $\Omega(N)$ is maximal by Corollary 13 and $|\Omega^*(N)| = |\Omega(N)|$ because $\lambda(\cdot, N^\dagger)$ is a bijection from N^\dagger to \bar{B}_{\max} , and $\bar{B}_{\max} \subseteq \Omega(N)$ while $(\Omega(N) \setminus \bar{B}_{\max}) \cap N^\dagger = \emptyset$.

It follows that the map Ω^* combinatorially projects the nestohedron $\text{Nest}(\mathbf{B})$ onto its face F .

Let N_0, \dots, N_ℓ be a path in the loaded flip graph $\mathcal{F}(\mathbf{B})$ whose endpoints N_0, N_ℓ lie in the face F , but which leaves the face F . In other words, $N^\dagger \subseteq N_0, N_\ell$ and there are $0 \leq i < j \leq \ell$ such that $N^\dagger \subseteq N_i, N_j$ while $N^\dagger \not\subseteq N_{i+1}, N_{j-1}$. We claim that

$$\Omega^*(N_0) = N_0, \quad \Omega^*(N_\ell) = N_\ell, \quad \Omega^*(N_i) = N_i = \Omega^*(N_{i+1}) \quad \text{and} \quad \Omega^*(N_{j-1}) = N_j = \Omega^*(N_j),$$

so that the path $N_0 = \Omega^*(N_0), \dots, \Omega^*(N_\ell) = N_\ell$ from N_0 to N_ℓ in F has length at most $\ell - 2$ after deletion of repetitions.

To prove our claim, consider a loaded nested set N on \mathbf{B} containing a maximal proper nested set \bar{N} on $\bar{\mathbf{B}}$. Then $\Omega(N) \supseteq \Omega(\bar{N}) = \bar{N}$ so that $\Omega(N) = \bar{N} \cup \bar{B}_{\max}$ by maximality of \bar{N} . This shows $\Omega^*(N) = \bar{N} \cup N^\dagger$. In particular, if $N = \bar{N} \cup N^\dagger$, then $\Omega^*(N) = N$. Moreover, if N' is adjacent to $N = \bar{N} \cup N^\dagger$ and does not contain N^\dagger , then N' contains \bar{N} and $\Omega^*(N') = N$. This shows the claim and concludes the proof. \square

Proposition 15 specializes in particular to the non-leaving-face and entering face properties for the upper set faces of graph associahedra.

Proposition 16. (i) *If T and T' are two maximal tubings on G , then any maximal tubing on a geodesic between T and T' in the flip graph $\mathcal{F}(G)$ contains any common upper set to the inclusion posets of T and T' .*

(ii) *If T , T' and T'' are three maximal tubings on G such that $\mathsf{T} \setminus \{\mathfrak{t}\} = \mathsf{T}' \setminus \{\mathfrak{t}'\}$ and \mathfrak{t}' belongs to the maximal common upper set to the inclusion poset of T' and T'' , then there is a geodesic between T and T'' starting by the flip from T to T' .*

Proof. Using Proposition 15, it is enough to show that the maximal common upper set T^\dagger to the inclusion posets of T and T' defines an upper ideal face of $\text{Asso}(G)$. For this, we use the characterization (ii) of Lemma 14. Consider an inclusion non-minimal tube \mathfrak{t} of T^\dagger . Let \mathfrak{t}' be a maximal tube of T^\dagger such that $\mathfrak{t}' \subsetneq \mathfrak{t}$. Then \mathfrak{t}' has a unique neighbor v in $G[\mathfrak{t}]$ and all connected components of $G[\mathfrak{t} \setminus \{v\}]$ are both in T and T' , thus in T^\dagger . Thus $\lambda(\mathfrak{t}, \mathsf{T}^\dagger) = \{v\}$. \square

Remark 17. For an arbitrary building set \mathbf{B} , the maximal common upper set N^\dagger to the inclusion poset of two maximal nested sets N, N' is not always an upper ideal face of $\text{Nest}(\mathbf{B})$. A minimal example is the building set $\mathbf{B} = \{\{1\}, \{2\}, \{3\}, \{1, 2, 3\}\}$ and the nested sets $N = \{\{1\}, \{2\}, \{1, 2, 3\}\}$ and $N' = \{\{2\}, \{3\}, \{1, 2, 3\}\}$. Their maximal common upper set $N^\dagger = \{\{2\}, \{1, 2, 3\}\}$ is not an upper ideal face of $\text{Nest}(\mathbf{B})$ since $\lambda(\{1, 2, 3\}, N^\dagger) = \{1, 3\}$ is not a singleton. Moreover, the face corresponding to N^\dagger does not satisfy SNLFP.

3.4. Diameter bounds. Using Theorem 6 and Proposition 16, the lower bound on the diameter of the associahedron [Pou14], the classical construction of graph associahedra of [CD06, Pos09] and the diameter of graphical zonotopes, we obtain the inequalities on the diameter $\delta(\mathcal{F}(G))$ of $\mathcal{F}(G)$.

Theorem 18. *For any connected graph G with $n + 1$ vertices and e edges, the diameter $\delta(\mathcal{F}(G))$ of the flip graph $\mathcal{F}(G)$ is bounded by*

$$\max(e, 2n - 18) \leq \delta(\mathcal{F}(G)) \leq \binom{n+1}{2}.$$

Proof. For the upper bound, we use that the diameter is non-decreasing (Theorem 6) and that the n -dimensional permutahedron has diameter $\binom{n+1}{2}$, the maximal number of inversions in a permutation of \mathfrak{S}_{n+1} .

The lower bound consists in two parts. For the first part, we know that the normal fan of the graph associahedron $\text{Asso}(G)$ refines the normal fan of the graphical zonotope of G (see e.g. [Zie95, Lect. 7] for a reference on zonotopes). Indeed, the graph associahedron of G can be constructed as a Minkowski sum of the faces of the standard simplex corresponding to tubes of G ([CD06, Pos09]) while the graphical zonotope of G is the Minkowski sum of the faces of the standard simplex corresponding only to edges of G . Since the diameter of the graphical zonotope of G is the




FIGURE 7. Decompositions of trees with 3 or 4 leaves.

number e of edges of G , we obtain that the diameter $\delta(\mathcal{F}(G))$ is at least e . For the second part of the lower bound, we use again Theorem 6 to restrict the argument to trees. Let T be a tree on $n + 1$ vertices. We first discard some basic cases:

- (i) If T has precisely two leaves, then T is a path and the graph associahedron $\text{Asso}(T)$ is the classical n -dimensional associahedron, whose diameter is known to be larger than $2n - 4$ by L. Pournin's result [Pou14].
- (ii) If T has precisely 3 leaves, then it consists in 3 paths attached by a 3-valent node v , see Figure 7 (left). Let w be a neighbor of v and P_1, P_2 denote the connected components of $T \setminus w$. Observe that P_1 and P_2 are both paths and denote by $p_1 + 1$ and $p_2 + 1$ their respective lengths. Let T'_1, T''_1 (resp. T'_2, T''_2) be a diametral pair of maximal tubings on P_1 (resp. on P_2), and consider the maximal tubings $T' = T'_1 \cup T'_2 \cup \{P_1, P_2\}$ and $T'' = T''_1 \cup T''_2 \cup \{P_1, P_2\}$ on the tree T . Finally, denote by \bar{T} the maximal tubing on T obtained by flipping P_1 in T' . Since $\{P_1, P_2\}$ is a common upper set to the inclusion posets of T' and T'' , Proposition 16 (ii) ensures that there exists a geodesic from \bar{T} to T'' that starts by the flip from \bar{T} to T' . Moreover, Proposition 16 (i) ensures that the distance between T' and T'' is realized by a path staying in the face of $\text{Asso}(T)$ corresponding to $\{P_1, P_2\}$, which is the product of a classical p_1 -dimensional associahedron by a classical p_2 -dimensional associahedron. We conclude that

$$\delta(\mathcal{F}(T)) \geq 1 + \delta(\mathcal{F}(P_1)) + \delta(\mathcal{F}(P_2)) \geq 1 + (2p_1 - 4) + (2p_2 - 4) = 2(p_1 + p_2 + 2) - 11 = 2n - 11.$$

- (iii) If T has precisely 4 leaves, it either contains a single 4-valent node v or precisely two 3-valent nodes u, v , see Figure 7 (middle and right). Define w to be a neighbor of v , not located in the path between u and v in the latter situation. Then w disconnects T into a path P on $p + 1$ nodes and a tree Y with $y + 1$ nodes and precisely 3 leaves. A similar argument as in (ii) shows that

$$\delta(\mathcal{F}(T)) \geq 1 + \delta(\mathcal{F}(P)) + \delta(\mathcal{F}(Y)) \geq 1 + (2p - 4) + (2y - 11) = 2(p + y + 2) - 18 = 2n - 18.$$

We can now assume that the tree T has $k \geq 5$ leaves l_1, \dots, l_k . Let $\bar{V} = V \setminus \{l_1, \dots, l_k\}$ and $\bar{T} = T[\bar{V}]$ denote the tree obtained by deletion of the leaves of T . By induction hypothesis, there exists two maximal tubings \bar{T} and \bar{T}' on \bar{T} at distance at least $2(n - k) - 18$. Define $t_i := V \setminus \{l_1, \dots, l_i\}$ for $i \in [k]$, and $t'_j := V \setminus \{l_j, \dots, l_k\}$ for $j \in [k]$. Consider the maximal tubings $T := \bar{T} \cup \{t_1, \dots, t_k\}$ and $T' := \bar{T}' \cup \{t'_1, \dots, t'_k\}$ on T . We claim that the distance between these tubings is at least $2n - 18$. To see it, consider the surjection Ω from the tubings on T onto that of $\bar{T} \cup \{l_1, \dots, l_k\}$ as defined in Section 3.1. It sends a path $T = T_0, \dots, T_\ell = T'$ in the flip graph $\mathcal{F}(T)$ to a path

$$\bar{T} \cup \{\{l_1\}, \dots, \{l_k\}\} = \Omega(T_0), \dots, \Omega(T_\ell) = \bar{T}' \cup \{\{l_1\}, \dots, \{l_k\}\}$$

in the flip graph $\mathcal{F}(\bar{T} \cup \{l_1, \dots, l_k\})$ with repeated entries. Since \bar{T} and \bar{T}' are at distance at least $2(n - k) - 18$ in the flip graph $\mathcal{F}(\bar{T})$, this path has at least $2(n - k) - 18$ non-trivial steps, so we must show that it has at least $2k$ repetitions. These repetitions appear whenever we flip a tube t_i or t'_j . Indeed, we observe that the image $\Omega(t)$ of any tube $t \in \{t_i \mid i \in [k]\} \cup \{t'_j \mid j \in [k]\}$ is composed by \bar{V} together with single leaves of T . Since all these tubes are connected components of \bar{T} , we have $\Omega(T \setminus \{t\}) = \Omega(T)$ for any maximal loaded tubing T containing t . To conclude, we distinguish three cases:

- (i) If the tube $\mathbf{t}_k = \bar{V} = \mathbf{t}'_1$ is never flipped along the path $\mathbf{T} = \mathbf{T}_0, \dots, \mathbf{T}_\ell = \mathbf{T}'$, then we need at least $\binom{k}{2}$ flips to transform $\{\mathbf{t}_1, \dots, \mathbf{t}_k\}$ into $\{\mathbf{t}'_1, \dots, \mathbf{t}'_k\}$. This can be seen for example from the description of the link of \mathbf{t}_k in $\mathcal{N}(\mathbf{T})$ in [CD06, Theorem 2.9]. Finally, we use that $\binom{k}{2} \geq 2k$ since $k \geq 5$.
- (ii) Otherwise, we need to flip all $\{\mathbf{t}_1, \dots, \mathbf{t}_k\}$ and then back all $\{\mathbf{t}'_1, \dots, \mathbf{t}'_k\}$. If no flip of a tube \mathbf{t}_i produces a tube \mathbf{t}'_j , we need at least $2k$ flips which produces repetitions in $\Omega(\mathbf{T}_0), \dots, \Omega(\mathbf{T}_\ell)$.
- (iii) Finally, assume that we flip precisely once all $\{\mathbf{t}_1, \dots, \mathbf{t}_k\}$ and then back all $\{\mathbf{t}'_1, \dots, \mathbf{t}'_k\}$, and that a tube \mathbf{t}_i is flipped into a tube \mathbf{t}'_j . According to the description of flips, we must have $i = k - 1$ and $j = 2$. If $p \in [\ell]$ denotes the position such that $\mathbf{T}_p \setminus \{\mathbf{t}_{k-1}\} = \mathbf{T}_{p+1} \setminus \{\mathbf{t}'_2\}$, we moreover know that $\mathbf{t}_{k-1} \in \mathbf{T}_q$ for $q \leq p$, that $\mathbf{t}'_2 \in \mathbf{T}_q$ for $q > p$, and that $\bar{V} \in \mathbf{T}_p \cap \mathbf{T}_{p+1}$. Applying the non-leaving-face property either to the upper set $\{\mathbf{t}_{k-1}, \mathbf{t}_k\}$ in $\text{Asso}(\mathcal{G}[\mathbf{t}_{k-1}])$ or to the upper set $\{\mathbf{t}'_1, \mathbf{t}'_2\}$ in $\text{Asso}(\mathcal{G}[\mathbf{t}'_2])$, we conclude that it would shorten the path $\mathbf{T}_0, \dots, \mathbf{T}_\ell$ to avoid the flip of $\mathbf{t}_k = \bar{V} = \mathbf{t}'_1$, which brings us back to Situation (i). \square

Remark 19. We note that although asymptotically optimal, our lower bound $2n - 18$ is certainly not sharp. We expect the correct lower bound to be the bound $2n - 4$ for the associahedron. Better upper bound can also be worked out for certain families of graphs. For example, L. Pournin investigates the cyclohedra, *i.e.* cycle associahedra. As far as trees are concerned, we understand better stars and their subdivisions. The diameter $\delta(\mathcal{F}(K_{1,n}))$ for the star $K_{1,n}$ is exactly $2n$ (for $n \geq 5$), see Figure 6. In fact, the diameter of the graph associahedron of any starlike tree (subdivision of a star) on $n + 1$ vertices is bounded by $2n$. To see it, we observe that any tubing is at distance at most n from the tubing \mathbf{T}_\circ consisting in all tubes adjacent to the central vertex. Indeed, we can always flip a tube in a tubing distinct from \mathbf{T}_\circ to create a new tube adjacent to the central vertex. This argument is not valid for non-starlike trees.

Remark 20. The lower bound in Theorem 18 shows that the diameter $\delta(\mathcal{F}(G))$ is at least the number of edges of G . In view of Theorem 1, it is tempting to guess that the diameter $\delta(\mathcal{F}(G))$ is of the same order as the number of edges of G . Adapting arguments from Remark 19, we can show that the diameter of any tree associahedron $\delta(\mathcal{F}(T))$ is of order at most $n \log n$. In any case, the following question remains open.

Question 21. Is there a family of trees T_n on n nodes such that $\delta(\mathcal{F}(T_n))$ is of order $n \log n$? Even more specifically, consider the family of trees illustrated in Figure 8: $T_1 = K_{1,3}$ (tripod) and T_{k+1} is obtained by grafting two leaves to each leaf of T_k . What is the order of the diameter $\delta(\mathcal{F}(T_k))$?




FIGURE 8. The family of trees T_k : the tree T_1 is the tripod and T_{k+1} is obtained from T_k by connecting two new nodes to each leaf of T_k .

Remark 22. The upper bound $\delta(\mathcal{F}(B)) \leq \binom{n+1}{2}$ holds for an arbitrary building set B by Theorem 6 and the fact that the permutahedron is the nestohedron on the complete building set. In contrast, the lower bound is not valid for arbitrary connected building sets. For example, the nestohedron on the trivial connected building set $\{\{1\}, \dots, \{n+1\}, \{1, \dots, n+1\}\}$ is the n -dimensional simplex, whose diameter is 1.

4. HAMILTONICITY

In this section, we prove that the flip graph $\mathcal{F}(G)$ is Hamiltonian for any graph G with at least 2 edges. This extends the result of H. Steinhaus [Ste64], S. M. Johnson [Joh63], and H. F. Trotter [Tro62] for the permutahedron, and of J. Lucas [Luc87] for the associahedron (see also [HN99]). For all the proof, it is more convenient to work with spines than with tubings (remind Sections 2.2 and 2.3). We first sketch the strategy of our proof.

4.1. **Strategy.** For any vertex v of G , we denote by $\mathcal{F}_v(G)$ the graph of flips on all spines on G where v is a root. We call *fixed-root subgraphs* of $\mathcal{F}(G)$ the subgraphs $\mathcal{F}_v(G)$ for $v \in V$. Note that the fixed-root subgraph $\mathcal{F}_v(G)$ is isomorphic to the flip graph $\mathcal{F}(G[\hat{v}])$, where $G[\hat{v}]$ is the subgraph of G induced by $\hat{v} := V \setminus \{v\}$.

We now distinguish two extreme types of flips. Consider two maximal tubings T, T' on G and tubes $t \in T$ and $t' \in T'$ such that $\bar{T} := T \setminus \{t\} = T' \setminus \{t'\}$. Let S, S' and \bar{S} denote the corresponding spines and $\{v\} = \lambda(t, T)$ and $\{v'\} = \lambda(t', T')$. We say that the flip \bar{T} (or equivalently \bar{S}) is

- (i) a *short flip* if both t and t' are singletons, that is, if $\{v, v'\}$ is a leaf of \bar{S} ;
- (ii) a *long flip* if t and t' are maximal proper tubes in T and T' , that is, if $\{v, v'\}$ is a root of \bar{S} .

Note that in a short flip, the vertices v, v' are necessarily adjacent in G . In the short flip \bar{S} , we call *short leaf* the leaf labeled by $\{v, v'\}$ of \bar{S} , *short root* the root of the tree of \bar{S} containing the short leaf, and *short child* the child w of the short root on the path to the short leaf. If the short leaf is already a child of the short root, then it coincides with the short child. Moreover, the short root, short child and short leaf all coincide if they form an isolated edge of G . In the long flip \bar{S} , we call *long root* the root labeled by $\{v, v'\}$.

We define a *bridge* to be a square B in the flip graph $\mathcal{F}(G)$ formed by two short and two long flips. We say that these two short (resp. long) flips are *parallel*, and we borrow the terms long root and short leaf for the bridge B . Figure 9 illustrates the notions of bridge, long flips and short flips.

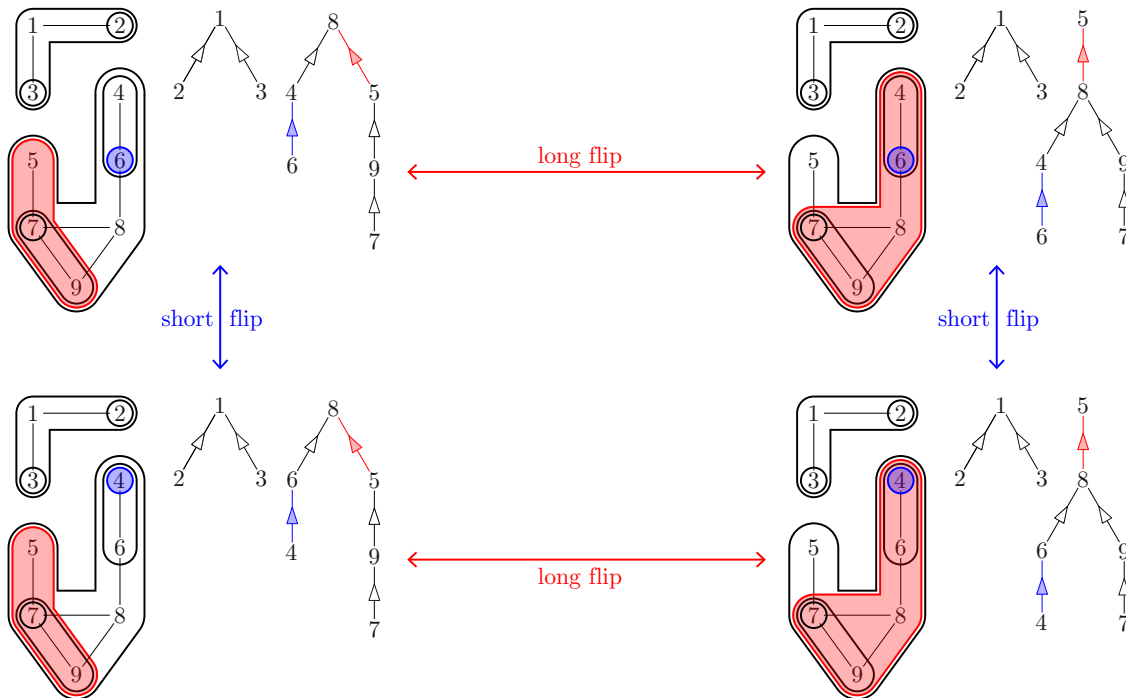



FIGURE 9. A bridge, with two long flips (red) and two short flips (blue).

In terms of spines, a bridge can equivalently be defined as a spine B of G where all labels are singletons, except the label $\{r, r'\}$ of a root and the label $\{s, s'\}$ of a leaf. We denote by $B[r]$ the short flip of B where r is a root, by $B[s]$ the long flip of B where s is a leaf, and by $B[rs]$ the maximal spine on G refining both $B[r]$ and $B[s]$, *i.e.* where r is a root and s a leaf. The flips $B[r']$ and $B[s']$ as well as the maximal spines $B[r's]$, $B[rs']$, and $B[r's']$ are defined similarly. These notations are summarized below



To obtain a Hamiltonian cycle \mathcal{H} of the flip graph $\mathcal{F}(G)$, we proceed as follows. The idea is to construct by induction a Hamiltonian cycle $\mathcal{H}_{\hat{v}}$ of each flip graph $\mathcal{F}(G[\hat{v}])$, which is isomorphic to a Hamiltonian cycle \mathcal{H}_v in each fixed-root subgraph $\mathcal{F}_v(G)$. We then select an ordering v_1, \dots, v_{n+1} of V , such that two consecutive Hamiltonian cycles \mathcal{H}_{v_i} and $\mathcal{H}_{v_{i+1}}$ meet the parallel short flips of a bridge B_i for all $i \in [n]$. The Hamiltonian cycle of $\mathcal{F}(G)$ is then obtained from the union of the cycles $\mathcal{H}_{v_1}, \dots, \mathcal{H}_{v_{n+1}}$ by exchanging the short flips with the long flips of all bridges B_1, \dots, B_n , as illustrated in Figure 10.

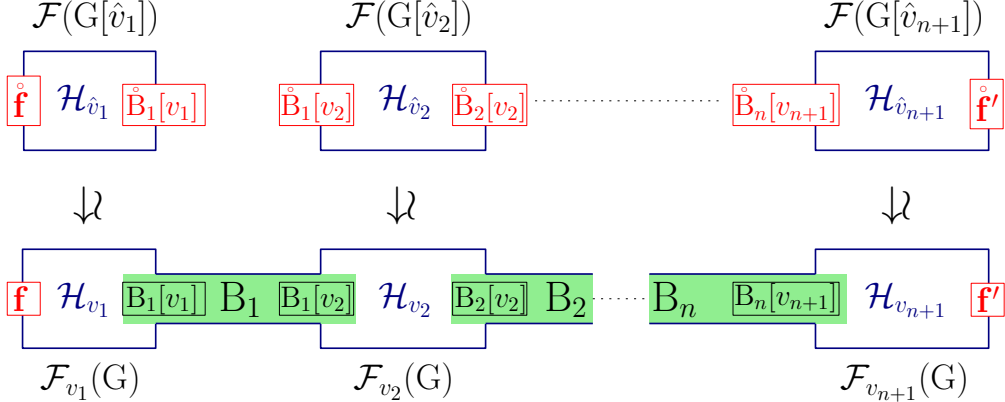


FIGURE 10. The strategy for the proof of the hamiltonicity of the flip graph $\mathcal{F}(G)$. The circles above the short flips in the flip graphs $\mathcal{F}(G[\hat{v}_i])$ on top indicate that they are obtained by deleting the root v_i in the corresponding short flip of the fixed-root subgraph $\mathcal{F}_{v_i}(G)$ on bottom. See also Theorem 23.

Of course, this description is a simplified and naive approach. The difficulty lies in that, given the Hamiltonian cycles \mathcal{H}_v of the fixed-root subgraphs $\mathcal{F}_v(G)$, the existence of a suitable ordering v_1, \dots, v_{n+1} of V and of the bridges B_1, \dots, B_n connecting the consecutive Hamiltonian cycles \mathcal{H}_{v_i} and $\mathcal{H}_{v_{i+1}}$ is not guaranteed. To overpass this issue, we need to impose the presence of two forced short flips in each Hamiltonian cycle \mathcal{H}_v . We include this condition in the induction hypothesis and prove the following sharper version of Theorem 4.

Theorem 23. *For any graph G , any pair of short flips of $\mathcal{F}(G)$ with distinct short roots is contained in a Hamiltonian cycle of the flip graph $\mathcal{F}(G)$.*

Note that for any graph G with at least 2 edges, the flip graph $\mathcal{F}(G)$ always contains two short flips with distinct short roots. Theorem 4 thus follows from the formulation of Theorem 23.

The issue in our inductive approach is that the fixed-root subgraphs of $\mathcal{F}(G)$ do not always contain two edges, and therefore cannot be treated by Theorem 23. Indeed, it can happen that:

- $G[\hat{v}]$ has a single edge and thus the fixed-root subgraph $\mathcal{F}_v(G) \sim \mathcal{F}(G[\hat{v}])$ is reduced to a single (short) flip. This case can still be treated with the same strategy: we consider this single flip $\mathcal{F}_v(G)$ as a degenerate Hamiltonian cycle and we can concatenate two bridges containing this short flip.
- $G[\hat{v}]$ has no edge and thus the fixed-root subgraph $\mathcal{F}_v(G) \sim \mathcal{F}(G[\hat{v}])$ is a point. This is the case when G is a star with central vertex v together with some isolated vertices. We need to make a special and independent treatment for this particular case. See Section 4.4.

4.2. Disconnected graphs. We first show how to restrict the proof to connected graphs using some basic results on products of cycles. We need the following lemmas.

Lemma 24. *For any two cycles $\mathcal{H}, \mathcal{H}'$ and any two edges e, e' of $\mathcal{H} \times \mathcal{H}'$, there exists a Hamiltonian cycle of $\mathcal{H} \times \mathcal{H}'$ containing both e and e' .*

Proof. The idea is illustrated in Figure 11. The precise proof is left to the reader. □




FIGURE 11. An idea for the proof of Lemma 24. Any pair of edges is contained in a Hamiltonian cycle similar to those. The pictures represent Cartesian products of the cycle \mathcal{H} with the path obtained by deleting one edge in \mathcal{H}' .

Lemma 25. *For any cycle \mathcal{H} , any isolated edge e_o and any two edges e, e' of $\mathcal{H} \times e_o$, there exists a Hamiltonian cycle containing both e and e' , as soon as one of the following conditions hold:*

- (1) *the edges e, e' are not both of the form $\{v\} \times e_o$ with $v \in \mathcal{H}$;*
- (2) *$e = \{v\} \times e_o$ and $e' = \{v'\} \times e_o$ where $\{v, v'\}$ is an edge of \mathcal{H} ;*
- (3) *\mathcal{H} has an even number of edges.*

Proof. The idea is illustrated in Figure 12. The precise proof is left to the reader. □




FIGURE 12. An idea for the proof of Lemma 25. The right picture only works for even cycles.

Corollary 26. *If two graphs G, G' both have the property that any pair of short flips of their flip graph with distinct short roots is contained in a Hamiltonian cycle of their flip graph, then $G \sqcup G'$ fulfills the same property.*

Proof. We have seen that the flip graph of the disjoint union of two graphs G_1 and G_2 is the product of their flip graphs $\mathcal{F}(G_1)$ and $\mathcal{F}(G_2)$. The statement thus follows from the previous lemmas. □

4.3. Generic proof. We now present an inductive proof of Theorem 23. Corollary 26 allows us to restrict to the case where G is connected. For technical reasons, the stars and the graphs with at most 6 vertices will be treated separately. We thus assume here that G is not a star and has at least 7 vertices, which ensures that any fixed root subgraph of the flip graph $\mathcal{F}(G)$ has at least one short flip. Fix two short flips \mathbf{f}, \mathbf{f}' of $\mathcal{F}(G)$ with distinct short roots v_1, v_{n+1} respectively.

We follow the strategy described in Section 4.1 and illustrated in Figure 10. To apply Theorem 23 by induction on $G[\hat{v}_i]$, the short flips $\mathbf{B}_{i-1}[v_i]$ and $\mathbf{B}_i[v_i]$ should have distinct short children. This forbids certain positions for v_{i+1} in $\mathbf{B}_{i-1}[v_i]$ illustrated in Figure 13, and motivates the following definition. We say that a vertex w and a short flip \mathbf{g} with root v are *in conflict* if either of the following happens:

- (A) $\{w\}$ is the short child of \mathbf{g} and all other children of v in \mathbf{g} are isolated in $G[\hat{v}]$;
- (B) the graph $G[\hat{v}]$ has at least three edges, the graph $G \setminus \{v, w\}$ has exactly one edge which is the short leaf of \mathbf{g} ;
- (C) the graph $G[\hat{v}]$ has exactly two edges, the graph $G \setminus \{v, w\}$ has exactly one edge which is the short leaf of \mathbf{g} , and w is a child of v .

It is immediate that a short flip is in conflict with at most one vertex. Observe also that if w is in the short leaf of \mathbf{g} , then w and \mathbf{g} cannot be in conflict.




FIGURE 13. Short flips in conflict with vertex w . The short leaf is shaded. The second short flip of Case (B) is in conflict with w only if the connected component of $G[\hat{v}]$ containing w is a star with central vertex w .

We now show how we order the vertices v_1, \dots, v_{n+1} such that for each $i \in [n]$ there exists a bridge \mathbf{B}_i connecting the fixed-root subgraphs $\mathcal{F}_{v_i}(G)$ and $\mathcal{F}_{v_{i+1}}(G)$.

Lemma 27. *There exists an ordering v_1, \dots, v_{n+1} of the vertices of G (provided $|V| \geq 7$) satisfying the following properties:*

- v_2 and \mathbf{f} are not in conflict, and v_n and \mathbf{f}' are not in conflict, and
- for any $i \in [n]$, the graph G contains an edge disjoint from $\{v_i, v_{i+1}\}$.

Proof. Let a and a' denote the vertices in conflict with \mathbf{f} and \mathbf{f}' if any. Let D denote the set of *totally disconnecting pairs* of G , i.e. of pairs $\{x, y\}$ such that $G \setminus \{x, y\}$ has no edge. We want to show that there exists an ordering on the vertices of G in which neither $\{v_1, a\}$ nor $\{a', v_{n+1}\}$, nor any pair of D are consecutive. For this, we prove that if G has at least 5 vertices and is not a star (i.e. all edges contain a central vertex), then $|D| \leq 2$ and the pairs in D are not disjoint.

Suppose by contradiction that D contains two disjoint pairs $\{x_1, y_1\}$ and $\{x_2, y_2\}$. Then any edge of G intersects both pairs, so that x_1, x_2, y_1, y_2 are the only vertices in G (by connectivity), contradicting that G has at least 5 vertices. Suppose now that D contains three pairwise distinct pairs $\{x, y_1\}$, $\{x, y_2\}$ and $\{x, y_3\}$. Then any edge of G contains x since it cannot contain y_1, y_2 and y_3 together. It follows that G is a star with central vertex x .

Since $|D| \leq 2$, at most 4 pairs of vertices of G cannot be consecutive in our ordering. It is thus clear that if there are enough other vertices, we can find a suitable ordering. In fact, it turns out that it is already possible as soon as G has 7 vertices. It is easy to prove by a boring case analysis. We just treat the worst case below.

Assume that $D = \{\{x, y\}, \{x, z\}\}$ where $x, y, z \notin \{v_1, v_{n+1}\}$ and that x is in conflict with both short flips \mathbf{f} and \mathbf{f}' . Since $|V| \geq 7$, there exists two distinct vertices $u, v \notin \{v_1, v_{n+1}, x, y, z\}$ and we set $v_2 = z, v_3 = y, v_4 = u, v_5 = x, v_6 = v$ and choose any ordering for the remaining vertices. This order satisfies the requested conditions. \square

Remark 28. In fact, using similar arguments, one can easily check that the result of Lemma 27 holds in the following situations:

- $|V| = 6$, and either $|D| \leq 1$ or $D = \{\{x, y\}, \{x, z\}\}$ where x is not in conflict with both \mathbf{f}, \mathbf{f}' .
- $|V| = 5$, and either $D = \emptyset$ or $D = \{\{x, y\}\}$ where neither x nor y is in conflict with both \mathbf{f}, \mathbf{f}' .
- $|V| = 5$, and $D = \{\{x, y\}, \{x, z\}\}$ and $|\{x, y, z\} \cap \{v_1, v_5\}| = 2$.
- $|V| = 5$, and $D = \{\{x, y\}, \{x, z\}\}$ and $|\{x, y, z\} \cap \{v_1, v_5\}| = 1$ and neither of x, y, z is in conflict with any of \mathbf{f} and \mathbf{f}' .

Given such an ordering v_1, \dots, v_{n+1} , we choose bridges B_1, \dots, B_n connecting the fixed-root subgraphs $\mathcal{F}_{v_1}(G), \dots, \mathcal{F}_{v_{n+1}}(G)$. We start with the choice of B_1 .

Lemma 29. *There exists a bridge B_1 with root $\{v_1, v_2\}$ such that*

- if $\mathcal{F}_{v_1}(G)$ is a square, the short flips \mathbf{f} and $B_1[v_1]$ are distinct,
- if $\mathcal{F}_{v_1}(G)$ is not reduced to a single flip nor to a square, the short flips \mathbf{f} and $B_1[v_1]$ have distinct short children,
- $B_1[v_2]$ and v_3 are not in conflict, and
- the singleton $\{v_3\}$ is a child of v_2 in $B_1[v_2]$ only if v_3 is isolated in $G \setminus \{v_1, v_2\}$.

Proof. The proof is an intricate case analysis. In each case, we will provide a suitable choice for B_1 , but the verification that this bridge exists and satisfies the conditions of the statement is immediate and left to the reader. We denote by κ the connected component of $G[\hat{v}_1]$ containing v_2 . The following cases cover all possibilities:

- ♠ $\kappa = \{v_2\}$:
 - ♡ $G \setminus \{v_1, v_2\}$ has only one edge: the fixed root subgraph $\mathcal{F}_{v_1}(G)$ is reduced to the short flip \mathbf{f} and the bridge obtained by contracting $\{v_1, v_2\}$ in \mathbf{f} suits for B_1 .
 - ♡ $G \setminus \{v_1, v_2\}$ has at least two edges: we choose for B_1 any bridge with root $\{v_1, v_2\}$ and with a short child different from that of \mathbf{f} .
- ♠ $\kappa \neq \{v_2\}$, so that κ has at least one edge:
 - ♡ $G[\hat{v}_1] \setminus \kappa$ has no edge: Condition (A) on \mathbf{f} and v_2 ensures that v_2 is not the short child of \mathbf{f} . Since the short leaf of \mathbf{f} has to be in κ , the short children of \mathbf{f} and $B_1[v_1]$ will automatically be different.
 - ◇ $v_3 \notin \kappa$: any bridge with root $\{v_1, v_2\}$ suits for B_1 .
 - ◇ $v_3 \in \kappa$:
 - ♣ v_3 is isolated in $\kappa \setminus \{v_2\}$: any bridge with root $\{v_1, v_2\}$ suits for B_1 .
 - ♣ v_3 is not isolated in $\kappa \setminus \{v_2\}$: we choose for B_1 a bridge with root $\{v_1, v_2\}$ and whose short leaf contains v_3 .
 - ♡ $G[\hat{v}_1] \setminus \kappa$ has precisely one edge e :
 - ◇ e is not the short leaf of \mathbf{f} : we choose for B_1 any bridge with root $\{v_1, v_2\}$, short leaf e and in which $\{v_3\}$ is a child of the root only if it is isolated in $G \setminus \{v_1, v_2\}$.
 - ◇ e is the short leaf of \mathbf{f} :
 - ♣ κ is a single edge: we choose for B_1 the bridge obtained by contracting $\{v_1, v_2\}$ in the short flip opposite to \mathbf{f} in the square $\mathcal{F}_{v_1}(G)$ (which suits by Condition (C)).
 - ♣ κ has at least two edges: Condition (B) ensures that $\kappa \setminus \{v_2\}$ has at least one edge.
 - $v_3 \notin \kappa$: any bridge with root $\{v_1, v_2\}$ and short leaf in κ suits for B_1 .
 - $v_3 \in \kappa$:
 - ★ v_3 is isolated in $\kappa \setminus \{v_2\}$: any bridge with root $\{v_1, v_2\}$ and short leaf in κ suits.
 - ★ v_3 is not isolated in $\kappa \setminus \{v_2\}$: we choose for B_1 a bridge with root $\{v_1, v_2\}$ and whose short leaf contains v_3 .
 - ♡ $G[\hat{v}_1] \setminus \kappa$ has at least two edges:

- ◇ $G[\hat{v}_1] \setminus \kappa$ has only one non-trivial connected component: we choose for B_1 a bridge with root $\{v_1, v_2\}$, with short leaf containing the non-isolated child of v_1 in \mathbf{f} which is not in κ , and in which $\{v_3\}$ is a child of the root only if it is either isolated in $G \setminus \{v_1, v_2\}$ or the short child of $B_1[v_1]$.
- ◇ $G[\hat{v}_1] \setminus \kappa$ has at least two non-trivial connected components: we choose for B_1 a bridge with root $\{v_1, v_2\}$, with short leaf in a connected component of $G[\hat{v}_1] \setminus \kappa$ not containing the short leaf of \mathbf{f} , and in which $\{v_3\}$ is a child of the root only if it is either isolated in $G \setminus \{v_1, v_2\}$ or the short child of $B_1[v_1]$. \square

The choice of B_n is similar to that of B_1 , replacing v_1, v_2, v_3 and \mathbf{f} by v_{n+1}, v_n, v_{n-1} and \mathbf{f}' respectively. For choosing the other bridges B_2, \dots, B_{n-1} , we first observe the existence of certain special vertices in G .

We say that a vertex distinct from v_1 and v_{n+1} which disconnects at most one vertex is an *almost leaf* of G . Observe that G contains at least one almost leaf: Consider a spanning tree T of G . If T is a path from v_1 to v_{n+1} , the neighbor of v_1 in T is an almost leaf of G . Otherwise, any leaf of T distinct from v_1 and v_{n+1} is an almost leaf of G .

Choose an almost leaf v_i of G which disconnects no vertex if possible, and any almost leaf otherwise. We sequentially construct the bridges B_2, \dots, B_{i-1} : once B_j is constructed, we choose B_{j+1} using Lemma 29 where we replace v_1, v_2, v_3 and \mathbf{f} by $v_{j+1}, v_{j+2}, v_{j+3}$ and $B_j[v_{j+1}]$. Similarly, we choose the bridges B_{n-1}, \dots, B_{i+1} : once B_{j+1} is constructed, we choose B_j using Lemma 29 where we replace v_1, v_2, v_3 and \mathbf{f} by v_{j+1}, v_j, v_{j-1} and $B_{j+1}[v_{j+1}]$. Note that the conditions on B_1 required in Lemma 29 ensure that the hypotheses in Lemma 27 can be propagated.

It remains to properly choose the last bridge B_i . This is done by the following statement.

Lemma 30. *Let \mathbf{g}, \mathbf{h} be two short flips on G with distinct roots v, w respectively. Assume that*

- (i) $G \setminus \{v, w\}$ has at least one edge;
- (ii) \mathbf{g} and w are not in conflict, and \mathbf{h} and v are not in conflict;
- (iii) $\{v\}$ is a child of w in \mathbf{h} only if v is isolated in $G[\hat{w}]$;
- (iv) v disconnects at most one vertex of G and this vertex is not w .

Then there exists a bridge B with root $\{v, w\}$ such that \mathbf{g} and $B[v]$ are distinct if $\mathcal{F}_v[G]$ is not reduced to a single flip and have distinct short children if $\mathcal{F}_v[G]$ is not a square, and similarly for \mathbf{h} and $B[w]$.

Proof. Condition (iv) implies that $\{w\}$ is the short child of $B[v]$ for any bridge B with root $\{v, w\}$. In contrast, Condition (iv) and Condition (A) for \mathbf{g} and w ensure that $\{w\}$ is not the short child of \mathbf{g} . Therefore, the conclusion of the lemma holds for \mathbf{g} and $B[v]$, for any bridge B with root $\{v, w\}$. The difficulty is to choose B in order to satisfy the conclusion for \mathbf{h} and $B[w]$. For this, we distinguish various cases, in a similar manner as in Lemma 29. Again, we provide in each case a suitable choice for B , but the verification that this bridge exists and satisfies the conditions of the statement is immediate and left to the reader.

- ♠ $G \setminus \{v, w\}$ has exactly one edge e : this edge e has to be the short leaf of any bridge with root $\{v, w\}$, thus Condition (A) for \mathbf{h} and v ensures that e is isolated in $G[\hat{w}]$.
- ♡ e is the short leaf of \mathbf{h} : Condition (B) for \mathbf{h} and v ensures that $\mathcal{F}_w(G)$ is either a single flip or a square (because v disconnects at most one vertex from G).
 - ◇ $\mathcal{F}_w(G)$ is a single flip: B is obtained by contracting $\{v, w\}$ in \mathbf{h} .
 - ◇ $\mathcal{F}_w(G)$ is a square: Condition (C) for \mathbf{h} and v ensures that v is not a child of w in \mathbf{h} and B is obtained by contracting $\{v, w\}$ in the short flip opposite to \mathbf{h} in the square $\mathcal{F}_w(G)$.
- ♡ e is not the short leaf of \mathbf{h} : we choose for B any bridge with root $\{v, w\}$ and short leaf e .
- ♠ $G \setminus \{v, w\}$ has at least two edges:
 - ♡ the short leaf and the short child of \mathbf{h} coincide: any bridge with root $\{v, w\}$ and a short leaf distinct from that of \mathbf{h} suits for B .
 - ♡ the short leaf and the short child of \mathbf{h} are distinct: we choose for B a bridge with root $\{v, w\}$ whose short leaf contains the short child of \mathbf{h} . \square

We have now chosen the order on the vertices v_1, \dots, v_{n+1} and chosen for each $i \in [n]$ a bridge B_i connecting the fixed-root subgraphs $\mathcal{F}_{v_i}(G)$ and $\mathcal{F}_{v_{i+1}}(G)$. Our choice forces the short flips $B_{i-1}[v_i]$ and $B_i[v_i]$ (as well as the short flips \mathbf{f} and $B_1[v_1]$ and the short flips $B_n[v_{n+1}]$ and \mathbf{f}') to be distinct if $\mathcal{F}_{v_i}[G]$ is not reduced to a single flip and have distinct short children if $\mathcal{F}_{v_i}[G]$ is not a square. We then construct a Hamiltonian cycle \mathcal{H}_i in each fixed-root subgraph $\mathcal{F}_{v_i}(G)$ such that \mathcal{H}_1 contains the short flips \mathbf{f} and $B_1[v_1]$, \mathcal{H}_{n+1} contains the short flips $B_n[v_{n+1}]$ and \mathbf{f}' , and \mathcal{H}_i contains the short flips $B_{i-1}[v_i]$ and $B_i[v_i]$ for all $2 \leq i \leq n$. Note that

- when $\mathcal{F}_{v_i}(G)$ is reduced to a single flip, we just set $\mathcal{H}_i = \mathcal{F}_{v_i}(G)$ and consider it as a degenerate Hamiltonian cycle;
- when $\mathcal{F}_{v_i}(G)$ is a square, it is already a cycle;
- otherwise, we apply Theorem 23 by induction to $G[\hat{v}_i]$ and obtain the Hamiltonian cycle \mathcal{H}_i . The theorem applies since the short flips $B_{i-1}[v_i]$ and $B_i[v_i]$ have distinct short children, so that the corresponding short flips in $\mathcal{F}(G[\hat{v}_i])$ have distinct short roots.

Finally, we obtain a Hamiltonian cycle of $\mathcal{F}(G)$ containing \mathbf{f} and \mathbf{f}' by gluing the cycles $\mathcal{H}_1, \dots, \mathcal{H}_{n+1}$ together using the bridges B_1, \dots, B_n as explained in Section 4.1. This is possible since the short flips $B_{i-1}[v_i]$ and $B_i[v_i]$ both belong to the Hamiltonian cycle \mathcal{H}_i , and are distinct when $\mathcal{F}_{v_i}(G)$ is not reduced to a single flip. This concludes the proof for all generic cases. The remaining of the paper deals with the special cases of stars and graphs with at most 6 vertices.

4.4. Stars. We now treat the particular case of stars. Consider a ground set V where a vertex $*$ is distinguished. The *star* on V is the tree X_V where all vertices of $V \setminus \{*\}$ are leaves connected to $*$. The flip graph $\mathcal{F}(X_V)$ has two kinds of fixed-root subgraphs:

- $\mathcal{F}_*(X_V)$ is reduced to a single spine \otimes with root $*$ and n leaves;
- for any other vertex $v \in V \setminus \{*\}$, the fixed-root subgraph $\mathcal{F}_v(X_V)$ is isomorphic to the flip graph $\mathcal{F}(X_{\hat{v}})$ of the star $X_{\hat{v}}$, where $*$ is still the distinguished vertex in $\hat{v} = V \setminus \{v\}$. For a spine $S \in \mathcal{F}_v(X_V)$, we denote by \hat{S} the unique subspine of S , and we write in column $S = \overset{v}{\hat{S}}$.

To find a Hamiltonian cycle passing through forced short flips and through the spine \otimes we need to refine again the induction hypothesis of Theorem 23 as follows.

Proposition 31. *Assume that $|V| \geq 3$, and fix two short flips \mathbf{f}, \mathbf{f}' of $\mathcal{F}(X_V)$ with distinct roots $r \neq r'$ and a long flip \mathbf{g} of $\mathcal{F}(X_V)$ with root $\{r'', *\}$. Then the flip graph $\mathcal{F}(X_V)$ has a Hamiltonian cycle containing $\mathbf{f}, \mathbf{f}', \mathbf{g}$.*

Proof. The proof works by induction on $|V|$. If $|V| = 3$, then X_V is a 3-path and its flip graph is a pentagon. The case $|V| = 4$ is solved by Figure 14 up to relabeling of V . Namely, whatever triple $\mathbf{f}, \mathbf{f}', \mathbf{g}$ is imposed, there is a permutation of the leaves of $X_{\{1,2,3,*\}}$ which sends the Hamiltonian cycle of Figure 14 to a Hamiltonian cycle passing through $\mathbf{f}, \mathbf{f}', \mathbf{g}$. Assume now that $|V| \geq 5$. We distinguish two cases.




FIGURE 14. A Hamiltonian cycle in the flip graph $\mathcal{F}(X_{\{1,2,3,*\}})$. Up to permutations of the leaves $\{1, 2, 3\}$, this cycle contains all possible triples $\mathbf{f}, \mathbf{f}', \mathbf{g}$ considered in Proposition 31.

Case 1: $r'' \in \{r, r'\}$, say for instance $r'' = r$. Let w' denote the child of r' in the short flip \mathbf{f}' . Let v_1, \dots, v_{n-2} be an arbitrary ordering of $V \setminus \{*, r, r'\}$ such that $v_1 \neq w'$ (this is possible since $|V| \geq 5$), and B_1, \dots, B_{n-2} any bridges such that the root of B_i is $\{v_{i-1}, v_i\}$ (where we set $v_0 = r'$). We now choose inductively a Hamiltonian cycle $\mathcal{H}_{\hat{v}}$ in each flip graph $\mathcal{F}(X_{\hat{v}})$ for all $v \in V \setminus \{*\}$ as follows.

- (i) In $\mathcal{F}(X_{\hat{r}})$, we choose a cycle $\mathcal{H}_{\hat{r}}$ containing the short flip $\mathring{\mathbf{f}}$ and the long flip $\circledast \leftrightarrow r'_{\circledast}$.
- (ii) In $\mathcal{F}(X_{\hat{r}'})$, we choose a cycle $\mathcal{H}_{\hat{r}'}$ containing the short flips $\mathring{\mathbf{f}}'$ and $\mathring{B}_1[r']$ and the long flip $\circledast \leftrightarrow r'_{\circledast}$.
- (iii) In $\mathcal{F}(X_{\hat{v}_i})$ for $i \in [n-3]$, we choose a cycle $\mathcal{H}_{\hat{v}_i}$ containing the short flips $\mathring{B}_i[v_i]$ and $\mathring{B}_{i+1}[v_i]$.
- (iv) In $\mathcal{F}(X_{\hat{v}_{n-2}})$, we choose a cycle $\mathcal{H}_{\hat{v}_{n-2}}$ containing the short flip $\mathring{B}_{n-2}[v_{n-2}]$.

Note that these Hamiltonian cycles exist by induction hypothesis. Indeed, the short flips $\mathring{B}_i[v_i]$ and $\mathring{B}_{i+1}[v_i]$ have distinct roots v_{i-1} and v_{i+1} . The only delicate case is thus Point (ii): the short flips $\mathring{\mathbf{f}}'$ and $\mathring{B}_1[r']$ have distinct roots since we forced v_1 to be different from w' . Each Hamiltonian cycle $\mathcal{H}_{\hat{v}}$ on $\mathcal{F}(X_{\hat{v}})$ induces a Hamiltonian cycle \mathcal{H}_v on $\mathcal{F}_v(X_V)$ (just add v at the root in all spines). From these Hamiltonian cycles, we construct a Hamiltonian cycle for $\mathcal{F}(X_V)$ as illustrated in Figure 15. We join \mathcal{H}_r with $\mathcal{H}_{r'}$ by deleting the flips $\overset{r}{\circledast} \leftrightarrow \overset{r'}{\circledast}$ and $\overset{r'}{\circledast} \leftrightarrow \overset{r}{\circledast}$ while inserting the long flips $\overset{r}{\circledast} \leftrightarrow \circledast \leftrightarrow \overset{r'}{\circledast}$ and $\overset{r'}{\circledast} \leftrightarrow \overset{r}{\circledast}$. Finally, we use the bridges B_1, \dots, B_{n-2} to connect the resulting cycle to the cycles $\mathcal{H}_{v_1}, \dots, \mathcal{H}_{v_{n-2}}$ by exchanging their short flips with their long flips.

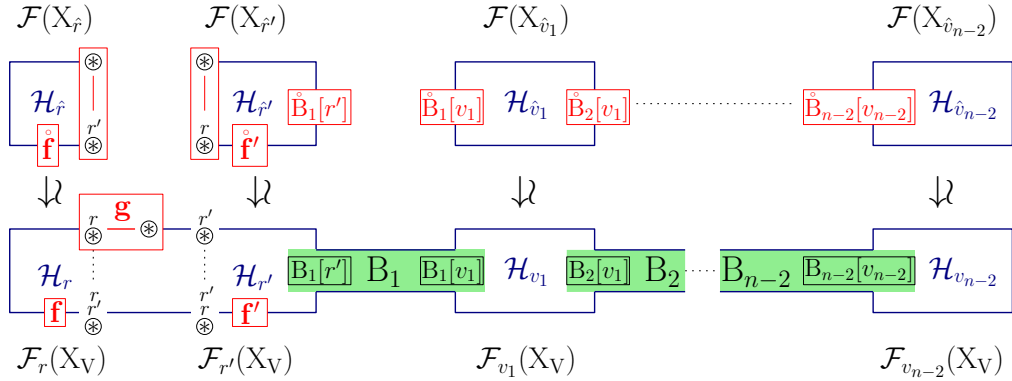


FIGURE 15. Construction of a Hamiltonian cycle in $\mathcal{F}(X_V)$ when $r'' = r$.

Case 2: $r'' \notin \{r, r'\}$. Let v_1, \dots, v_{n-3} be an arbitrary ordering of $V \setminus \{*, r, r', r''\}$, and B_1, \dots, B_{n-3} any bridges such that the root of B_i is $\{v_{i-1}, v_i\}$ (where we set $v_0 = r''$). We now choose inductively a Hamiltonian cycle $\mathcal{H}_{\hat{v}}$ in each flip graph $\mathcal{F}(X_{\hat{v}})$ for all $v \in V \setminus \{*, r'\}$ as follows.

- (i) In $\mathcal{F}(X_{\hat{r}})$, we choose a cycle $\mathcal{H}_{\hat{r}}$ containing the short flip $\mathring{\mathbf{f}}$ and the long flip $\circledast \leftrightarrow r'_{\circledast}$.
- (ii) In $\mathcal{F}(X_{\hat{r}''})$, we choose a cycle $\mathcal{H}_{\hat{r}''}$ containing a short flip $\mathring{\mathbf{h}}$ with root r' , the short flip $\mathring{B}_1[r'']$ and the long flip $\circledast \leftrightarrow r'_{\circledast}$.
- (iii) In $\mathcal{F}(X_{\hat{v}_i})$ for $i \in [n-4]$, we choose a cycle $\mathcal{H}_{\hat{v}_i}$ containing the short flips $\mathring{B}_i[v_i]$ and $\mathring{B}_{i+1}[v_i]$.
- (iv) In $\mathcal{F}(X_{\hat{v}_{n-3}})$, we choose a cycle $\mathcal{H}_{\hat{v}_{n-3}}$ containing the short flip $\mathring{B}_{n-3}[v_{n-3}]$ and a short flip $\mathring{\mathbf{k}}$ with root r' .

Each Hamiltonian cycle $\mathcal{H}_{\hat{v}}$ on $\mathcal{F}(X_{\hat{v}})$ induces a Hamiltonian cycle \mathcal{H}_v on $\mathcal{F}_v(X_V)$ (just add v at the root in all spines). From these Hamiltonian cycles, we construct the cycle illustrated in Figure 16. We still have to enlarge this cycle to cover $\mathcal{F}_{r'}(X_V)$. Let \mathbf{h}' and \mathbf{k}' denote the short flips in $\mathcal{F}_{r'}(X_V)$ parallel to the short flips \mathbf{h} and \mathbf{k} respectively. Since $r'' \neq v_{n-3}$, the root w' of $\mathring{\mathbf{f}}'$ cannot coincide with both. Assume for example that $w' \neq r''$. By induction, we can then find a Hamiltonian cycle $\mathcal{H}_{\hat{r}'}$ of $\mathcal{F}(X_{\hat{r}'})$ containing both $\mathring{\mathbf{f}}'$ and $\mathring{\mathbf{h}}'$. This cycle induces a Hamiltonian cycle $\mathcal{H}_{r'}$ of $\mathcal{F}_{r'}(X_V)$ passing through \mathbf{f}' and \mathbf{h}' . We can then connect this cycle to the cycle of

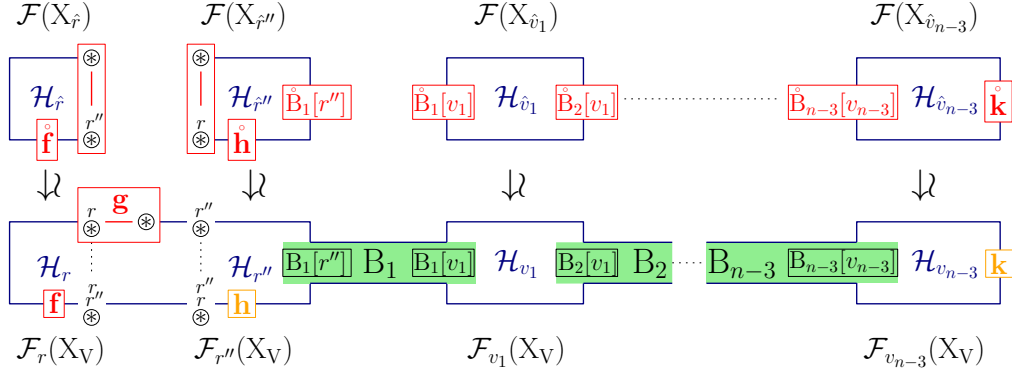


FIGURE 16. Construction of a Hamiltonian cycle in $\mathcal{F}(X_V)$ when $r'' \notin \{r, r'\}$.

Figure 16 by exchanging the parallel short flips \mathbf{h} and \mathbf{h}' by the corresponding parallel long flips. In the situation when $w' = r''$, we have $w' \neq v_{n-3}$ and we argue similarly by attaching $\mathcal{F}_{r'}(X_V)$ to \mathbf{k} instead of \mathbf{h} . \square

4.5. **Graph with at most 6 vertices.** Again we will focus on connected graphs because of Corollary 26. The analysis for graphs with at most 3 vertices is immediate. We now treat separately the graphs with 4, 5 and 6 vertices, which are not stars (stars have been treated in the previous section).

4.5.1. *Graphs with 4 vertices.* We consider all possible connected graphs on 4 vertices and exhibit explicit Hamiltonian cycles of their flip graphs. To do so, we could draw a cycle of spines as in Figure 14 (middle). Instead, we rather draw the Hamiltonian cycle on the flip graph $\mathcal{F}(G)$ represented as the 1-skeleton of the graph associahedron $\text{Asso}(G)$ as in Figure 14 (right). Let us remind from [CD06] that the graph associahedron $\text{Asso}(G)$ is obtained from the standard simplex $\Delta_V := \text{conv}\{e_v \mid v \in V\}$ (where $(e_v)_{v \in V}$ denotes the canonical basis of \mathbb{R}^V) by successive truncations of the faces $\Delta_{V \setminus \mathbf{t}} = \text{conv}\{e_v \mid v \in V \setminus \mathbf{t}\}$ for the tubes \mathbf{t} of G , in decreasing order of dimension. Each tube \mathbf{t} of G corresponds to a facet $F_{\mathbf{t}}$ of $\text{Asso}(G)$, and each maximal tubing \mathbf{T} corresponds to the vertex of $\text{Asso}(G)$ which belongs to all facets $F_{\mathbf{t}}$ for $\mathbf{t} \in \mathbf{T}$. In Figure 17 (right), we label the positions of the vertices of Δ_V before the truncations. The fixed-root subgraphs appear as the 1-skeleta of the four shaded faces of G , and the bridges are the five thin parallelograms (the short flips correspond to their short sides, and the long flips correspond to their long sides).




FIGURE 17. Correspondence between vertices of $\text{Asso}(G)$ and spines on G .

Using these conventions, Figure 18 represents Hamiltonian cycles for the flip graphs on all connected graphs on 4 vertices (the 4-star was already treated in Figure 14). The Hamiltonian cycles, together with their orbits under the action of the isomorphism group of the corresponding graph, prove the following statements, which imply Theorem 23 for all graphs on 4 vertices.





FIGURE 18. Hamiltonian cycles showing Proposition 32. Each vertex of the graph associahedra corresponds to a spine as explained in Figure 17.

- Proposition 32.** (a) For any graph G on at most 4 vertices, any pair of short flips (even with the same root) is contained in a Hamiltonian cycle of $\mathcal{F}(G)$.
- (b) For the stars on 3 and 4 vertices, each triple consisting of two short flips (even with the same root) and one long flip as in Proposition 31 is contained in a Hamiltonian cycle of $\mathcal{F}(G)$.
- (c) For the classical 3-dimensional (path) associahedron, there exists a Hamiltonian cycle containing simultaneously all short flips.
- (d) For all connected graphs on 4 vertices, there exist a Hamiltonian cycle of $\mathcal{F}(G)$ containing at least one short flip in each fixed-root subgraph. We can even preserve this property if we impose the Hamiltonian cycle to pass through one distinguished short flip.

4.5.2. *Graphs with 5 vertices.* Graphs on 5 vertices are treated by a case analysis. As in the proof of Lemma 27, we will denote by D the set of totally disconnecting pairs of G , i.e. pairs $\{x, y\}$ of vertices of G such that $G \setminus \{x, y\}$ has no edge. Recall from the proof of Lemma 27 that D has at most two elements and that they are not disjoint.


Consider now a graph G on 5 vertices. According to Remark 28, the proof of Section 4.3 applies in various configurations. We treat here the remaining cases. As we observed in Proposition 32 (a) that for any connected graph G on at most 4 vertices, any pair of short flips (even with the same root) is contained in a Hamiltonian cycle of $\mathcal{F}(G)$, we can ignore Condition (A) in the definition of conflict. We therefore say that a vertex w and a short flip \mathbf{g} with root v are *in conflict* if $G \setminus \{v, w\}$ has a single edge which is the short leaf of \mathbf{g} , and w is a child of v . With this definition, there is only one bridge connecting $\mathcal{F}_v(G)$ and $\mathcal{F}_w(G)$, but we cannot use it if we want the short flip \mathbf{g} to belong to the Hamiltonian cycle. One can check that the conclusions of Lemmas 29 and 30 still hold in this situation.

We first suppose that $D = \{\{x, y\}\}$ is a singleton and that either x or y is in conflict with both \mathbf{f} and \mathbf{f}' . Checking all connected graphs on five vertices, we see that this situation can only happen for the following graphs:




For each one, we explain how to prove Theorem 23.

$G = G_1$: The only possible conflicts are between x and a short flip with root v or w . Thus, up to isomorphism of the graph, the only instance of Theorem 23 fitting to the configuration we are looking at is given by



Observe that there exists bridges B_v, B_u, B_y with respective roots $\{w, v\}, \{w, u\}, \{w, y\}$ and a bridge B with root $\{u, x\}$. Notice that the fixed root subgraph $\mathcal{F}_w(G_1)$ is isomorphic to the classical (path) associahedron so that Proposition 32(c) ensures that there exists a Hamiltonian cycle \mathcal{H}_w of the flip graph $\mathcal{F}_w(G_1)$ containing all the short flips $\mathbf{f}, B_v[w], B_u[w], B_y[w]$. Moreover Proposition 32(a) ensures that there exists a Hamiltonian cycle \mathcal{H}_y (resp. \mathcal{H}_x) of the flip graph $\mathcal{F}_y(G_1)$ (resp. $\mathcal{F}_x(G_1)$) containing the short flip $B_y[y]$ (resp. $B[x]$). Proposition 32(a) again gives us a Hamiltonian cycle \mathcal{H}_u (resp. \mathcal{H}_v) of the flip graph $\mathcal{F}_u(G_1)$ (resp. $\mathcal{F}_v(G_1)$) containing the two short flips $B_u[u]$ and $B[u]$ (resp. \mathbf{f}' and $B_v[v]$). Note that the short flips of the bridges are all distinct since u, v and w do not disconnect the graph. Gluing all the Hamiltonian cycles of the fixed root subgraphs along the bridges as explained in Section 4.1 gives a Hamiltonian cycle of $\mathcal{F}(G_1)$ containing \mathbf{f} and \mathbf{f}' .

$G = G_2$: The only possible conflicts are between x and a short flip with root v or w . Thus, up to isomorphism of the graph, the only instance of Theorem 23 fitting to the configuration we are looking at is given by



Observe that there exists a bridge B with root $\{u, x\}$. Notice that the fixed root subgraph $\mathcal{F}_w(G_1)$ is isomorphic to the graph associahedron of a connected graph on 4 vertices so that Proposition 32(d) ensures that there exists a Hamiltonian cycle \mathcal{H}_w of the flip graph $\mathcal{F}_w(G_2)$ containing the short flip \mathbf{f} and three short flips $B_v[w], B_u[w], B_y[w]$ of some bridges B_v, B_u, B_y whose respective roots are $\{w, v\}, \{w, u\}, \{w, y\}$. Moreover Proposition 32(a) ensures that there exists a Hamiltonian cycle \mathcal{H}_y (resp. \mathcal{H}_x) of the flip graph $\mathcal{F}_y(G_2)$ (resp. $\mathcal{F}_x(G_2)$) containing the short flip $B_y[y]$ (resp. $B[x]$). Proposition 32(a) again gives us a Hamiltonian cycle \mathcal{H}_u (resp. \mathcal{H}_v) of the flip graph $\mathcal{F}_u(G_2)$ (resp. $\mathcal{F}_v(G_2)$) containing the two short flips $B_u[u]$ and $B[u]$ (resp. \mathbf{f}' and $B_v[v]$). Note that the short flips of the bridges are all distinct since u, v and w do not disconnect the graph. Gluing all the Hamiltonian cycles of the fixed root subgraphs along the bridges as explained in Section 4.1 gives a Hamiltonian cycle of $\mathcal{F}(G_2)$ containing \mathbf{f} and \mathbf{f}' .

$G = G_3$: The analysis is identical to the case $G = G_1$.

$G = G_4$: The analysis is identical to the case $G = G_2$.

We now suppose that G has 5 vertices and that $D = \{\{x, y\}, \{x, z\}\}$. Since all edges either contain x or both y and z , G is one of the following graphs:

$$G_5 = \begin{array}{c} u & & v \\ & \diagdown & / \\ & x & \\ & / & \diagdown \\ y & \text{---} & z \end{array} \quad G_6 = \begin{array}{c} u & & v \\ & \diagdown & / \\ & x & \\ & / & \diagdown \\ y & \text{---} & z \end{array} .$$

We note that in both of them, the only possible conflicts are between x and short flips with root either u or v . Indeed, $\{x, u\}$ and $\{x, v\}$ are the only pairs of vertices disjoint from exactly one edge, and the fixed-root subgraphs $\mathcal{F}_x(G_5)$ and $\mathcal{F}_x(G_6)$ are reduced to single flips. Using Remark 28, we can restrict to the cases in which $x \notin \{v_1, v_5\}$. Again we treat separately the two graphs:

$G = G_5$: Notice that the fixed-root subgraphs $\mathcal{F}_y(G_5)$ and $\mathcal{F}_z(G_5)$ both are isomorphic to the flip graph of a star on 4 vertices with central vertex x . So given a short flip \mathbf{h} (resp. \mathbf{k}) with roots y (resp. z), Proposition 31 provides us with a Hamiltonian cycle \mathcal{H}_y (resp. \mathcal{H}_z) of $\mathcal{F}_y(G_5)$ (resp. $\mathcal{F}_z(G_5)$) containing \mathbf{h} (resp. \mathbf{k}) and the flip of $\mathcal{F}_y(G_5)$ (resp. $\mathcal{F}_z(G_5)$) corresponding to the long flip of $\mathcal{F}(G_5[\hat{y}])$ (resp. $\mathcal{F}(G_5[\hat{z}])$) with root $\{x, z\}$ (resp. $\{x, y\}$). Then gluing together the cycles \mathcal{H}_y and \mathcal{H}_z and the fixed-root subgraph $\mathcal{F}_x(G_5)$ as in Figure 19 gives a tool to deal with the remaining configurations, always with the strategy of gluing Hamiltonian cycles of the fixed-root subgraphs along bridges.

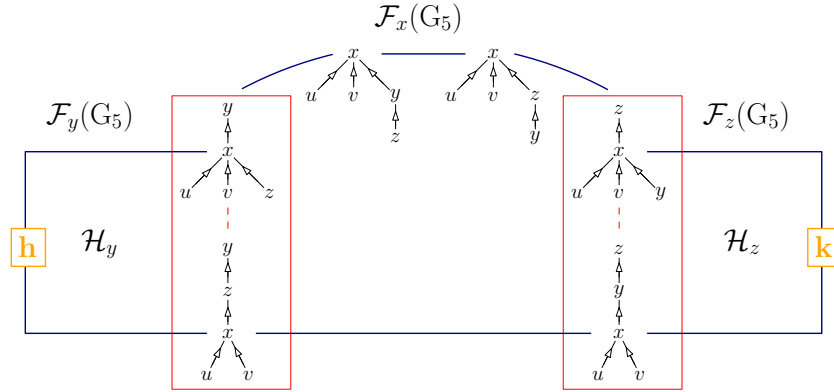


FIGURE 19. How to glue together the flip graphs $\mathcal{F}_y(G_5)$, $\mathcal{F}_x(G_5)$ and $\mathcal{F}_z(G_5)$.

$G = G_6$: Observe that both fixed-root subgraphs $\mathcal{F}_u(G_6)$ and $\mathcal{F}_v(G_6)$ are isomorphic to the classical (path) associahedron. Thus as soon as one of the short flips \mathbf{f} and \mathbf{f}' is not in conflict with x , one can find an arrangement of the vertices in the same way as when we treated the graph G_2 and G_4 (without the intermediary of the vertex u) which always makes our strategy work. We thus only need to deal with the case where x is in conflict with both \mathbf{f} and \mathbf{f}' , which corresponds to a single instance of Theorem 23, checked by hand in Figure 20.

4.5.3. *Graphs with 6 vertices.* To finish, we need to deal with the case where G has 6 vertices, $D = \{\{x, y\}, \{x, z\}\}$ and x is in conflict with both \mathbf{f} and \mathbf{f}' . Again G can only be one of the two following graphs:

$$G_7 = \begin{array}{c} u & v & w \\ & \diagdown & / \\ & x & \\ & / & \diagdown \\ y & \text{---} & z \end{array} \quad G_8 = \begin{array}{c} u & v & w \\ & \diagdown & / \\ & x & \\ & / & \diagdown \\ y & \text{---} & z \end{array} .$$

The graph G_7 is treated exactly as G_5 , using Remark 28 instead of Proposition 32 to restrict the number of cases to analyze. In the case of G_8 , there is again a single difficult instance which can be treated by hand (since the graph associahedron $\text{Asso}(G_8)$ has 236 vertices, we do not include here the resulting picture).




FIGURE 20. The flip graph $\mathcal{F}(G_6)$ represented as the 1-skeleton of the graph associahedron $\text{Asso}(G_6)$, visualized by its Schlegel diagram. The (blue) Hamiltonian cycle passes through the only two short flips in conflict with x (in red).

ACKNOWLEDGMENTS

We thank F. Santos for pointing out that the diameter of the graphical zonotope is a lower bound for that of the graph associahedron, which led to the lower bound in Theorem 18.

REFERENCES

[BFS90] Louis J. Billera, Paul Filliman, and Bernd Sturmfels. Constructions and complexity of secondary polytopes. *Adv. Math.*, 83(2):155–179, 1990.

[CD06] Michael P. Carr and Satyan L. Devadoss. Coxeter complexes and graph-associahedra. *Topology Appl.*, 153(12):2155–2168, 2006.

[CFZ02] Frédéric Chapoton, Sergey Fomin, and Andrei Zelevinsky. Polytopal realizations of generalized associahedra. *Canad. Math. Bull.*, 45(4):537–566, 2002.

[CP16] C. Ceballos and V. Pilaud. The diameter of type D associahedra and the non-leaving-face property. *European J. Combin.*, 51:109–124, 2016.

[CSZ15] Cesar Ceballos, Francisco Santos, and Günter M. Ziegler. Many non-equivalent realizations of the associahedron. *Combinatorica*, 2015. DOI:10.1007/s00493-014-2959-9.

[Deh10] Patrick Dehornoy. On the rotation distance between binary trees. *Adv. Math.*, 223(4):1316–1355, 2010.

[Dev09] Satyan L. Devadoss. A realization of graph associahedra. *Discrete Math.*, 309(1):271–276, 2009.

- [FS05] Eva Maria Feichtner and Bernd Sturmfels. Matroid polytopes, nested sets and Bergman fans. *Port. Math. (N.S.)*, 62(4):437–468, 2005.
- [GKZ08] Israel Gelfand, Mikhail M. M. Kapranov, and Andrei Zelevinsky. *Discriminants, resultants and multidimensional determinants*. Modern Birkhäuser Classics. Birkhäuser Boston Inc., Boston, MA, 2008. Reprint of the 1994 edition.
- [HL07] Christophe Hohlweg and Carsten Lange. Realizations of the associahedron and cyclohedron. *Discrete Comput. Geom.*, 37(4):517–543, 2007.
- [HLT11] Christophe Hohlweg, Carsten Lange, and Hugh Thomas. Permutahedra and generalized associahedra. *Adv. Math.*, 226(1):608–640, 2011.
- [HN99] Ferran Hurtado and Marc Noy. Graph of triangulations of a convex polygon and tree of triangulations. *Comput. Geom.*, 13(3):179–188, 1999.
- [Joh63] Selmer M. Johnson. Generation of permutations by adjacent transposition. *Math. Comp.*, 17:282–285, 1963.
- [Lee89] Carl W. Lee. The associahedron and triangulations of the n -gon. *European J. Combin.*, 10(6):551–560, 1989.
- [Lod04] Jean-Louis Loday. Realization of the Stasheff polytope. *Arch. Math. (Basel)*, 83(3):267–278, 2004.
- [LR98] Jean-Louis Loday and María O. Ronco. Hopf algebra of the planar binary trees. *Adv. Math.*, 139(2):293–309, 1998.
- [Luc87] Joan M. Lucas. The rotation graph of binary trees is Hamiltonian. *J. Algorithms*, 8(4):503–535, 1987.
- [Pos09] Alexander Postnikov. Permutohedra, associahedra, and beyond. *Int. Math. Res. Not. IMRN*, (6):1026–1106, 2009.
- [Pou14] Lionel Pournin. The diameter of associahedra. *Adv. Math.*, 259:13–42, 2014.
- [PS12] Vincent Pilaud and Francisco Santos. The brick polytope of a sorting network. *European J. Combin.*, 33(4):632–662, 2012.
- [PS15] V. Pilaud and C. Stump. Brick polytopes of spherical subword complexes and generalized associahedra. *Adv. Math.*, 276:1–61, 2015.
- [Sta63] Jim Stasheff. Homotopy associativity of H-spaces I, II. *Trans. Amer. Math. Soc.*, 108(2):293–312, 1963.
- [Ste64] Hugo Steinhaus. *One hundred problems in elementary mathematics*. Basic Books Inc. Publishers, New York, 1964.
- [STT88] Daniel D. Sleator, Robert E. Tarjan, and William P. Thurston. Rotation distance, triangulations, and hyperbolic geometry. *J. Amer. Math. Soc.*, 1(3):647–681, 1988.
- [Tro62] H. F. Trotter. Algorithm 115: Perm. *Commun. ACM*, 5(8):434–435, 1962.
- [Zel06] Andrei Zelevinsky. Nested complexes and their polyhedral realizations. *Pure Appl. Math. Q.*, 2(3):655–671, 2006.
- [Zie95] Günter M. Ziegler. *Lectures on polytopes*, volume 152 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1995.

LIX, ÉCOLE POLYTECHNIQUE, PALAISEAU
E-mail address: thibault.manneville@lix.polytechnique.fr
URL: <http://www.lix.polytechnique.fr/~manneville/>

CNRS & LIX, ÉCOLE POLYTECHNIQUE, PALAISEAU
E-mail address: vincent.pilaud@lix.polytechnique.fr
URL: <http://www.lix.polytechnique.fr/~pilaud/>