3 Max-slope pivot rule polytopes

J'ai appris que la voie du progrès n'était ni rapide ni facile.
Dans la vie, rien n'est à craindre, tout est à comprendre.
– Marie Curie

3.1 Max-slope pivot rule and max-slope pivot polytope

For a linear program $(\mathsf{P}, \boldsymbol{c})$, the choice of a $\boldsymbol{c}$-improving neighbor of a vertex $\boldsymbol{v} \in V(\mathsf{P})$ is determined by the pivot rule adopted by the simplex method, see Section 1.3. Beside varying the objective function of a linear program, one can wonder about the behavior of different pivot rules. Recall that the pivot rule is called *memoryless* if the choice is deterministic and based only on the knowledge of $\boldsymbol{v}$. Among pivot rules, max-slope pivot rules are of a great theoretical and practical importance. First observe that linear programming is easy in dimension 2: for a polygon in the plane, the monotone path chosen by the simplex method is either the upper path of the polygon or its lower path, see the exterior of Figure 25(Left). Thus, for a linear program $(\mathsf{P}, \boldsymbol{c})$, the idea of a *max-slope pivot rule* is to choose a secondary vector $\boldsymbol{\omega}$, linearly independent with $\boldsymbol{c}$, then to project P onto the plane defined by $(\boldsymbol{c}, \boldsymbol{\omega})$. The path chosen by the simplex method is defined, by convention, to be the upper path of this 2-dimensional projection of P.

Thereby, a max-slope pivot rule depends on one parameter $\boldsymbol{\omega}$, but several $\boldsymbol{\omega}$ can give the same monotone path. A monotone path $\mathcal{P}$ on P is said to be coherent if there exists $\boldsymbol{\omega}$ such that $\mathcal{P}$ is the path followed by the simplex method with the max-slope pivot rule associated to $\boldsymbol{\omega}$, see Figure 25(Left). The set of coherent paths on P can be seen as the set of vertices of a polytope, called its *monotone path polytope* $\mathsf{M}_{\boldsymbol{c}}(\mathsf{P})$, we will discuss this construction in more details in Section 4.2.1.

When studying the monotone path polytope, one focuses on the behavior of the pivot rule only on the monotone path that the simplex method walks on when starting from the worst vertex (the vertex $\boldsymbol{v}_0$ minimizing $\langle \boldsymbol{v}, \boldsymbol{c}\rangle$), and going towards the optimal vertex (the vertex $\boldsymbol{v}_{\text{opt}}$ maximizing $\langle \boldsymbol{v}, \boldsymbol{c}\rangle$). But besides this, one can also look at all the monotone paths at once, *i.e.* the combinatorial behavior the pivot rule has on all the vertices of the polytope, see [BDLLS22]: each vertex points towards the one of its neighbor improving $\langle \boldsymbol{v}, \boldsymbol{c}\rangle$ that maximizes the slope $\frac{\langle \boldsymbol{\omega}, \boldsymbol{u}-\boldsymbol{v}\rangle}{\langle \boldsymbol{c}, \boldsymbol{u}-\boldsymbol{v}\rangle}$. As depicted in Figure 25(Right), for a given linear program $(\mathsf{P}, \boldsymbol{c})$, and a secondary direction $\boldsymbol{\omega}$, the associated arborescence is a function $A^{\boldsymbol{\omega}} : V(\mathsf{P}) \smallsetminus \{\boldsymbol{v}_{\text{opt}}\} \to V(\mathsf{P})$ defined by (where "argmax" designate the unique maximizer of the studied quantity):

$$A^{\boldsymbol{\omega}}(\boldsymbol{v}) = \operatorname{argmax}\left\{ \frac{\langle \boldsymbol{\omega}, \boldsymbol{u}-\boldsymbol{v}\rangle}{\langle \boldsymbol{c}, \boldsymbol{u}-\boldsymbol{v}\rangle} \ ;\ \boldsymbol{u} \text{ improving neighbor of } \boldsymbol{v} \right\}$$

Conversely, a function $A : V(\mathsf{P}) \smallsetminus \{\boldsymbol{v}_{\text{opt}}\} \to V(\mathsf{P})$ is said to be a *coherent arborescence* or a *max-slope arborescence* when there exists $\boldsymbol{\omega}$ such that $A = A^{\boldsymbol{\omega}}$. Note that coherent arborescences are necessary monotone in the sense that $\langle A(\boldsymbol{v}), \boldsymbol{c}\rangle > \langle \boldsymbol{v}, \boldsymbol{c}\rangle$ for all $\boldsymbol{v} \in V(P) \smallsetminus \{\boldsymbol{v}_{\text{opt}}\}$, we call *arborescence* a function that satisfies this monotonicity property, and extend A to $V(\mathsf{P})$ by setting $A(\boldsymbol{v}_{\text{opt}}) = \boldsymbol{v}_{\text{opt}}$ when convenient. As for coherent monotone paths, the set of coherent arborescences can be embedded as the vertices of a polytope. We give several ways to construct this polytope. To begin with, one can construct a fan whose maximal cones are $\mathsf{C}_A = \{\boldsymbol{\omega} \ ;\ A^{\boldsymbol{\omega}} = A\}$ for A a (coherent) arborescence on P: Figure 26 shows how to gather all coherent arborescences for the case of the 3-dimensional simplex, while Figure 27 pictures the resulting fan.

For a fixed linear program $(\mathsf{P}, \boldsymbol{c})$ and an arborescence $A : V(\mathsf{P}) \to V(\mathsf{P})$ define

$$\Psi(A) \ :=\ \sum_{v \neq \boldsymbol{v}_{\text{opt}}} \frac{1}{\langle \boldsymbol{c}, A(v) - v\rangle}(A(v) - v)\,.$$

The *max-slope pivot rule polytope* is defined as

$$\Pi(\mathsf{P}, \boldsymbol{c}) \ :=\ \operatorname{conv}\{\Psi(A) : A \text{ arborescence of } (\mathsf{P}, \boldsymbol{c})\}\,.$$

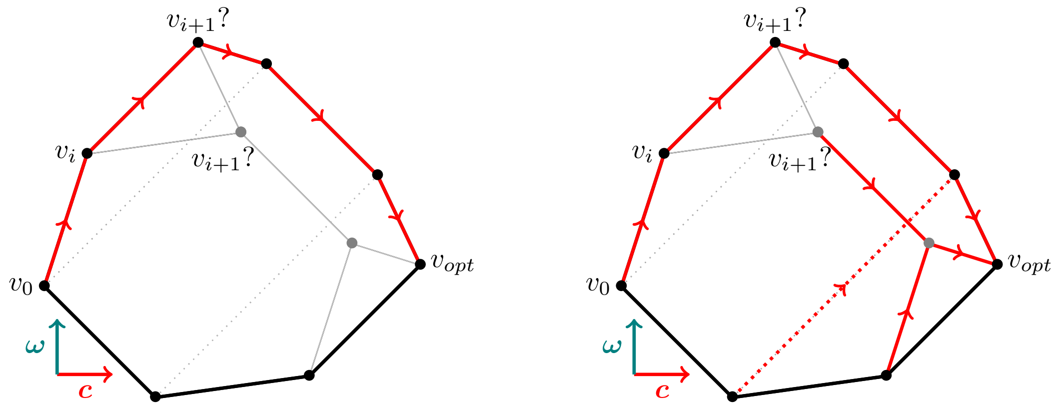

Figure 25: (Left) In red, the coherent monotone path associated to the parameter $\boldsymbol{\omega}$ for the linear problem $(\mathsf{P}, \boldsymbol{c})$. (Right) In red, the coherent arborescence associated to the same parameter. In both figures, to each vertex, the pivot rule associates the one of its neighbors that maximizes the slope in the plane $(\boldsymbol{c}, \boldsymbol{\omega})$.

The following special case of [BDLLS22, Theorem 1.4] states that the pivot rule polytope captures the combinatorics of max-slope arborescences, as its normal fan is the fan constructed above. For a polytope $\mathsf{Q} \subset \mathbb{R}^d$ and $\boldsymbol{\omega} \in \mathbb{R}^d$, we denote as usual (see Section 1.2) $\mathsf{Q}^{\boldsymbol{\omega}} = \{x \in \mathsf{Q} \;;\; \langle \boldsymbol{\omega}, x\rangle \geq \langle \boldsymbol{\omega}, y\rangle \text{ for all } y \in \mathsf{Q}\}$ the face of Q that maximizes $\boldsymbol{\omega}$.

Theorem 3.1 ([BDLLS22])**.** *Let $\mathsf{P} \subset \mathbb{R}^d$ be a polytope of dimension k and $\boldsymbol{c} \in \mathbb{R}^d$ a generic objective function. The polytope $\Pi(\mathsf{P}, \boldsymbol{c})$ is of dimension $k-1$ and for any generic $\boldsymbol{\omega} \in \mathbb{R}^d$: $\Pi(\mathsf{P}, \boldsymbol{c})^{\boldsymbol{\omega}} = \{\Psi(A^{\boldsymbol{\omega}})\}$.*

In particular, the vertices of $\Pi(\mathsf{P}, \boldsymbol{c})$ are in bijection to max-slope arborescences of $(\mathsf{P}, \boldsymbol{c})$. The faces of $\Pi(\mathsf{P}, \boldsymbol{c})$ are in correspondence to certain *multi-arborescences*, that is, maps $\mathcal{A} : V(\mathsf{P}) \to 2^{V(\mathsf{P})}$ such that for all $\boldsymbol{v} \neq \boldsymbol{v}_{\text{opt}}$, $\mathcal{A}(\boldsymbol{v})$ is a nonempty subset of $\boldsymbol{c}$-improving neighbors of $\boldsymbol{v}$, see Section 3.3.

The max-slope pivot polytope can also be constructed as a sum of sections of P, see Figure 28: for each vertex $\boldsymbol{v} \in V(\mathsf{P})$, consider the convex hull $\operatorname{conv}\big(\{\boldsymbol{v}\} \cup \{\boldsymbol{u} \;;\; \boldsymbol{u} \text{ improving neighbor of } \boldsymbol{v}\}\big)$ and take a section orthogonal to $\boldsymbol{c}$ close to $\boldsymbol{v}$ (which correspond to the vertex figure at $\boldsymbol{v}$). The Minkowski sum of these sections for $\boldsymbol{v} \in V(\mathsf{P}) \setminus \{\boldsymbol{v}_{\text{opt}}\}$ is (a dilate of) $\Pi(\mathsf{P}, \boldsymbol{c})$. For instance, Figure 28 illustrates the fact that the max-slope pivot polytope of any simplex (for any generic objective function) is an associahedron [BDLLSon], see also Section 3.3.1 for a self-contained proof of this fact. We give here a first description of this result. Let S be a simplex of dimension $n-1$. We may assume that the vertices $\boldsymbol{v}_1, \boldsymbol{v}_2, \dots, \boldsymbol{v}_n$ are labelled in such a way that $\langle \boldsymbol{c}, \boldsymbol{v}_1\rangle < \langle \boldsymbol{c}, \boldsymbol{v}_2\rangle < \dots < \langle \boldsymbol{c}, \boldsymbol{v}_n\rangle$. As $V(\mathsf{S})$ is in bijection with $[n]$, an arborescence of $(\mathsf{S}, \boldsymbol{c})$ is encoded by a map $A : [n] \to [n]$ such that $A(n) = n$ and $A(i) > i$ for $i = 1, \dots, n-1$. We sometimes identify A with the collection of pairs $(i, A(i))$ and write $(i,k) \in A$ if $A(i) = k$. We call an arborescence $A : [n] \to [n]$ *non-crossing* if for all $i < j$ if $j < A(i)$, then $A(j) \leq A(i)$. In other words, there are no $i < a < j < b$ such that $(i,j), (a,b) \in A$, see Figure 32(Right) for an example. Non-crossing arborescences form a Catalan family in the sense of Section 1.2.4, we detail here some of their properties. It is straightforward, that for any polytope P whose graph is the complete graph, all coherent arborescences on P are non-crossing, see Figure 29 for an illustration.

Theorem 3.2 ([BDLLSon])**.** *Let S be an $(n-1)$-simplex and $\boldsymbol{c}$ a generic objective function. An arborescence A is a max-slope arborescence for $(\mathsf{S}, \boldsymbol{c})$ if and only if A is non-crossing. Moreover, $\Pi(\mathsf{S}, \boldsymbol{c})$ is combinatorially isomorphic to Asso_{n-2}.*

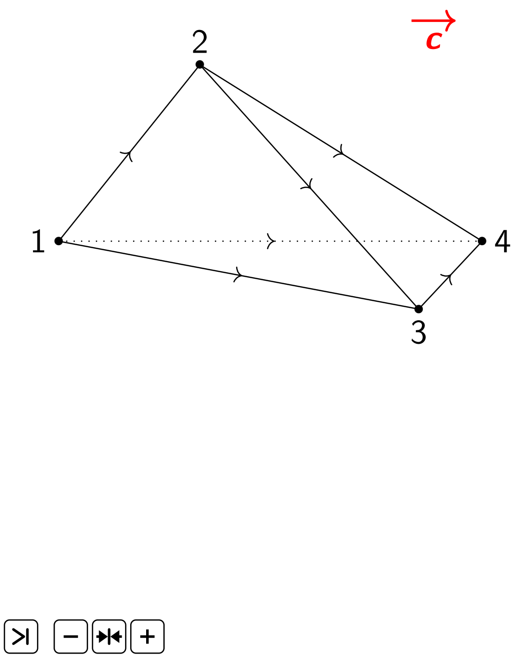

Figure 26: Animated construction of the normal fan of the max-slope pivot polytope of the 3-dimensional simplex. For each $\boldsymbol{\omega} \in \mathbb{R}^3$ orthogonal to $\boldsymbol{c}$, we project Δ_3 onto the plane $(\boldsymbol{c}, \boldsymbol{\omega})$ (Left), and record the corresponding coherent arborescence (Right). *(Animated figures obviously do not display on paper, and some PDF readers do not support the format: it is advised to use Adobe Acrobat Reader. If no solution is suitable, the animation can be found on my website or asked by email.)*

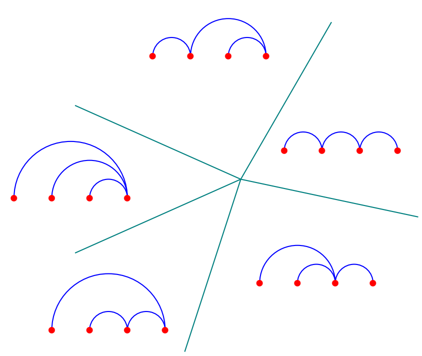

Figure 27: Normal fan of $\Pi(\Delta_3, \boldsymbol{c})$ where each maximal cone is labelled by the corresponding coherent arborescence.

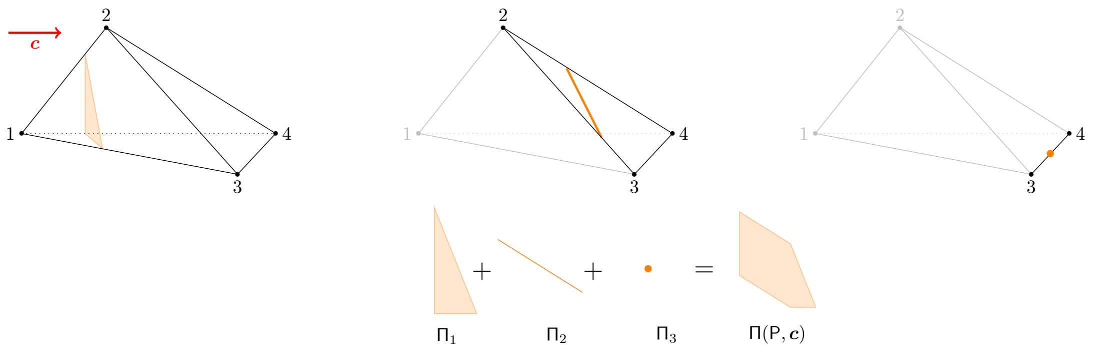

Figure 28: The max-slope pivot rule polytope $\Pi(\mathsf{P}, \boldsymbol{c})$ can be obtained as (a dilate of) the Minkowski sum of sections, here illustrated for the tetrahedron $\mathsf{P} = \Delta_3$. For each vertex $\boldsymbol{v} \in V(\mathsf{P})$, consider the convex hull that $\boldsymbol{v}$ forms with its improving neighbors, and take a section of it orthogonal to $\boldsymbol{c}$ (close to $\boldsymbol{v}$). Their sum is normally equivalent to $\Pi(\mathsf{P}, \boldsymbol{c})$.

Example 4.3 in [BDLLS22] illustrates the max-slope polytope of a simplex. Theorem 3.2 is discussed in the more general context of pivot associahedra in [BDLLSon]. Here, we give the main results and ideas. We will use the following simple decomposition of non-crossing arborescences.

Proposition 3.3. *For a non-crossing arborescence $A : [n] \to [n]$ define $r(A)$ as the minimal $r \geq 1$ such that $A(r) = n$. Restricting A to $[r]$ and to $[r+1, n] := \{r+1, \dots, n\}$ yields two non-crossing arborescences A' and A'' and A is uniquely determined by (r, A', A'').*

Proposition 3.3 immediately gives a recursive way to compute the number of non-crossing arborescences (identical to the recursive formula of Section 1.2.4), which shows that there are Catalan-many non-crossing arborescences. To check that $\Pi(\mathsf{P}, \boldsymbol{c})$ is indeed combinatorially isomorphic to the associahedron, it suffices to determine the graph of $\Pi(\mathsf{P}, \boldsymbol{c})$ and use the fact that simple polytopes are determined by their graph [BML87]; see also [Zie98, Chapter 3.4]. We call an $i < n$ *forward-sliding* if $A(i) \neq n$ and there is no $j < i$ with $A(j) = A(i)$. We call i *backward-sliding* if $A(i) \neq i+1$. If i is forward-sliding, then define $F_i A$ by $F_i A(i) := A(A(i))$ and $F_i A(k) := A(k)$ for all $k \neq i$. Likewise, if i is backward-sliding, then define $B_i A(i) := j$ where $j > i$ is minimal with $A(j) = A(i)$ and $B_i A(k) = A(k)$ for $k \neq i$. A quick scribble reveals that both $B_i A$ and $F_i A$ are non-crossing, and that $B_i F_i A = A$ and $F_i B_i A = A$. As for all Catalan families, we say that $B_i A$ and $F_i A$ differ from A by a *flip*. To summarize, non-crossing arborescences form a Catalan family, and flips in this family are forward- or backward-slide.

The following proposition is adapted from [BDLLSon] or Theorem 3.69.

Proposition 3.4. *Let $A_1, A_2 : [n] \to [n]$ be non-crossing arborescences. Then $[\Psi(A_1), \Psi(A_2)]$ is an edge of $\Pi(\mathsf{S}, \boldsymbol{c})$ if and only if A_1, A_2 differ by a flip.*

As for the proof's strategy of Theorem 3.2, let us first note that up to linear transformation, we may assume that $\mathsf{S} = \Delta_{n-1} := \operatorname{conv}(\boldsymbol{e}_1, \dots, \boldsymbol{e}_n) \subset \mathbb{R}^n$. For a given $\boldsymbol{\omega} \in \mathbb{R}^n$, define $p_i = (c_i, \omega_i)$ for $i \in [n]$ and the *slope* $\tau(i,j) = \frac{\omega_j - \omega_i}{c_j - c_i}$ for all $1 \leq i < j \leq n$. Then A is a max-slope arborescence on S if and only if there is $\boldsymbol{\omega} \in \mathbb{R}^n$ such that

$$\tau(i, A(i)) \;>\; \tau(i,k) \quad \text{for all } k > i \text{ and } k \neq A(i)$$

Pictorially, consider the points $p_1, \dots, p_n \in \mathbb{R}^2$. Then $A(i) = j$ if all points p_k for $i < k \neq j$ are strictly below the line $\overline{p_i p_j}$, see Figure 32(Left). This perspective yields the following key insight:

Lemma 3.5. *For $1 \leq r < s < t \leq n$*

$$\tau(r,t) > \tau(r,s) \iff \tau(s,t) > \tau(r,t) \quad \text{and} \quad \tau(r,t) < \tau(r,s) \iff \tau(s,t) < \tau(r,t)$$

Proof. The following convex combination gives the result: $\tau(r,t) = \frac{c_r - c_s}{c_t - c_r}\tau(r,s) + \frac{c_t - c_r}{c_t - c_r}\tau(s,t)$. $\square$

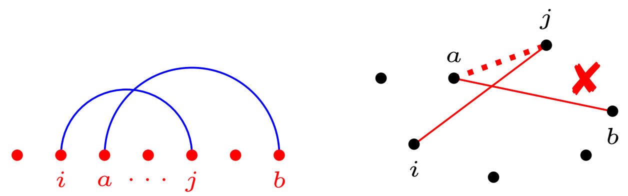

Figure 29: Coherent arborescences on a polytope whose graph is complete are non-crossing. For $i < a < j < b$ with $A(i) = j$, then Lemma 3.5 ensures that $\frac{\langle \boldsymbol{\omega}, \boldsymbol{v}_j - \boldsymbol{v}_a \rangle}{\langle \boldsymbol{\omega}, \boldsymbol{v}_j - \boldsymbol{v}_a \rangle} > \frac{\langle \boldsymbol{\omega}, \boldsymbol{v}_b - \boldsymbol{v}_a \rangle}{\langle \boldsymbol{\omega}, \boldsymbol{v}_b - \boldsymbol{v}_a \rangle}$, so $A(a) = b$ is impossible.

$$\begin{tikzcd}
\mathsf{P} \arrow[rr, "\pi"] \arrow[dr, "{\langle \cdot, \pi^*(\boldsymbol{c}) \rangle}"'] & & \pi(\mathsf{P}) \arrow[dl, "{\langle \cdot, \boldsymbol{c} \rangle}"] \\
& {[x, y]} &
\end{tikzcd}$$

$$\Pi(\mathsf{P}, \pi^*(\boldsymbol{c})) \xrightarrow{\quad \pi \quad} \Pi(\pi(\mathsf{P}), \boldsymbol{c})$$

Figure 30: **When the graph of P is isomorphic to the graph of** $\pi(\mathsf{P})$, then the max-slope polytope of $\pi(\mathsf{P})$ is a projection of the one of P. We denote $x = \langle \boldsymbol{v}_0, \boldsymbol{c} \rangle$ and $y = \langle \boldsymbol{v}_{\text{opt}}, \boldsymbol{c} \rangle$.

The monotone path polytope construction behaves functorially with respect to linear projections, see Section 4.1, that is, if $\pi : \mathsf{P} \to \mathsf{Q}$ is a projection of polytopes, then the monotone path polytope of Q is a projection of the one of P. This does not hold for max-slope pivot polytopes in general. However, it does hold in the special case when graphs are retained under projection, see Figure 30. This first new result will help us link the max-slope pivot rule polytope of cyclic polytopes with the associahedron (as the max-slope pivot rule polytope of the simplex).

Theorem 3.6. *Let $\mathsf{P} \subset \mathbb{R}^d$ be a polytope and $\pi : \mathbb{R}^d \to \mathbb{R}^e$ a linear projection. If $\pi(\mathsf{P})$ has the same graph as P, then for every linear function $\boldsymbol{c}$ that is generic for $\pi(\mathsf{P})$, denoting π^* the adjoint of π:*

$$\Pi(\pi(\mathsf{P}), \boldsymbol{c}) \;=\; \pi(\Pi(\mathsf{P}, \pi^*(\boldsymbol{c})))$$

Proof. Let $\mathsf{P}' = \pi(\mathsf{P})$. If P and P' have the same graph, then $\pi : V(\mathsf{P}) \to V(\mathsf{P}')$ is a bijection and we write $\pi(v) = v'$. In particular, A is an arborescence of P if and only if $A' := \pi \circ A \circ \pi^{-1}$ is an arborescence of P'. For an arborescence A on P we compute

$$\pi(\Psi(A)) = \sum_{v \in V(\mathsf{P})} \langle \pi^*(\boldsymbol{c}), A(v) - v \rangle^{-1} \pi(A(v) - v) = \sum_{v \in V(\mathsf{P})} \langle \boldsymbol{c}, A'(v') - v' \rangle^{-1} (A'(v') - v') = \Psi(A')$$

and hence $\pi(\Pi(\mathsf{P}, \pi^*(\boldsymbol{c}))) = \operatorname{conv}(\Psi(A') : A) = \Pi(\mathsf{P}', \boldsymbol{c})$. □

A polytope P is called *2-neighborly* if any two vertices of P share an edge, that is to say when the graph of P is the complete graph.

Corollary 3.7. *Let P be a 2-neighborly polytope. Then, for any generic linear function $\boldsymbol{c}$, $\Pi(\mathsf{P}, \boldsymbol{c})$ is the projection of an associahedron.*

Proof. Every polytope $\mathsf{P} = \operatorname{conv}(\boldsymbol{v}_1, \dots, \boldsymbol{v}_n) \subset \mathbb{R}^d$ is the projection of Δ_{n-1} under the linear map $\pi : \mathbb{R}^n \to \mathbb{R}^d$ with $\pi(\boldsymbol{e}_i) = \boldsymbol{v}_i$ for $i \in [n]$. Since P is 2-neighborly, the projection π retains the graph of Δ_{n-1} and the result follows from Theorem 3.6 and Theorem 3.2. □

This result motivates the next section. We will study the max-slope pivot polytopes of a well-known family of 2-neighborly polytopes: cyclic polytopes.

3.2 Max-slope pivot polytope of cyclic polytopes

This section is a joint work with Aenne Benjes and Raman Sanyal. An article is in preparation, containing this section together with Section 4.3.

For a linear program $(\mathsf{P}, \boldsymbol{c})$, we have seen that when the graph of P is the complete graph on n nodes, that is, if P is *2-neighborly*, then its max-lope pivot polytope is a projection of the associahedron Asso_{n-2}. This implies that the arborescences corresponding to the vertices of $\Pi(\mathsf{P}, \boldsymbol{c})$ will be *non-crossing*, meaning that there are no $i < j < k < l$ such that $A(i) = k$ and $A(j) = l$. Note that not all non-crossing arborescences will show up as a vertex of $\Pi(\mathsf{P}, \boldsymbol{c})$, but only a sub-family of them. Stronger even, Corollary 3.7 shows that when the vertices of P are sufficiently generic, then the faces of $\Pi(\mathsf{P}, \boldsymbol{c})$ are products of associahedra.

For $n > d \geq 4$ and $\boldsymbol{t} = (t_1 < t_2 < \cdots < t_n) \in \mathbb{R}^n$, the d-dimensional *cyclic polytope* is $\mathsf{Cyc}_d(\boldsymbol{t}) := \mathrm{conv}\{(t_i, t_i^2, \ldots, t_i^d) : i = 1, \ldots, n\}$, see Figure 31. Cyclic polytopes play the main role in the Upper Bound Theorem for polytopes [McM70] and they exhibit a number of remarkable properties. In particular, for $d \geq 4$, $\mathsf{Cyc}_d(\boldsymbol{t})$ is 2-neighborly and simplicial. Moreover, its vertices are in general position: no $d + 1$ of them belong to the same hyperplane. The linear function $\boldsymbol{x} \mapsto \langle \boldsymbol{e}_1, \boldsymbol{x} \rangle = x_1$ is generic with respect to $\mathsf{Cyc}_d(\boldsymbol{t})$. Corollary 3.7 yields that $\Pi(\mathsf{Cyc}_d(\boldsymbol{t}), \boldsymbol{e}_1)$ is a projection of Asso_{n-2} defined in terms of $\boldsymbol{t}$. Thus, for generic $\boldsymbol{t}$, its combinatorial structure corresponds to a subposet of the lattice of faces of Asso_{n-2}, and justifies the following definition.

Definition 3.8. For $n > d \geq 4$ and $\boldsymbol{t} = (t_1 < t_2 < \cdots < t_n)$, the $(d-1)$-dimensional *cyclic associahedron* $\mathsf{Asso}_{d-1}(\boldsymbol{t})$ is the max-slope pivot polytope $\Pi(\mathsf{Cyc}_d(\boldsymbol{t}), \boldsymbol{e}_1)$.

The following section is devoted to the study of cyclic associahedra. In particular, a quick numerical implementation reveals that the combinatorics of $\mathsf{Asso}_{d-1}(\boldsymbol{t})$ depends on $\boldsymbol{t}$, whereas Athanasiadis, De Loera, Reiner and Santos have shown in [ALRS00] that the combinatorics of the monotone path polytope of the cyclic polytope does not depend on $\boldsymbol{t}$. A first result (Corollary 3.13) determines the vertices of $\mathsf{Asso}_{d-1}(\boldsymbol{t})$ *as if it were not depending on* $\boldsymbol{t}$, that is we indicate which non-crossing arborescence appears as a vertex of $\mathsf{Asso}_{d-1}(\boldsymbol{t})$ for at least a $\boldsymbol{t}$. Then, in Theorem 3.16, we give a general but complex criterion for a non-crossing arborescence to appear as a vertex of $\mathsf{Asso}_{d-1}(\boldsymbol{t})$ for a given $\boldsymbol{t}$. These two theorems allow a full description of the case $d = 3$: intriguingly, the number of vertices of $\mathsf{Asso}_2(\boldsymbol{t})$ does not depend on $\boldsymbol{t}$, see Corollaries 3.37 and 3.49.

To determine which non-crossing arborescences correspond to a vertex of $\mathsf{Asso}_{d-1}(\boldsymbol{t})$, we propose a general perspective in elementary geometric terms: For a univariate polynomial $P(t) = w_1 t + w_2 t^2 + \cdots + w_d t^d$, consider the n points $p_i = (t_i, P(t_i)) \in \mathbb{R}^2$. Define $A : [n] \to [n]$ by the property that $A(i) = j$ if $j > i$ and all points p_k for $i < k \neq j$ are below the line $\overline{p_i p_j}$; see Figure 32 for an illustration. This defines a non-crossing arborescence, and we say that P *captures* A on $\boldsymbol{t}$. Thus, if $d \geq 4$, a non-crossing arborescence A correspond to a vertex of $\mathsf{Asso}_{d-1}(\boldsymbol{t})$ if and only if it can be captured on $\boldsymbol{t}$ by a polynomial of degree at most d.

Since $\mathsf{Cyc}_n(\boldsymbol{t})$ is a $(n-1)$-simplex, every non-crossing arborescence is captured by some polynomial on $\boldsymbol{t}$. We define the *degree* $\mu(A, \boldsymbol{t})$ of a non-crossing arborescence A as the minimal degree of a polynomial P that captures A on $\boldsymbol{t}$. In general, the degree of A depends on the choice of $\boldsymbol{t} = (t_1 < \cdots < t_n)$. We define the *intrinsic degree* of A as $\mu(A) := \min_{\boldsymbol{t}} \mu(A, \boldsymbol{t})$. The intrinsic degree defines a natural complexity measure on non-crossing arborescences and hence on Catalan families. We show that $\mu(A)$ can be determined directly from the combinatorics of A, see Corollary 3.13.

Even though our motivation originally comes from the study of the geometry of pivot rules, it also advocates for new ways to think about the associahedron. In fact, (combinatorial) understandings and realizations of the associahedron already naturally arise in a wide variety of domains: the associahedron occurs as a secondary polytope [Lee89], as a generalized permutahedron [Lod04, Pos09], or as a polytope of expansive motions [RSS03]. In this section, the realizations of the associahedron we introduce are parametrized by $\boldsymbol{t} \in \mathbb{R}^n$, and generalized to more convoluted structures whose genesis prompts a natural complexity parameter on Catalan families (parenthe-

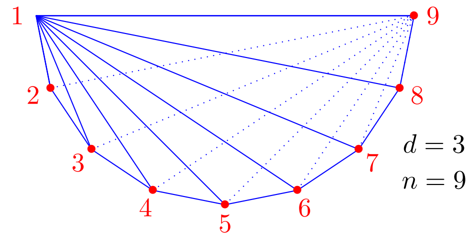

Figure 31: The cyclic polytope $\mathsf{Cyc}_3(\boldsymbol{t})$ for $n=9$. Note that its graph is not complete (the graph of $\mathsf{Cyc}_d(\boldsymbol{t})$ is complete for $d \geq 4$).

sations, binary search trees, triangulations of polygons...). Moreover, the tools developed here will be of prime importance for the study of fiber polytopes in Section 4.3.

If A is a max-slope arborescence of $(\mathsf{P}, \boldsymbol{c})$, then the leading path from the minimal to the maximal vertex of $\boldsymbol{c}$ over P is a coherent $\boldsymbol{c}$-monotone path in the sense of Billera–Sturmfels [BS92]. The monotone path polytope $\mathsf{M}_{\boldsymbol{c}}(\mathsf{P})$ gives a geometric model for coherent monotone paths, and one can prove it captures the homotopy type of the Baues poset [BKS94]. It was shown in [BDLLS22] that $\mathsf{M}_{\boldsymbol{c}}(\mathsf{P})$ is a weak Minkowski summand of $\Pi(\mathsf{P}, \boldsymbol{c})$: the monotone path polytope is a deformation of the max-slope pivot polytope, see Section 4.2.1 for the details. On the combinatorial level, this implies that $\mathsf{Asso}_{d-1}(\boldsymbol{t})$ refines the combinatorics of $\mathsf{M}_{\boldsymbol{e}_1}(\mathsf{Cyc}_d(\boldsymbol{t}))$. The latter was studied by Athanasiadis, De Loera, Reiner and Santos in [ALRS00]: the proof of Corollary 3.13 is related to the relative locations of roots of $P''(t)$ induced by the combinatorics of A and is inspired by their work. There, the authors show that the combinatorics of $\mathsf{M}_{\boldsymbol{e}_1}(\mathsf{Cyc}_d(\boldsymbol{t}))$ is that of cyclic zonotopes and, in particular, independent of $\boldsymbol{t}$. In our case, the combinatorics of the polytope $\mathsf{Asso}_{d-1}(\boldsymbol{t})$ will depend on $\boldsymbol{t}$. This motivates the notion of universal non-crossing arborescences.

A non-crossing arborescence A is *universal* if for any $\boldsymbol{t} = (t_1 < t_2 < \cdots < t_n)$ there is a polynomial of degree $\mu(A)$ that captures A on $\boldsymbol{t}$ (that is $\mu(A, \boldsymbol{t}) = \mu(A)$ for all $\boldsymbol{t}$). For $d \geq \max(4, \mu(A))$ this implies that A is always a vertex of $\mathsf{Asso}_{d-1}(\boldsymbol{t})$ whatever $\boldsymbol{t}$. In Section 3.2.2 we completely classify universal arborescences of intrinsic degree 3 and smaller. To that end, we study realization sets of A, that is, the set of $\boldsymbol{t} = (t_1 < \cdots < t_n)$ such that A is captured by a polynomial of a given degree. Moreover, we prove in Corollary 3.37 that the number of non-crossing arborescence A captured on $\boldsymbol{t}$ in degree 3 or smaller is $\binom{n}{2} - 1$, independently of $\boldsymbol{t}$.

Note however that when $d \leq 3$, the max-slope pivot polytope $\mathsf{Asso}_{d-1}(\boldsymbol{t})$ is not a projection of the associahedron, the special cases $\mathsf{Asso}_1(\boldsymbol{t})$ and $\mathsf{Asso}_2(\boldsymbol{t})$ will be studied in Section 3.2.3. We prove there that the number of vertices of $\mathsf{Asso}_1(\boldsymbol{t})$ and $\mathsf{Asso}_2(\boldsymbol{t})$ are independent from $\boldsymbol{t}$: they are respectively 2 and $3n-7$, see Theorem 3.39 and Corollary 3.49.

3.2.1 Cyclic associahedra and the intrinsic degree

We start by fixing the notations of what we have swiftly introduced above.

For $d \geq 2$ and $\boldsymbol{t} = (t_1, t_2, \ldots, t_n) \in \mathbb{R}^n$ with $t_1 < t_2 < \cdots < t_n$, the d-dimensional *cyclic polytope* $\mathsf{Cyc}_d(\boldsymbol{t})$ is defined as the convex hull of $\boldsymbol{\gamma}_d(t_1), \ldots, \boldsymbol{\gamma}_d(t_n)$ where $\boldsymbol{\gamma}_d(t) := (t, t^2, \ldots, t^d)$. It is well-known that for $d \geq 4$, the cyclic polytope $\mathsf{Cyc}_d(\boldsymbol{t})$ is 2-neighborly [Zie98, Cor. 0.8]. For $\boldsymbol{c} = \boldsymbol{e}_1 = (1, 0, \ldots, 0)$, we have $\langle \boldsymbol{e}_1, \boldsymbol{\gamma}_d(t) \rangle = t$ for all $t \in \mathbb{R}$ and hence $\boldsymbol{e}_1$ is a generic linear function for $\mathsf{Cyc}_d(\boldsymbol{t})$. For $d \geq 4$, we call $\mathsf{Asso}_{d-1}(\boldsymbol{t}) := \Pi(\mathsf{Cyc}_d(\boldsymbol{t}), \boldsymbol{e}_1)$ a *cyclic associahedron*.

Proposition 3.9. *For any $d \geq 4$, if $\boldsymbol{t}$ is sufficiently generic, then the faces of $\mathsf{Asso}_{d-1}(\boldsymbol{t})$ are combinatorially isomorphic to products of associahedra.*

Proof. When $d \geq 4$, $\mathsf{Cyc}_d(\boldsymbol{t})$ is 2-neighborly for any $\boldsymbol{t}$. By Corollary 3.7, $\mathsf{Asso}_{d-1}(\boldsymbol{t})$ is the image of $\Pi(\Delta_{n-1}, \boldsymbol{t}) = \mathsf{Asso}_{n-2}$ under the projection $\pi(\boldsymbol{e}_i) = \boldsymbol{\gamma}_d(t_i)$ for $i = 1, \ldots, n$. For every $(d-2)$-face $G \subset \mathsf{Asso}_{n-2}$ let U_G be the $(d-2)$-dimensional linear subspace parallel to G. This gives a finite collection of $(d-2)$-linear subspaces. The collection of $\boldsymbol{t}$ such that π is not injective on the union

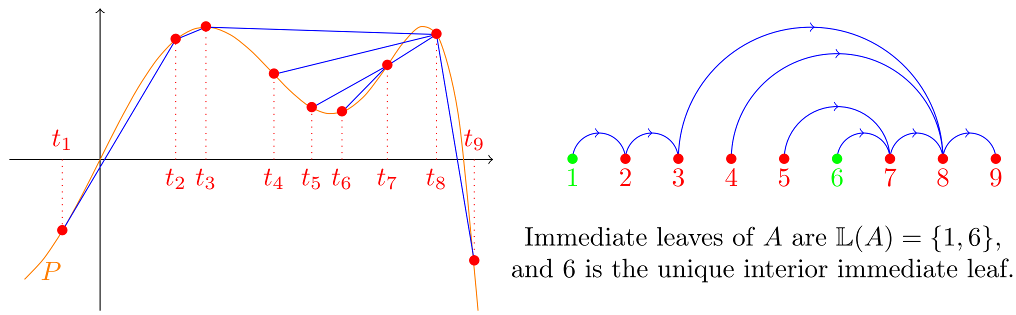

Figure 32: (Left) A 5-degree polynomial with the points p_i in red and the (pieces of) lines $\overline{p_i p_{A(i)}}$ in blue; (Right) the non-crossing arborescence A captured by P on $\boldsymbol{t}$.

of these subspaces is Zariski closed. For any $\boldsymbol{t}$ in the complement, for any facet $F \subset \mathsf{Asso}_{d-1}(\boldsymbol{t})$, the face $\pi^{-1}(F) \subset \mathsf{Asso}_{n-2}$ is linearly isomorphic to F. This proves the claim. $\square$

Proposition 3.9 also implies that for generic $\boldsymbol{t}$ and $d \geq 4$, the vertices of $\mathsf{Asso}_{d-1}(\boldsymbol{t})$ correspond to certain non-crossing arborescences $A : [n] \to [n]$ that depend on $\boldsymbol{t}$ and d. For a given (generic) $\boldsymbol{w} = (w_1, \ldots, w_d)$, there is a simple way to determine the associated non-crossing arborescence. We note that $\langle \boldsymbol{w}, \boldsymbol{\gamma}_d(t) \rangle = w_1 t + w_2 t^2 + \cdots + w_d t^d =: P_{\boldsymbol{w}}(t)$ is a univariate polynomial in t of degree at most d. Consider the graph of the function $P_{\boldsymbol{w}}(t)$ with the marked points $p_i(\boldsymbol{t}) = (t_i, P(t_i))$ for $i = 1, \ldots, n$. As in the previous section, we note that $A(i) = j$ if and only if $p_k(\boldsymbol{t})$ is below the line $\overline{p_i(\boldsymbol{t}) p_j(\boldsymbol{t})}$ for all $k > i$ with $k \neq j$. Figure 32 illustrates the construction. We say that the non-crossing arborescence A is *captured* by the polynomial $P_{\boldsymbol{w}}$ on $\boldsymbol{t}$. For $d \geq 4$, A is captured by a polynomial of degree d on $\boldsymbol{t}$ if and only if $\Psi(A)$ appears as a vertex of $\mathsf{Asso}_{d-1}(\boldsymbol{t})$. The following fact is immediate from this geometric perspective, as 'being below' a line is an open condition.

Corollary 3.10. *Let A be a non-crossing arborescence that is captured by P on $\boldsymbol{t}$. Then there is an $\varepsilon > 0$ such that A is captured by P on $\boldsymbol{t}'$ for all $\boldsymbol{t}'$ with $\|\boldsymbol{t} - \boldsymbol{t}'\|_\infty < \varepsilon$.*

Recall that the *degree* of a non-crossing arborescence A with respect to $\boldsymbol{t}$ is

$$\mu(A, \boldsymbol{t}) \ := \ \min\{\deg P \ ; \ A \text{ is captured by } P \text{ on } \boldsymbol{t}\} \, .$$

For fixed $\boldsymbol{t}$, the degree defines a filtration on non-crossing arborescences. However, as we will see in the next section, the degree of A depends on $\boldsymbol{t}$. This motivates the definition of the *intrinsic degree* of A as

$$\mu(A) \ := \ \min\{\mu(A, \boldsymbol{t}) \ ; \ \boldsymbol{t} = (t_1 < t_2 < \cdots < t_n) \in \mathbb{R}^n\} \, .$$

In the remainder of this section, we prove that the intrinsic degree can be computed efficiently and easily from the non-crossing arborescence, see Corollary 3.13. A $j \in [n-1]$ is a *leaf* of A if there is no i with $A(i) = j$. We call j an *immediate leaf* if, additionally, $A(j) = j + 1$ and denote by $\mathbb{L}(A)$ the set of immediate leaves. An immediate leaf $j \in \mathbb{L}(A)$ is *interior* if $1 < j < n - 1$, and we write $\mathbb{L}^\circ(A)$ for the interior leafs. We first prove that a lower bound on the degree $\mu(A)$:

Theorem 3.11. *Let A be a non-crossing arborescence. Then*

$$\mu(A) \ \geq \ |\mathbb{L}(A)| + |\mathbb{L}^\circ(A)| + 1 \, .$$

Proof. Let P be a polynomial and $\boldsymbol{t} = (t_1 < t_2 < \cdots < t_n)$ so that P captures A on $\boldsymbol{t}$. Recall that for $i < j$, $\tau(i, j) = \frac{P(t_j) - P(t_i)}{t_j - t_i}$ is the slope of the line connecting $(t_i, P(t_i))$ to $(t_j, P(t_j))$. Applying the mean-value theorem to $t \mapsto P(t)$, we get that for every $i = 1, \ldots, n-1$, there is $t_i < \theta_i < t_{i+1}$ with $P'(\theta_i) = \tau(i, i+1)$.

Let $i \in \mathbb{L}(A)$ with $i > 1$ and set $j = A(i-1)$. Since i is a leaf, we know that $j \geq i + 1$ and $\tau(i-1, j) > \tau(i-1, i)$. From Lemma 3.5 we obtain $\tau(i, j) > \tau(i-1, j)$. From $A(i) = i + 1$, we

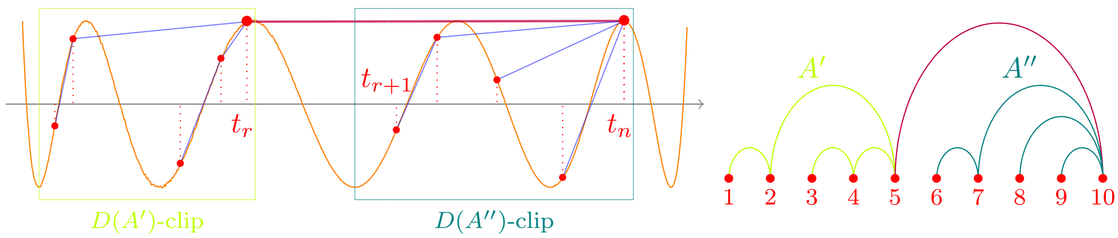

Figure 33: Decomposition of a $D(A)$-clip into $D(A')$-clip and $D(A'')$-clip for $r+1 \in \mathbb{L}(A)$.

deduce that $\tau(i,i+1) \geq \tau(i,j)$ and therefore $\tau(i,i+1) > \tau(i-1,i)$. Applying the mean-value theorem to $t \mapsto P'(t)$, we find $\theta_{i-1} < \alpha_i < \theta_i$ with $P''(\alpha_i) = \frac{P'(\theta_i)-P'(\theta_{i-1})}{\theta_i-\theta_{i-1}} = \frac{\tau(i,i+1)-\tau(i-1,i)}{\theta_i-\theta_{i-1}} > 0$.

On the other hand, for every $i < n-1$, if $A(i) = i+1$, then $\tau(i,i+1) > \tau(i,i+2)$. Lemma 3.5 yields $\tau(i,i+1) > \tau(i,i+2) > \tau(i+1,i+2)$ and there is $\theta_i < \beta_i < \theta_{i+1}$ such that $P''(\beta_i) < 0$.

If $i,j \in \mathbb{L}(A)$ with $i < j$, then $P''(\beta_i) < 0 < P''(\alpha_j)$ and, since $i \leq j-2$, we have $\beta_i < \theta_{j-1} < \alpha_j$. If there is no immediate leaf between i and j, then there is a root in the open interval (β_i, α_j). So far, this gives $|\mathbb{L}(A)| - 1$ roots of P''.

For every interior immediate leaf $1 < j < n-1$, we have $P''(\beta_j) < 0 < P''(\alpha_j)$ and $\alpha_j < \theta_j < \beta_j$ and P'' has a root in (α_j, β_j). This gives an additional $|\mathbb{L}^\circ(A)|$ roots of P''. To finish the proof, we compute $\deg P = \deg P'' + 2 \geq |\mathrm{roots}(P'')| + 2 \geq |\mathbb{L}(A)| + |\mathbb{L}^\circ(A)| + 1$. $\square$

The idea for the proof was inspired by the proof of Theorem 3.1 in [ALRS00], where the authors study monotone path polytopes of cyclic polytopes. The coherence of a $\boldsymbol{e}_1$-monotone path of $\mathsf{Cyc}_d(\boldsymbol{t})$ depends on the degree, but it was shown that it does not depend on the choice of $\boldsymbol{t}$. Unfortunately, the degree of A depends on the choice of $\boldsymbol{t}$, and in order to prove $\mu(A) \leq |\mathbb{L}(A)| + |\mathbb{L}^\circ(A)| + 1$ we need to exhibit a concrete polynomial P and $\boldsymbol{t}$ to capture A.

For $d \geq 1$, the *Chebychev polynomial* T_d of the first kind is the polynomial of degree d with the property that $T_d(\cos(\alpha)) = \cos(d\alpha)$ for all $\alpha \in [0,\pi]$. It follows that all d roots of T_d are distinct and real, and lie in the open interval $(-1,1)$. Moreover, T_d has $d+1$ extrema in the interval $[-1,1]$ and the extrema alternate between -1 and 1. In particular, $T_d(1) = 1$ and $T_d(-1) = (-1)^d$. We also note that $T_d''(1) > 0$ and $(-1)^d T_d''(-1) > 0$. Hence $(-1)^d T_d(t)$ and $T_d(t)$ are convex in a neighborhood of $t = -1$ and $t = 1$, respectively.

For $D \geq 0$, a *D-clip* of T_d is an interval $[m, M] \subseteq [-1,1]$ that contains $D+1$ extrema including m and M and M is a maximum ($T_d(M) = 1$). We call the D-clip *concave* if T_d is concave on $[M-\varepsilon, M]$ for some $\varepsilon > 0$. Any D-clip with $M < 1$ is concave.

For a non-crossing arborescence on $n \geq 1$ nodes we define $D(A) := 2|\mathbb{L}(A)|$ if $1 \notin \mathbb{L}(A)$ and $D(A) := 2|\mathbb{L}(A)| - 1$ if $1 \in \mathbb{L}(A)$. If $n = 1$, then set $D(A) := 0$.

Proposition 3.12. *Let A be a non-crossing arborescence on n nodes. For any concave $D(A)$-clip $[m, M]$ of T_d there are $m \leq t_1 < t_2 < \cdots < t_n = M$ such that A is captured by T_d on $\boldsymbol{t} = (t_1, \ldots, t_n)$ and $T_d(t_i) < 1 = T_d(t_n)$ for $i = 1, \ldots, n-1$.*

Proof. We prove the claim by induction on n. Let $[m, M]$ be a fixed concave $D(A)$-clip of T_d. For $n \leq 2$, this is clearly true. Let $n \geq 3$ and consider the decomposition of A into (r, A', A'') of Proposition 3.3. We distinguish three cases.

If $r = 1$, then A' has a single node. Moreover, $D(A) = 2|\mathbb{L}(A)|$ is even and thus m is a maximum. If $2 \in \mathbb{L}(A)$, then $D(A'') = D(A) - 1$ is odd. Let m'' be the first minimum of the concave $D(A)$-clip $[m, M]$. By induction, A'' can be captured on $[m'', M]$. The condition $T_d(t_i) < 1$ for $i < n$ ensures that for $t_1 = m + \varepsilon$ with $\varepsilon > 0$ sufficiently small, the points $p_i(\boldsymbol{t}) = (t_i, T_d(t_i))$ for $1 < i < n$ are below the line $\overline{p_1(\boldsymbol{t})p_n(\boldsymbol{t})}$. If $2 \notin \mathbb{L}(A)$, then $D(A'') = D(A)$ and A'' can be captured on $[m, M]$. Since $D(A)$ is even, m is a maximum and $m < t_2$. The same argument as before shows that choosing $t_1 \in (m, t_2)$ close enough to m suffices.

If $r = n-1$, then $D(A') = D(A)$ and A' can be captured on the concave $D(A)$-clip $[m, M]$ with $m \le t_1 < t_2 < \cdots < t_{n-1} = M$. By Corollary 3.10, we can change t_{n-1} to $t_{n-1} - \varepsilon$ for some small $\varepsilon > 0$. For $t_n = M$, the line $\overline{p_{n-1}(\boldsymbol{t})p_n(\boldsymbol{t})}$ is close to the tangent of T_d at t_n and by the concavity of the maximum, $p_n(\boldsymbol{t})$ is below all lines $\overline{p_i(\boldsymbol{t})p_{n-1}(\boldsymbol{t})}$ for $i < n-1$.

Thus, we can assume $1 < r < n-1$. If $r+1 \notin \mathbb{L}(A)$, then $D(A) = D(A') + D(A'')$ and $D(A'')$ is even. Let $y \in [m, M]$ such that $[m, y]$ is a concave $D(A')$-clip and $[y, M]$ is a concave $D(A'')$-clip. By induction there are $m \le t_1 < \cdots < t_r = y < \cdots < t_n = M$ that capture A' and A'' on their clips. Again by Corollary 3.10, $t_1 < \cdots < t_r = y - \varepsilon$ still captures A' for $\varepsilon > 0$ sufficiently small. Since $T_d(t_j) < 1$ for all $j < r$, the points $p_i(\boldsymbol{t})$ for $i > r$ are all below the lines $\overline{p_j(\boldsymbol{t})p_r(\boldsymbol{t})}$ as well as below the line $\overline{p_r(\boldsymbol{t})p_n(\boldsymbol{t})}$. If $r+1 \in \mathbb{L}(A)$, then $D(A) = D(A') + D(A'') + 1$ and $D(A'')$ is odd, see Figure 33. Choose $y \in [m, M]$ so that $[y, M]$ is a concave $D(A'')$-clip. Since $D(A'')$ is odd, y is a minimum and let x be the maximum before y. We can capture A' on $[m, x]$ and again changing t_r to $t_r = x - \varepsilon$, the resulting $t_1 < \cdots < t_n$ capture A on the concave $D(A)$-clip $[m, M]$. □

Corollary 3.13. *Let $A : [n] \to [n]$ be a non-crossing arborescence. Then*

$$\mu(A) \ = \ |\mathbb{L}(A)| + |\mathbb{L}^\circ(A)| + 1\,.$$

Proof. Let $A : [n] \to [n]$ be a non-crossing arborescence. By Theorem 3.11 it suffices to prove that A is captured by a polynomial P of degree $d = |\mathbb{L}(A)| + |\mathbb{L}^\circ(A)| + 1$ on some $\boldsymbol{t} = (t_1 < \cdots < t_n)$. If $n-1$ is not an immediate leaf of A, then $D(A) = |\mathbb{L}(A)| + |\mathbb{L}^\circ(A)| = d-1$ and $-T_d$ has a concave $D(A)$-clip. If $n-1 \in \mathbb{L}(A)$, then $D(A) = |\mathbb{L}(A)| + |\mathbb{L}^\circ(A)| + 1 = d$. We note that if $n-1$ is an immediate leaf, then in a $D(A)$-clip $[m, M]$ the maximum M does not have to be concave and we can choose $M = 1$ to capture A on T_d. □

3.2.2 Realization sets and universal arborescence

In this section, we now investigate the collection of vectors $\boldsymbol{t} = (t_1 < t_2 < \cdots < t_n)$ for which a non-crossing arborescence A can be captured on $\boldsymbol{t}$ by a polynomial P of degree at most d. We use these realization sets to characterize universal non-crossing arborescences A (i.e. $\mu(A, \boldsymbol{t}) = \mu(A)$ for all $\boldsymbol{t}$) of intrinsic degree $\mu(A) \le 3$. These universal arborescences will correspond to vertices of the cyclic associahedron $\mathsf{Asso}_{d-1}(\boldsymbol{t})$ for every $d \ge \max(\mu(A), 4)$.

Realization sets

Definition 3.14. For a non-crossing arborescence $A : [n] \to [n]$ and any d, we define the *realization set $\mathcal{T}_d^\circ(A)$* of A as the collection of $\boldsymbol{t} = (t_1 < t_2 < \cdots < t_n) \in \mathbb{R}^n$ such that A can be captured on $\boldsymbol{t}$ by some polynomial P of degree at most d.

If A is captured on $\boldsymbol{t}$ by P, then for $\lambda > 0$, A is captured on $\lambda\boldsymbol{t}$ by $P(\frac{t}{\lambda})$. Likewise, A is captured on $(c, \ldots, c) + \boldsymbol{t}$ by $P(t - c)$. With Corollary 3.10, this shows the closure $\mathcal{T}_d(A)$ of $\mathcal{T}_d^\circ(A)$ is a (generally non-convex) full-dimensional subcone of the *order cone* $\mathsf{O}_n = \{\boldsymbol{t} \in \mathbb{R}^n : t_1 \le \cdots \le t_n\}$.

In particular, when convenient, we can assume $t_1 = 0$ and $t_n = 1$. By definition, remark that

$$\mathcal{T}_1(A) \subseteq \mathcal{T}_2(A) \subseteq \cdots \subseteq \mathcal{T}_n(A) = \mathsf{O}_n.$$

In order to give a description of $\mathcal{T}_d(A)$, let $\mathcal{I}_A^b, \mathcal{I}_A^f \subseteq [n-2]$ be the sets of backward-sliding and forward-sliding nodes of A. We start with a description of the collection of polynomials P that capture A on a given $\boldsymbol{t}$. If $i \in \mathcal{I}_A^f$ is forward-sliding, we write $i^* := A(i)$. If $i \in \mathcal{I}_A^b$ is backward-sliding, we write i^* for minimal $j > i$ with $A(j) = A(i)$. See Figure 34.

Lemma 3.15. *Let $A : [n] \to [n]$ be a non-crossing arborescence and $\boldsymbol{t} \in \mathsf{O}_n^\circ$. A polynomial P captures A on $\boldsymbol{t}$ if and only if for all forward-sliding $i \in \mathcal{I}_A^f$*

$$(t_{A(i^*)} - t_i)(P(t_{i^*}) - P(t_i)) - (t_{i^*} - t_i)(P(t_{A(i^*)}) - P(t_i)) \ > \ 0$$

and for all backward-sliding $i \in \mathcal{I}_A^b$

$$(t_{A(i^*)} - t_i)(P(t_{i^*}) - P(t_i)) - (t_{i^*} - t_i)(P(t_{A(i^*)}) - P(t_i)) \ < \ 0\,,$$

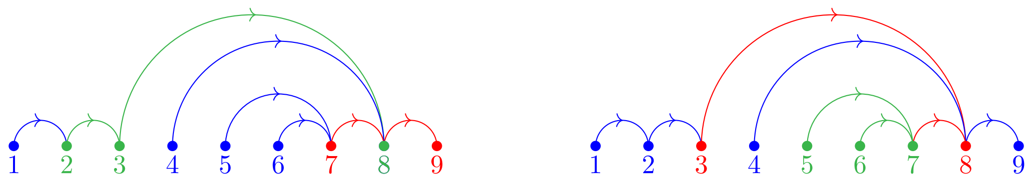

2 is forward-sliding.
7 is not forward-sliding.

5 is backward sliding.
3 is backward-sliding, but $3^* \neq 7$

Figure 34: Here, $\mathcal{I}_A^f = \{1, 2, 3, 5\}$ is the set of forward-sliding nodes, and $\mathcal{I}_A^b = \{3, 4, 5\}$ is the set of backward-sliding nodes.

Proof. Let $P(t) = w_d t^d + w_{d-1} t^{d-1} + \cdots + w_1 t = \langle \boldsymbol{w}, \boldsymbol{\gamma}_d(t) \rangle$. Then P captures A on $\boldsymbol{t}$ if and only if $\langle \boldsymbol{w}, \Psi(A) \rangle > \langle \boldsymbol{w}, \Psi(A') \rangle$ for all non-crossing arborescences $A' \neq A$. By convexity, it suffices to consider only those A' such that $[\Psi(A), \Psi(A')]$ is an edge. By Proposition 3.4 this boils down to

$$\langle \boldsymbol{w}, \Psi(A) \rangle - \langle \boldsymbol{w}, \Psi(F_i A) \rangle \;=\; \frac{P(t_{i^*}) - P(t_i)}{t_{i^*} - t_i} - \frac{P(t_{A(i^*)}) - P(t_i)}{t_{A(i^*)} - t_i} \;>\; 0\,,$$

when $i \in \mathcal{I}_A^f$; and when $i \in \mathcal{I}_A^b$ to

$$\langle \boldsymbol{w}, \Psi(A) \rangle - \langle \boldsymbol{w}, \Psi(B_i A) \rangle \;=\; \frac{P(t_{A(i)}) - P(t_i)}{t_{A(i)} - t_i} - \frac{P(t_{i^*}) - P(t_i)}{t_{i^*} - t_i} \;>\; 0$$

□

We can write Lemma 3.15 as follows. Define the *complete symmetric polynomial* of degree s:

$$h_s(x_1, x_2, x_3) \;:=\; \sum_{a+b+c=s} x_1^a x_2^b x_3^c$$

For $i \in \mathcal{I}_A^b \cup \mathcal{I}_A^f$, we construct $\Phi_i^d(\boldsymbol{t}) \in \mathbb{R}^d$ with $\Phi_i^d(\boldsymbol{t})_j = h_{j-2}\left(t_i, t_{i^*}, t_{A(i^*)}\right)$.

Observe that $\Phi_i^d(\boldsymbol{t})_1 = 0$ and $\Phi_i^d(\boldsymbol{t})_2 = 1$. We set $\overline{\Phi}_i^d(\boldsymbol{t}) := \left(h_j(t_i, t_{i^*}, t_{A(i^*)})\right)_{j=1,\ldots,d-2} \in \mathbb{R}^{d-2}$.

Theorem 3.16. *Let $A : [n] \to [n]$ be a non-crossing arborescence. For $\boldsymbol{t} \in \mathsf{O}_n^\circ$ and $d \geq 2$, define the polytopes*

$$\mathsf{P}_d^f(A, \boldsymbol{t}) \;:=\; \operatorname{conv}\left\{\overline{\Phi}_i^d(\boldsymbol{t}) : i \in \mathcal{I}_A^f\right\} \qquad \text{and} \qquad \mathsf{P}_d^b(A, \boldsymbol{t}) \;:=\; \operatorname{conv}\left\{\overline{\Phi}_i^d(\boldsymbol{t}) : i \in \mathcal{I}_A^b\right\}.$$

Then A is captured on $\boldsymbol{t}$ by some P of degree at most d if and only if $\mathsf{P}_d^f(A, \boldsymbol{t}) \cap \mathsf{P}_d^b(A, \boldsymbol{t}) = \varnothing$.

Proof. Let $P(t) = w_1 t + \cdots + w_d t^d = \langle \boldsymbol{w}, \boldsymbol{\gamma}_d(t) \rangle$, and $\boldsymbol{t} \in \mathsf{O}_n^\circ$. By Lemma 3.15, A is captured on $\boldsymbol{t}$ by some polynomial P of degree at most d if and only if $\langle \boldsymbol{w}, \Psi(A) - \Psi(F_i A) \rangle > 0$ for all $i \in \mathcal{I}_A^f$ and $\langle \boldsymbol{w}, \Psi(A) - \Psi(B_i A) \rangle > 0$ for all $i \in \mathcal{I}_A^b$. For $i \in \mathcal{I}_A^f$, let $j = A(i)$ and $k = A(j)$. We compute

$$\Psi(F_i A)_{r+1} - \Psi(A)_{r+1} \;=\; \frac{t_k^{r+1} - t_i^{r+1}}{t_k - t_i} - \frac{t_j^{r+1} - t_i^{r+1}}{t_j - t_i} \;=\; \sum_{s=0}^{r} t_k^s t_i^{r-s} - \sum_{s=0}^{r} t_j^s t_i^{r-s} \;=\; \sum_{s=0}^{r} (t_k^s - t_j^s) t_i^{r-s}\,.$$

This implies that $\Phi_i^d(\boldsymbol{t}) = \frac{1}{t_k - t_j}(\Psi(F_i A) - \Psi(A))$, and as $t_j < t_k$, the inequality $\langle \boldsymbol{w}, \Psi(A) - \Psi(F_i A) \rangle > 0$ is equivalent to $\left\langle \boldsymbol{w}, \Phi_i^d(\boldsymbol{t}) \right\rangle < 0$. For $i \in \mathcal{I}_A^b$ we can prove analogously, that $\langle \boldsymbol{w}, \Psi(A) - \Psi(B_i A) \rangle > 0$ is equivalent to $\left\langle \boldsymbol{w}, \Phi_i^d(\boldsymbol{t}) \right\rangle > 0$. This gives us a system of strict linear inequalities that by Gordan's lemma, a variant of Farkas' lemma (cf. [Sch98, Sect. 7.8]), has a solution if and only if there are $\lambda_i \geq 0$ for $i \in \mathcal{I}_A^f$ and $\mu_j \geq 0$ for $j \in \mathcal{I}_A^b$ not identically zero and

$$\sum_{i \in \mathcal{I}_A^f} \lambda_i \Phi_i^d(\boldsymbol{t}) \;=\; \sum_{i \in \mathcal{I}_A^b} \mu_i \Phi_i^d(\boldsymbol{t})\,.$$

Since $\Phi_i^d(\boldsymbol{t})_2 = 1$, it follows that $\Lambda := \sum_{i \in \mathcal{I}_A^f} \lambda_i = \sum_{i \in \mathcal{I}_A^b} \mu_i > 0$. Dividing both sides of the above equality by Λ yields a point in $\mathsf{P}_d^f(A, \boldsymbol{t}) \cap \mathsf{P}_d^b(A, \boldsymbol{t})$. □

Example 3.17. For $n = 5$, one can consider $\boldsymbol{t} = (1, 2, 3, 4, 5)$. In Figure 35 are drawn $\mathsf{P}^b_d(A, \boldsymbol{t})$ and $\mathsf{P}^f_d(A, \boldsymbol{t})$ for $d = 3$ and $d = 4$ for the two given non-crossing arborescences. None of the polytopes intersects for $d = 4$, indicating that both arborescences are captured in degree 4 (what was already known, as $d = 4 = n - 1$, and all non-crossing arborescences on n nodes are captured in degree d). However, for $d = 3$, the left arborescence is not captured while the right one is.

The situation is inverted for $\boldsymbol{t} = (-1, 2, 3, 4, 5)$, as shown in Figure 36.

Non-crossing arborescences with $\mu(A) \leq 3$ In this section, we will characterize realization sets of non-crossing arborescences of intrinsic degree at most 3. We call a non-crossing arborescence A *quadratic* if $\mu(A) = 2$, and *cubic* if $\mu(A) = 3$. We start with a classification of quadratic and cubic arborescences, obtained from Corollary 3.13. Quadratic non-crossing arborescences have exactly 1 exterior immediate leave (and no interior one), while cubic ones have either 1 interior and no exterior, or 2 exterior ones. Consequently, their non-crossing property gives:

Corollary 3.18. *The two non-crossing arborescences of intrinsic degree 2 on $n \geq 3$ nodes are A_m and A_M, defined by $A_m(i) = i + 1$ for $1 \leq i \leq n - 1$ and $A_M(i) = n$ for $1 \leq i \leq n - 1$. See Figure 37.*

Corollary 3.19. *For $n \geq 4$, there are precisely $2^{n-2} + n - 5$ non-crossing arborescences A on n nodes with $\mu(A) = 3$. They come in two types:*

(i) *For $1 < k < n - 1$, define $A(i) = i + 1$ for $1 \leq i < k$ and $A(i) = n$ for $k \leq i < n$. These are $n - 3$ non-crossing arborescences with $\mathbb{L}(A) = \{1, n - 1\}$.*

(ii) *For $1 < k < n - 1$ and $n \geq j_1 \geq j_2 \geq \cdots \geq j_{k-1} > k$, define $A(i) = j_i$ for $1 \leq i < k$ and $A(i) = i + 1$ for $k \leq i < n$. These are $2^{n-2} - 2$ non-crossing arborescences with $\mathbb{L}(A) = \{k\}$.*

We call a non-crossing arborescence $A : [n] \to [n]$ *universal* if A can be captured by a polynomial of degree $\mu(A)$ on all $\boldsymbol{t} \in \mathsf{O}^\circ_n$, that is: $\mathcal{T}_{\mu(A)}(A) = \mathsf{O}_n$. Note that this implies that $\mathcal{T}_d(A) = \mathsf{O}_n$ for all $d \geq \mu(A)$. We next determine all universal arborescences A with $\mu(A) \leq 3$, giving a description of the realization set of any non-crossing arborescence A with $\mu(A) \leq 3$.

Lemma 3.20. *Let $A : [n] \to [n]$ be a non-crossing arborescence with $\mu(A) = 3$ and $\mathbb{L}(A) = \{k\}$, $1 < k < n - 1$. Then for all $\boldsymbol{t} \in \mathsf{O}^\circ_n$,*

$$\min\{t_i + t_{i^*} + t_{A(i^*)} : i \in \mathcal{I}^b_A\} < \max\{t_i + t_{i^*} + t_{A(i^*)} : i \in \mathcal{I}^f_A\}\,.$$

Proof. As $k \neq n - 1$ there is an $i \in \mathcal{I}^f_A \cap [n - 2]$ with $i^* = n - 1$. Recall that A is of the form Corollary 3.19(ii). If $j_1 \leq n - 2$, then for all $j \in \mathcal{I}^b_A$ we have $A(j^*) \leq n - 2$ and the statement follows. Otherwise, either $j \in \mathcal{I}^b_A$ exists with $j^* = n - 1$ and $j < i$, or $j^* = n - 2$ and $i \leq j$. In both cases, for all $\boldsymbol{t} \in \mathsf{O}^\circ_n$, we have $\min\{t_i + t_{i^*} + t_{A(i^*)} : i \in \mathcal{I}^b_A\} \leq t_j + t_{j^*} + t_{A(j^*)} < t_i + t_{i^*} + t_{A(i^*)} \leq \max\{t_i + t_{i^*} + t_{A(i^*)} : i \in \mathcal{I}^f_A\}$. □

Theorem 3.21. *Let $A : [n] \to [n]$ be a non-crossing arborescence with $\mu(A) \leq 3$. Then A is universal if and only if $\mu(A) = 2$, or if $\mu(A) = 3$ and*

(a) *$\mathbb{L}(A) = \{1, n - 1\}$, or*

(b) *$\mathbb{L}(A) = \{n - 2\}$, $A(i) = n$ for $i = 1, \ldots, n - 4$, and $A(n - 3) \in \{n - 1, n\}$, or*

(c) *$\mathbb{L}(A) = \{2\}$ and $A(1) \in \{3, 4\}$.*

Proof. Using Theorem 3.16 it suffices to show that $\mathsf{P}^f_d(A, \boldsymbol{t}) \cap \mathsf{P}^b_d(A, \boldsymbol{t}) = \varnothing$ for all choices of $\boldsymbol{t}$ holds in precisely the situations stipulated above.

When $\mu(A) = 2$, $\mathsf{P}^f_d(A, \boldsymbol{t})$ and $\mathsf{P}^b_d(A, \boldsymbol{t})$ are polytopes in $\mathbb{R}^0$. Hence, the claim holds if and only if $\mathcal{I}^b_A = \varnothing$ or $\mathcal{I}^f_A = \varnothing$. By Corollary 3.18 this is the case for both A_m and A_M.

If $\mu(A) = 3$, then $\mathsf{P}^f_d(A, \boldsymbol{t}) = [x_f, y_f]$, $\mathsf{P}^b_d(A, \boldsymbol{t}) = [x_b, y_b] \subset \mathbb{R}$ with (x_f, y_f) the minimum and maximum of $\{t_i + t_{i^*} + t_{A(i^*)} : i \in \mathcal{I}^f_A\}$ and likewise for (x_b, y_b).

Let A an arborescence with $\mu(A) = 3$. If $\mathbb{L}(A) = \{1, n - 1\}$, then A satisfies Corollary 3.19 (i) for some $1 < k < n - 1$. Thus $\mathcal{I}^f_A = \{1, 2, \ldots, k - 1\}$ and $\mathcal{I}^b_A = \{k, \ldots, n - 2\}$. It follows that

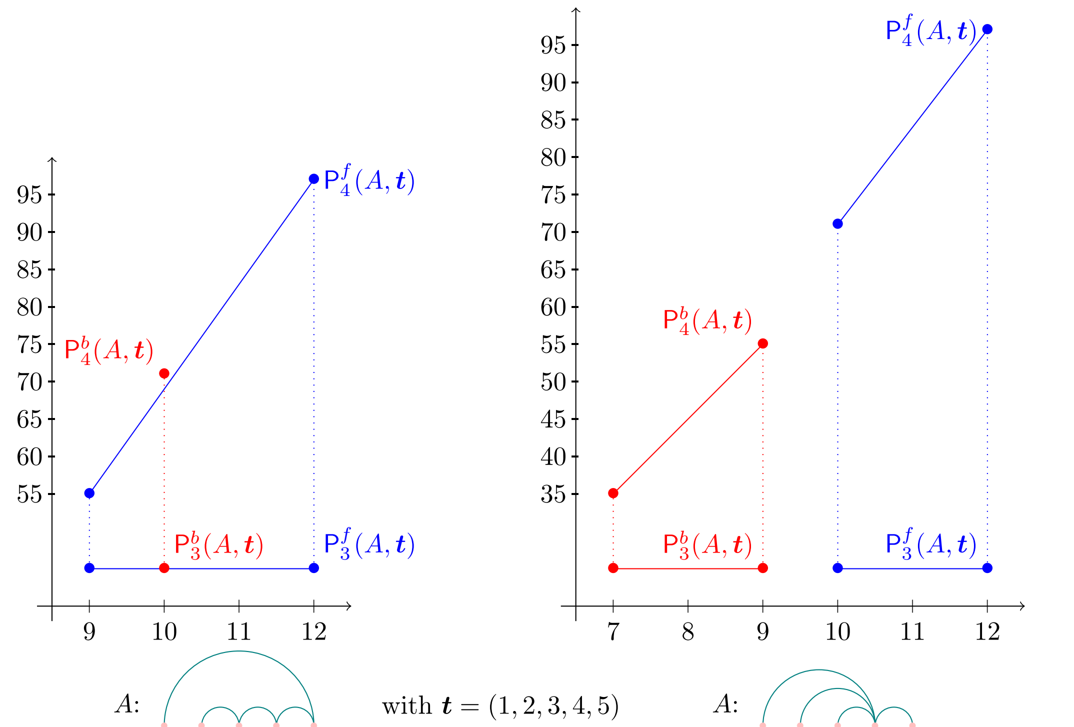

Figure 35: The polytopes $\mathsf{P}_d^b(A, \boldsymbol{t})$ and $\mathsf{P}_d^f(A, \boldsymbol{t})$ with $\boldsymbol{t} = (1, 2, 3, 4, 5)$, for $d = 3$ (Bottom, 1-dimensional drawing) and $d = 4$ (Top, 2-dimensional drawing), for the two non-crossing arborescences drawn (Left and Right).

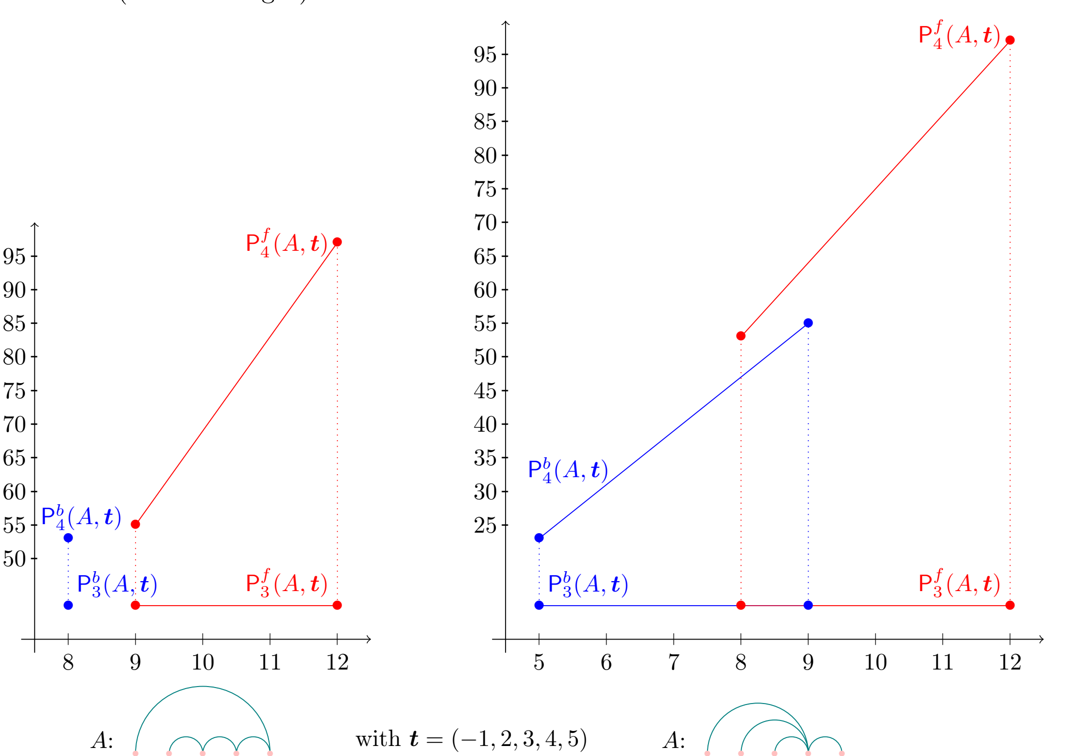

Figure 36: The polytopes $\mathsf{P}_d^b(A, \boldsymbol{t})$ and $\mathsf{P}_d^f(A, \boldsymbol{t})$ with $\boldsymbol{t} = (-1, 2, 3, 4, 5)$, for $d = 3$ (Bottom) and $d = 4$ (Top), for the two non-crossing arborescences drawn (Left and Right).

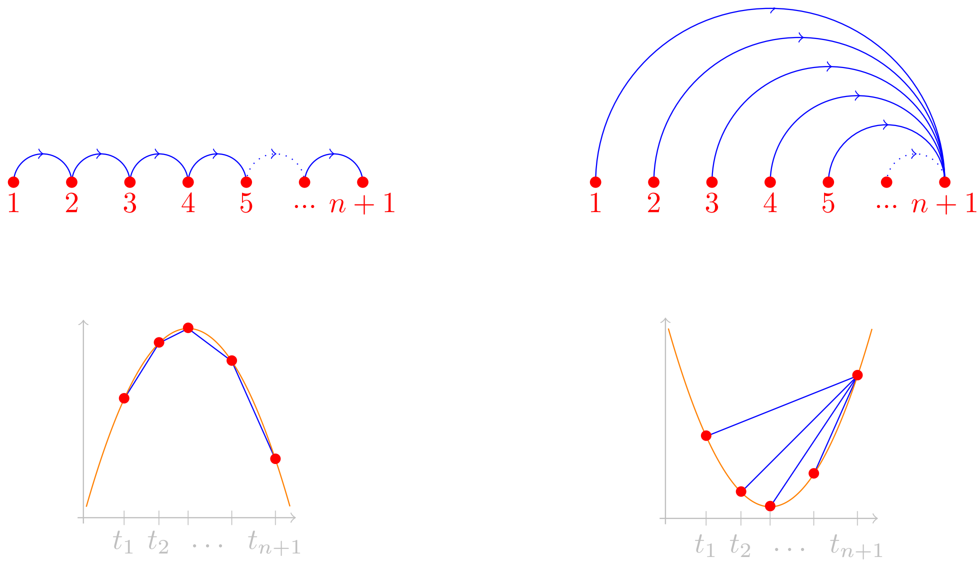

Figure 37: The two (universal) non-crossing arborescences of intrinsic degree 2, and degree 2 polynomials capturing them.

$y_f = t_{k-1} + t_k + t_{n+1} < t_k + t_{k+1} + t_n = x_b$. This proves (a). Otherwise, A satisfies Corollary 3.19 (ii) for some $1 < k < n-1$ and $n \geq j_1 \geq j_2 \geq \cdots \geq j_{k-1} > k$. Note that $\mathcal{I}_A^b \subseteq [k-1]$ and they are all leaves. For (b), $k = n-2$ and $j_1 = \cdots = j_{n-4} = n$. If $j_{n-3} = n$, then $n-2$ is the only forward-sliding node and $x_f = y_f = t_{n-2} + t_{n-1} + t_n > t_{n-3} + t_{n-1} + t_n = y_b$. If $j_{n-3} = n-1$, then $x_f = y_f = t_{n-3} + t_{n-1} + t_n > \max(t_{n-3} + t_{n-1} + t_{n-2},\ t_{n-4} + t_{n-1} + t_n) = y_f$. Likewise, for (c) we have $k = 2$. If $A(1) = 3$, then $x_b = y_b = t_1 + t_2 + t_3 < t_1 + t_3 + t_4 = x_f$. If $A(1) = 4$, then $x_b = y_b = t_1 + t_3 + t_4 < \min(t_2 + t_3 + t_4, t_1 + t_4 + t_5) = x_f$.

Assume A satisfies Corollary 3.19 (ii) with $\mathbb{L}(A) = \{k\}$, but neither (b) nor (c). In this case, we will find $\boldsymbol{t}$ with $\mathsf{P}_d^f(A, \boldsymbol{t}) \cap \mathsf{P}_d^b(A, \boldsymbol{t}) \neq \varnothing$, which proves A is not universal. By Lemma 3.20, $\mathsf{P}_d^f(A, \boldsymbol{t}) \cap \mathsf{P}_d^b(A, \boldsymbol{t}) = \varnothing$ if and only if $y_b < x_f$.

If k is forward-sliding, then $k^* = k+1$ and either there is $i < k$ with $j_i > A(k^*) = k+2$, or $j_1 = \cdots = j_{k-1} = k+2 < n$ and $1 < k-1$. In the first case, choose the maximal $i < k$ with $j_i > k+2$: i is backward-sliding and $i^* = A(i^*) - 1 = j_i - 1 > k+1$. For any $\boldsymbol{t} \in \mathsf{O}_n^\circ$ with $t_{i^*} > t_k + t_{k+1}$ it follows $x_f \leq t_k + t_{k+1} + t_{k+2} < t_i + t_{i^*} + t_{A(i^*)} \leq y_b$. In the second case, $i = 1$ is forward-sliding with $i^* = k+2$, and $k-1$ is backward-sliding with $(k-1)^* = k+1$. Choosing $\boldsymbol{t} \in \mathsf{O}_n^\circ$ with t_1 small enough, ensures $x_f \leq t_1 + t_{1^*} + t_{A(1^*)} < t_{k-1} + t_{k-1^*} + t_{A(k-1^*)} \leq y_b$.

If k is not forward-sliding, then $k-1$ is backward-sliding with $(k-1)^* = k$ and there is a minimal forward-sliding $i < k$ with $i^* = k+1$. If $i < k-1$ then we can find $\boldsymbol{t} \in \mathsf{O}_n^\circ$ with t_i small enough so that $x_f \leq t_i + t_{k+1} + t_{k+2} < t_{k-1} + t_k + t_{k+1} \leq y_b$. If $i = k-1$, then $1 < k-1$ and we can choose $j < k-1$ forward-sliding. Similarly, we can find $\boldsymbol{t} \in \mathsf{O}_n^\circ$, with t_j small enough and thus $x_f \leq t_j + t_{j^*} + t_{A(j^*)} < t_{k-1} + t_k + t_{k+1} \leq y_b$. ☐

Corollary 3.22. *The number of universal non-crossing arborescences $A : [n] \to [n]$ with $\mu(A) \leq 3$ is $n+3$.*

Proof. The universal non-crossing arborescences on $[n]$ are characterized by Theorem 3.21. There are exactly two non-crossing arborescences with $\mu(A) = 2$. Moreover, there are $n-3$ arborescences of the form (a), 2 of the form (b) and 2 of the form (c). ☐

Example 3.23. In Figure 38(Left), you can see all non-crossing arborescences $A : [5] \to [5]$ with $\mu(A) \leq 3$. Two arborescences A, A' are connected by an edge if and only if they differ by a flip.

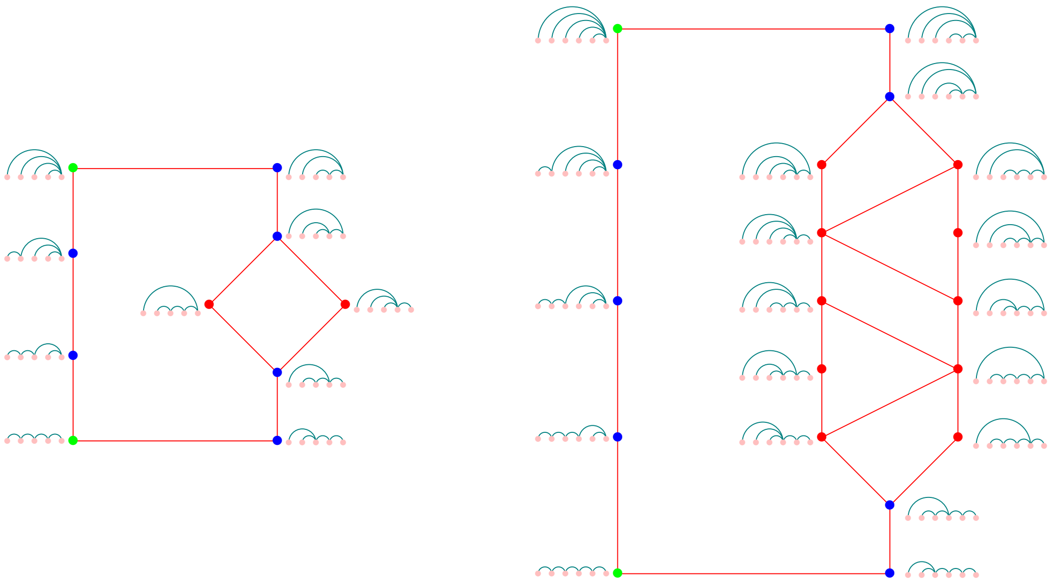

Figure 38: All non-crossing arborescences A on $n = 5$ (Left) and $n = 6$ (Right) nodes with $\mu(A) \leq 3$. Green and blue dots represent universal arborescences, red dots non-universal ones.

The arborescences at the green dots are the two universal arborescences of intrinsic degree 2, the ones at the blue dots are the universal arborescences of intrinsic degree 3. The two remaining arborescences at the red dots are the two non-universal arborescences of intrinsic degree 3, see Example 3.17.

One can construct the same graph for $n = 6$, see Figure 38(Right), and $n = 7$, see Figure 39. Note that these graphs are oriented from top to bottom by the Tamari orientation of flips. For $d \geq 4$ and fixed $\boldsymbol{t}$, one can consider the graph $G_{\boldsymbol{t}}$ whose set of vertices is the set of A with $\mu(A, \boldsymbol{t}) \leq 3$, and the edges are the flips between them. This forms a sub-graph of the graph of $\mathsf{Asso}_{d-1}(\boldsymbol{t})$, and hence a sub-graph of the graph of the associahedron. The projection principle of Proposition 3.9 ensures that $G_{\boldsymbol{t}}$ is the graph of a polygon, thus $G_{\boldsymbol{t}}$ is a cycle. The idea behind the rest of this section will be to prove that $G_{\boldsymbol{t}}$ is a *great cycle* in Figures 38 and 39 (and of the corresponding graph for greater n), meaning that $G_{\boldsymbol{t}}$ is composed by two paths from A_M to A_m: the left path of universal arborescences, and a right path that is increasing for the Tamari orientation. Not all increasing right paths will correspond to a $G_{\boldsymbol{t}}$ for some $\boldsymbol{t}$, but this will allow us to prove that the number of vertices of $G_{\boldsymbol{t}}$ is independent from $\boldsymbol{t}$.

Remark 3.24. Applying the bijection between non-crossing arborescences and triangulations to Figure 38 we obtain the graph of the fiber polytope for the canonical projection from $\mathsf{Cyc}_4(\boldsymbol{t})$ to $\mathsf{Cyc}_2(\boldsymbol{t})$ for $\boldsymbol{t} \in \mathsf{O}_n^\circ$ as pictured in Figure [ALRS00, Figure 1]. The study of this phenomenon will be at the heart of Section 4.3.

In the remaining of the section, we are going to use the properties of cubic arborescences and the first property we have given of their realization set in order to count the number of cubic arborescences that can be captured for a given $\boldsymbol{t} \in \mathsf{O}_n^\circ$. Even though it will be notationally heavy, most of it will boil down to proving that what our previous drawings indicates holds in general.

Definition 3.25. For a non-crossing arborescence A, a forward-sliding i is called *minimal* when i is a leaf, and a backward-sliding i is called *maximal* when $i^* + 1 = A(i^*)$.

Lemma 3.26. *Let $A : [n] \to [n]$ be a non-universal cubic arborescence, and $\mathbb{L}^\circ(A) = \{k\}$. For any minimal forward-sliding i, $i^* + 1 = A(i^*)$ with $i \leq k$ and $i^* > k$. Any maximal backward-sliding j is a leaf, $j < k$ and $j^* \geq k$.*

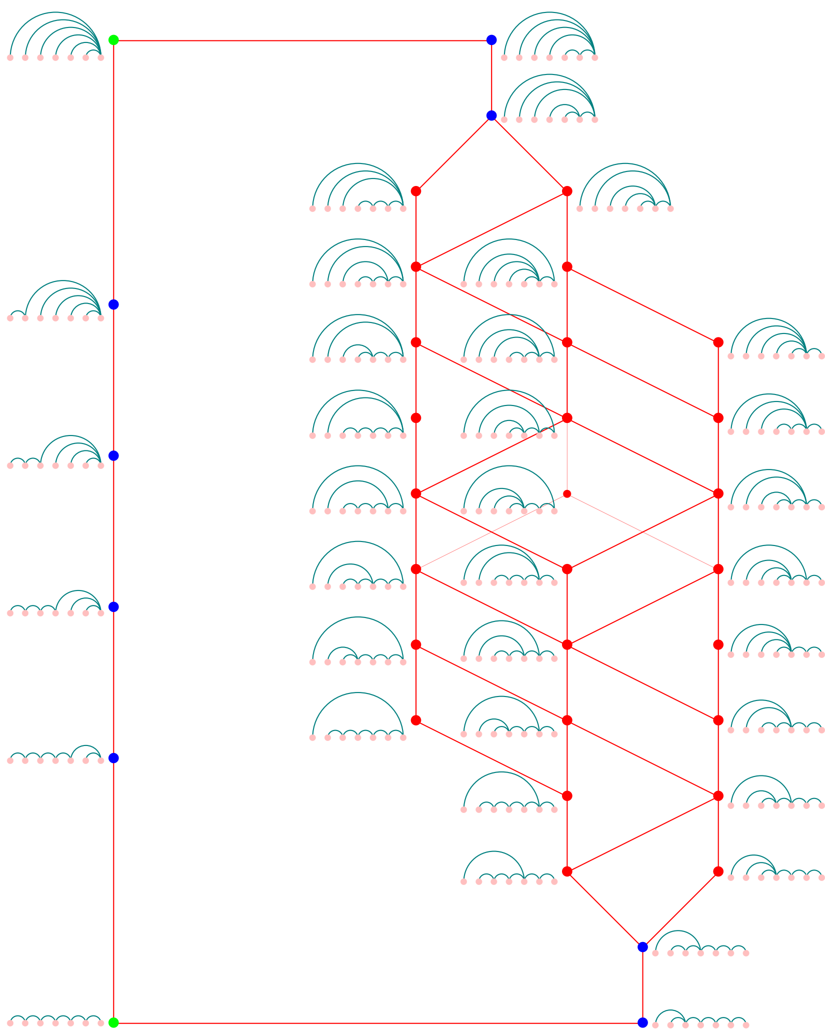

Figure 39: All non-crossing arborescences A on $n = 7$ nodes with $\mu(A) \leq 3$. Green and blue dots represent universal arborescences, red dots non-universal ones.

Proof. Since A is non-universal and cubic, A is of the form described in Corollary 3.19(ii). For any $i > k$, one has $A(i-1) = i$; thus any minimal forward-sliding i has to be smaller or equal to k. Moreover this ensures that $i^* > k$ and $A(i^*) = i^* + 1$.

Suppose j is backward-sliding and maximal. The fact that $A(j) > j+1$, forces $j < k$. Thus j is a leaf. By Corollary 3.19(ii), $j^* > k$. □

Theorem 3.27. *Let $A : [n] \to [n]$ be a cubic arborescence and $\boldsymbol{t} \in \mathsf{O}_n^\circ$. If A is non-universal, then $\boldsymbol{t} \in \mathcal{T}_3^\circ(A)$ if and only if*

(i) $t_i < t_{i+1}$ for all $1 \leq i < n$,

(ii) $t_j + t_{j^*} + t_{A(j^*)} < t_i + t_{i^*} + t_{A(i^*)}$ for all $j \in \mathcal{I}_A^b$ maximal and $i \in \mathcal{I}_A^f$ minimal.

Proof. Let A be cubic non-universal. By Theorem 3.16, it suffices to show that $\mathsf{P}_d^f(A, \boldsymbol{t}) \cap \mathsf{P}_d^b(A, \boldsymbol{t}) = \varnothing$ and $\boldsymbol{t} \in \mathsf{O}_n^\circ$ precisely when $\boldsymbol{t}$ satisfies the conditions (i) and (ii). Let $\mathsf{P}_d^f(A, \boldsymbol{t}) = [x_f, y_f]$ and $\mathsf{P}_d^b(A, \boldsymbol{t}) = [x_b, y_b]$. By Lemma 3.20 we conclude that for $\boldsymbol{t} \in \mathsf{O}_n^\circ$, $\mathsf{P}_d^f(A, \boldsymbol{t}) \cap \mathsf{P}_d^b(A, \boldsymbol{t}) = \varnothing$ if and only if $y_b < x_f$. As a consequence any $\boldsymbol{t} \in \mathcal{T}_3^\circ(A)$ satisfies (i) and (ii).

Conversely, let $\boldsymbol{t} \in \mathbb{R}^n$ fulfil (i) and (ii). Obviously $t \in \mathsf{O}_n^\circ$.

Assume that $i \in \mathcal{I}_A^f$ is not minimal. Then there is a forward-sliding leaf $a < i$ and $m > 1$ such that $A^m(a) = i$. This shows $x_f \leq t_a + t_{a^*} + t_{A(a^*)} < t_i + t_{i^*} + t_{A(i^*)}$. Hence $t_j + t_{j^*} + t_{A(j^*)} < t_i + t_{i^*} + t_{A(i^*)}$ is implied by $t_j + t_{j^*} + t_{A(j^*)} < t_a + t_{a^*} + t_{A(a^*)}$.

Assume that $j \in \mathcal{I}_A^b$ is not maximal. Then $j^* < A(j^*) - 1$ implies that there exists a backward-sliding b with $j^* \leq b < A(j^*)$ and $A(j^*) = A(b^*) = b^* + 1$. This shows $t_i + t_{i^*} + t_{A(i^*)} < t_j + t_{j^*} + t_{A(j^*)} \leq y_b$. We conclude that $t_j + t_{j^*} + t_{A(j^*)} < t_i + t_{i^*} + t_{A(i^*)}$ is implied by $t_b + t_{b^*} + t_{A(b^*)} < t_i + t_{i^*} + t_{A(i^*)}$.

□

Example 3.28. For $n = 5$, there are 10 non-crossing arborescences A with $\mu(A) \leq 3$, see Figure 38(Left). Among them, 8 are universal and 2 are not. We have discussed the 2 non-universal ones in Example 3.17. Theorem 3.27 allows us to compute the realizations sets of theses 2 non-universal arborescences: besides facets of O_5, they have a facet on the hyperplane $\{\boldsymbol{t} \in \mathbb{R}^5 \ ; \ t_2 + t_3 + t_4 = t_1 + t_4 + t_5\}$. In Figure 40(Bottom) are drawn these two realization sets, embedded inside the order cone O_5. As this cone is 5-dimensional, we intersect it with the two hyperplanes $\{\boldsymbol{t} \in \mathbb{R}^5 \ ; \ t_1 = 0\}$ and $\{\boldsymbol{t} \in \mathbb{R}^5 \ ; \ t_5 = 1\}$, making the picture 3-dimensional. For each realization sets, we have drawn in blue $\mathsf{Asso}_{d-1}(\boldsymbol{t})$ for $d = 4$ which is an associahedron, and highlighted in red the non-crossing arborescences A with $\mu(A, \boldsymbol{t}) \leq 3$ for the corresponding $\boldsymbol{t}$.

Remark 3.29. Note that different cases of Theorem 3.27 lead to the same inequality (see Figure 41):

(i) For example, if $i \in \mathcal{I}_A^b \cap [1, k-1]$, $i+1 \in \mathcal{I}_A^f$ and $A(i+1) < i^*$. Since $i+1 \leq k$, $i+1$ is a leaf and thus minimal. By $A(i+1) < i^*$ and thus $i^* > k$ follows that i is maximal. Moreover, $i < i+1 < k < A(i+1) < A(A(i+1)) \leq i^* < A(i^*)$. By (ii), $t_{A(i+1)} < t_{A(A(i+1))} \leq t_{i^*} < t_{A(i^*)}$ and so (iii) implies $t_i < t_{i+1}$.

(ii) If $i \in \mathcal{I}_A^b$, $j \in \mathcal{I}_A^f$ and $i = j$, then $A(i^*) = j^*$ and thus the inequality $t_i + t_{i^*} + t_{A(i^*)} < t_j + t_{j^*} + t_{A(j^*)}$ is implied by $t_{i^*} < t_{A(j^*)}$.

(iii) If $i \in \mathcal{I}_A^b$, $j \in \mathcal{I}_A^f$ and $i^* = j^*$, then $j < i$ as A is non-crossing. As i is forward-sliding and j is maximal, $i = j+1$ and the inequality corresponding to i and j follows from $t_j < t_i$.

Now we want a closer look at the realization set and describe, which of the inequalities in Theorem 3.27 of the form (ii) give a facet for $\mathcal{T}_3(A)$.

Let $A : [n] \to [n]$ be a non-crossing arborescence, $i \in [n-2]$ forward-sliding and $j \in [n-2]$ backward-sliding. Recall the definitions of flips (F_iA and B_jA are non-crossing arborescences):

$$F_iA(k) = \begin{cases} A(i^*) & k = i \\ A(k) & k \neq i \end{cases} \qquad \text{and} \qquad B_jA(k) = \begin{cases} j^* & k = j \\ A(k) & k \neq i. \end{cases}$$

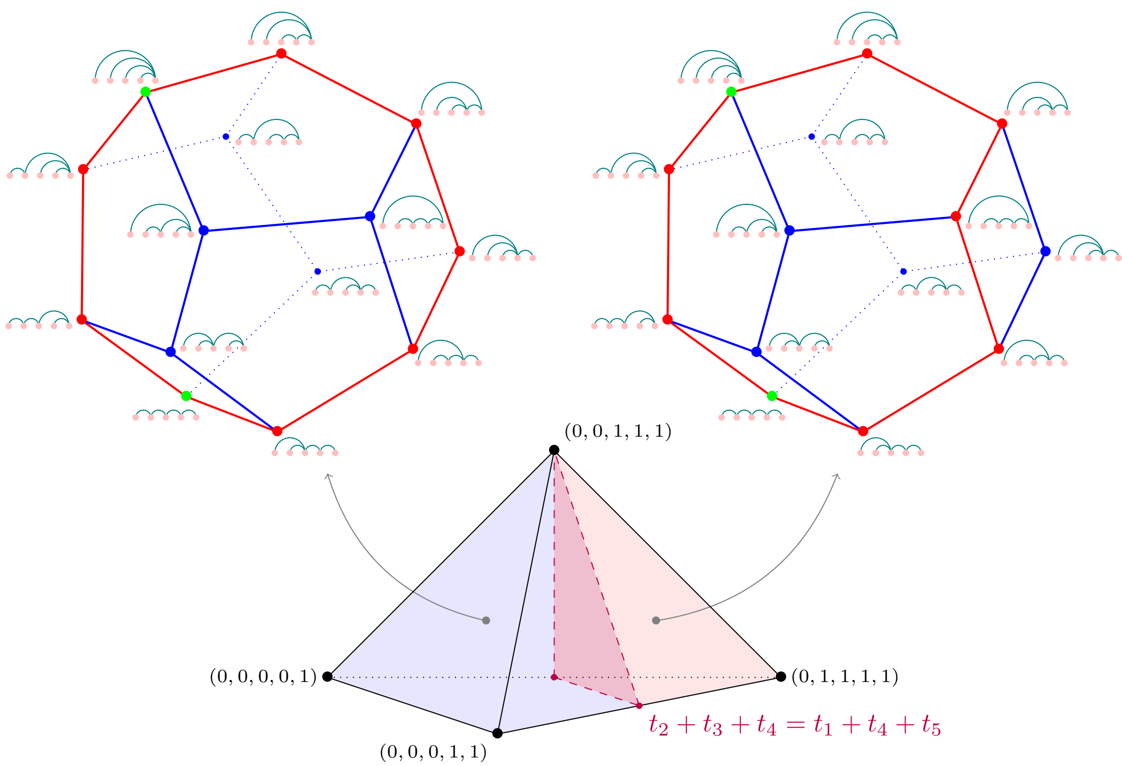

Figure 40: The order cone O_5 intersected by the hyperplanes $\{\boldsymbol{t}\ ;\ t_1 = 0\}$ and $\{\boldsymbol{t}\ ;\ t_5 = 1\}$, and subdivided into the different realization sets for non-crossing arborescences A with $\mu(A) \leq 3$. For each realization sets, the cyclic associahedron $\mathsf{Asso}_3(\boldsymbol{t})$ with, highlighted in red, the non-crossing arborescences A with $\mu(A, \boldsymbol{t}) \leq 3$.

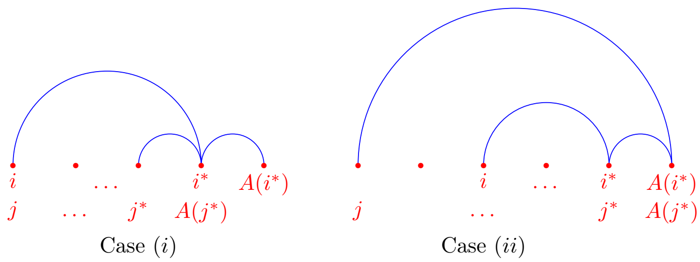

Figure 41: Some diagonal switches are forbidden due to the relative positions of i, j, i^* and j^*.

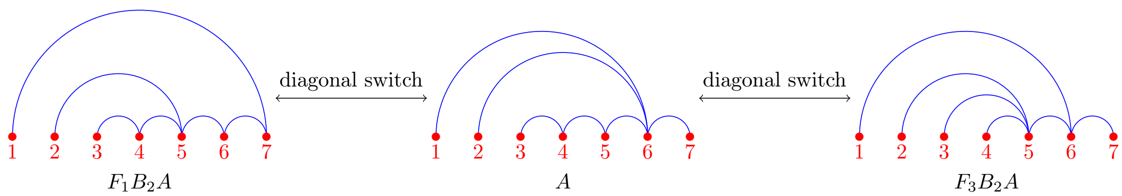

Figure 42: A non-crossing arborescence $A : [7] \to [7]$ and two diagonal switches.

If j is backward-sliding in F_iA and i is forward-sliding in B_jA then we could consider the combination of the two flips. We say that A and A' differ by a *diagonal switch* if $A' = F_iB_jA = B_jF_iA \neq A$, equivalently we say that we perform a diagonal switch with respect to i and j on A.

The notion *diagonal* switch is motivated by the fact that A, F_iA, B_jA and F_iB_jA form a square face[8] of Asso_{n-2}, and switching from A to F_iB_jA corresponds to switching along the diagonal of the square that contains A. In the directed graph $G_{\boldsymbol{t}}$ defined in Example 3.23, a diagonal flip amounts to travelling through two edges: one respecting the Tamari orientation and the other not (in Figure 38 and Figure 39, two arborescences are linked by a diagonal switch when there are at the same height and at distance 2).

Lemma 3.30. *Let $A : [n] \to [n]$ be a non-crossing arborescence, i forward-sliding and j backward-sliding. We can perform a diagonal switch with respect to i and j on A, if and only if the following conditions are fulfilled, see Figure 41:*

(i) $i \neq j$
(ii) $j^* \neq i^*$
(iii) $j \neq i^*$

Proof. Suppose one of the conditions is not fulfilled. If $i = j$, then $F_iB_jA = B_jF_iA = A$. If $j^* = i^*$, then $A(j^*) = A(i^*)$ and as $A(j) = A(j^*) > A(i) = i^*$, we need $j < i$. Hence $B_jA(j) = j^* = i^*$, and thus i is not forward-sliding in B_jA. If $j = i^*$, then $F_iB_jA(i) = j^* \neq B_jF_iA(i) = A(i^*)$.

Suppose $(i) - (iii)$ are fulfilled. Notice that (i) implies $i^* \neq A(j^*)$, (ii) and (iii) imply $A(i^*) \neq A(j^*)$. It cannot happen that $j^* = i$, as then i would not be forward-sliding. If $i = A(j^*), j = A(i^*)$ or $j^* = A(i^*)$, then the diagonal switch can be performed. Otherwise, $\{i, i^*, A(i^*)\} \cap \{j, j^*, A(j^*)\} = \varnothing$ and thus the diagonal switch can be performed as well. $\square$

Corollary 3.31. *Let $A : [n] \to [n]$ be a non-crossing arborescence, $1 \leq i \leq n-2$ forward-sliding and $1 \leq j \leq n-2$ backward-sliding. If $j = i^*$, then either $\mathbb{L}(A) = \{1, n-1\}$ or $\mu(A) > 3$.*

Proof. From $j = i^*$ follows that $A(i^*) = A(j^*)$. As $i^* < A(i^*) - 1$, there is an immediate leaf $i^* < \ell_1 < A(i^*) \leq n-1$. Furthermore either $A(1) = 2$ and thus $1 \in \mathbb{L}(A)$, or there is an interior immediate leaf $1 < \ell_2 < i^*$. $\square$

[8] All induced 3-, 4- and 5-cycles in the graph of a simple polytope are 2-faces, thus the four vertices associated to A, F_iA, B_jA and F_iB_jA form a square **face** of Asso_{n-2}.

Proposition 3.32. *Let $A : [n] \to [n]$ be a cubic non-universal non-crossing arborescence. Let $1 \leq i \leq n-2$ be forward-sliding and minimal and $1 \leq j \leq n-2$ backward-sliding and maximal. If $i \neq j$ and $i^* \neq j^*$, then $A' = F_i B_j A = B_j F_i A$ is cubic and non-universal.*

Proof. As A is non-universal and cubic, by Theorem 3.21 A satisfies $\mathbb{L}(A) = \{\ell\}$, $1 < \ell < n-1$. If i is minimal, then, by Corollary 3.19 and Lemma 3.26, either $i = \ell$ or $i < \ell$. If $i = \ell$, then this implies $i^* = \ell + 1$ and $A(i^*) = \ell + 2$ and thus $\mathbb{L}(F_i A) = \mathbb{L}(F_l A) = \{\ell + 1\}$. If $i < \ell$, then as $A(i^*) = i^* + 1 \geq \ell + 1$, it follows that $\mathbb{L}(F_i A) = \mathbb{L}(A) = \{\ell\}$.

If j is maximal, then, by Corollary 3.19, either $j = \ell - 1$ or $j < \ell < j^*$. In the first case, $j^* = \ell$, $A(j^*) = \ell + 1$ and $\mathbb{L}(B_j A) = \{\ell - 1\}$. In the second case, $\mathbb{L}(B_j A) = \mathbb{L}(A) = \{\ell\}$.

Moreover, we observe, that ℓ can not be forward-sliding and $\ell - 1$ is backward-sliding at the same time. We conclude, that $\mathbb{L}(F_i B_j A) \subseteq [2, n-2]$ and thus $F_i B_j A$ is cubic. If $3 \leq \ell \leq n-3$ this almost shows that $F_i B_j A$ is non-universal. We only have to consider two special cases. Suppose $\ell = 3$ and $A(1) = 4$. Then only 2 is backward-sliding and maximal and only 1 is forward-sliding and minimal. Thus $\mathbb{L}(B_2 F_1 A) = \{2\}$, but with $A(1) = 5$, implying that $B_2 F_1 A$ is non-universal. Similarly, if $\ell = n-3$ and $A(k) = n$ for all $k < n-3$, then as only $n-3$ is forward-sliding and minimal and only $n-4$ is backward-sliding and maximal. We have $\mathbb{L}(B_{n-4} F_{n-3} A) = \{n-2\}$. Moreover, $B_{n-4} F_{n-3} A(n-4) = n-1$ and thus $B_{n-4} F_{n-3} A$ is non-universal.

If $\ell = 2$, as A is non-universal, we have $A(1) > 4$. Thus only 1 is backward-sliding and maximal, and only 2 is forward-sliding and minimal, which leads to $\mathbb{L}(B_1 F_2 A) = 3$ and hence $B_1 F_2 A$ is non-universal. If $\ell = n-2$, as A is non-universal, we have $A(n-4) = A(n-3) = n-1$. Then the smallest $i < n-4$ such that $A(i) = n-1$ is the only forward-sliding and minimal choice, and only $n-3$ is backward-sliding and maximal. This leads to $\mathbb{L}(F_i B_{n-3} A) = \{n-3\}$ and thus A is non-universal. □

Now we want to highlight some inequalities of Theorem 3.27. Let $A : [n] \to [n]$ be a cubic arborescence. We call a facet of $\mathcal{T}_3(A)$ *internal*, if it is not contained in any facet of O_n. By definition, universal arborescences have no internal facets.

Theorem 3.33. *Let A be a cubic, non-crossing and non-universal arborescence. Any internal facet of $\mathcal{T}_3(A)$ is of the form*

$$\mathcal{T}_3(A) \cap \{\boldsymbol{t} \in \mathbb{R}^n \ ; \ t_j + t_{j^*} + t_{A(j^*)} = t_i + t_{i^*} + t_{A(i^*)}\},$$

where i is forward-sliding and minimal and j is backward-sliding and maximal, and a diagonal switch can be performed on A with respect to i and j.

Proof. In the inequality description given by Theorem 3.27, the inequalities of the form (i) come from facets of O_n. By Lemma 3.30 and Corollary 3.31, a diagonal switch can be performed with respect to i and j if and only if $i \neq j$ and $i^* \neq j^*$. If $i = j$ or $i^* = j^*$, then the associated inequality gives rise to a facet of O_n, Remark 3.29(ii) and (iii). □

Remark 3.34. The above Theorem 3.33 almost gives a facet-description of $\mathcal{T}_3(A)$. Indeed, we know that its facets are associated to pairs of compatible **minimal** forward-sliding and **maximal** backward-sliding nodes, but it is not mandatory that all such couples are associated to a facet. Nonetheless, the collection of such couples is far smaller than the collection of all couples of forward-sliding and backward-sliding nodes, and will reveal to be far more manageable thanks to the diagonal flip.

Definition 3.35. The *switching arrangement $\mathcal{H}_n$* is the collection of hyperplanes

$$H_{(i,j)} = \{\boldsymbol{t} \in \mathbb{R}^n : t_j + t_{j^*} + t_{A(j^*)} = t_i + t_{i^*} + t_{A(i^*)}\}$$

for all couples (i, j) such that there exists a non-crossing non-universal cubic arborescence $A : [n] \to [n]$ with i forward-sliding and minimal, and j backward-sliding and maximal and a diagonal switch can be performed with respect to i and j.

For $\boldsymbol{t} \in \mathbb{R}^n$, let $\mathcal{A}(\boldsymbol{t})$ be the collection of non-crossing arborescences $A : [n] \to [n]$ such that $\boldsymbol{t} \in \mathcal{T}_3^\circ(A)$. Observe, that $A \in \mathcal{A}(\boldsymbol{t})$ for any universal $A : [n] \to [n]$ and any $\boldsymbol{t} \in \mathsf{O}_n^\circ$. We can prove the last main theorem of this section:

Theorem 3.36. *For all $\boldsymbol{t}, \boldsymbol{t}' \in \mathsf{O}_n^\circ \setminus \bigcup_{H \in \mathcal{H}_n} H$, one has $|\mathcal{A}(\boldsymbol{t})| = |\mathcal{A}(\boldsymbol{t}')|$.*

Proof. By Theorem 3.33, if $\boldsymbol{t}$ and $\boldsymbol{t}'$ belong to the same maximal cone of $\mathsf{O}_n^\circ \setminus \bigcup_{H \in \mathcal{H}_n} H$, then $\mathcal{A}(\boldsymbol{t}) = \mathcal{A}(\boldsymbol{t}')$.

Suppose $\boldsymbol{t} \in \mathsf{C}$ and $\boldsymbol{t}' \in \mathsf{C}'$ where C and C' are two adjacent maximal cones of $\mathsf{O}_n^\circ \setminus \bigcup_{H \in \mathcal{H}_n} H$. Then C and C' are separated by a hyperplane $H_{(i,j)}$.

For $A \in \mathcal{A}(\boldsymbol{t})$, if (i,j) is not a couple with i minimal forward-sliding in A and j maximal backward-sliding in A such that a diagonal switch can be performed, then $A \in \mathcal{A}(\boldsymbol{t}')$, as the segment $[\boldsymbol{t}, \boldsymbol{t}']$ does cross any facet of $\mathcal{T}_3(A)$.

Suppose the converse. Then $A \notin \mathcal{A}(\boldsymbol{t}')$, but we can perform a diagonal switch on A with respect to i and j to obtain $A' = F_i B_j A$. We are going to prove that $A' \in \mathcal{A}(\boldsymbol{t}')$. For a minimal forward-sliding in A', then one can list the possibilities: either $a = j$, or $a^* = i$, or $a^* = j^*$, or a is forward-sliding in A (all other possibilities contradict the minimality of a or the fact that $\mu(A') = 3$). In all these cases, as A is captured on $\boldsymbol{t}$, we have $t_i + t_{i^*} + t_{A(i^*)} \leq t_a + t_{a^*} + t_{A'(a^*)}$. By adjacency of C and C', the segment $[\boldsymbol{t}, \boldsymbol{t}']$ does not cross any hyperplane of $\mathcal{H}_n$ other than $H_{(i,j)}$, so $t'_j + t'_{j^*} + t'_{A'(j^*)} \leq t'_a + t'_{a^*} + t'_{A'(a^*)}$. The same arguments ensure that $t'_i + t'_{i^*} + t'_{A'(i^*)} \leq t'_b + t'_{b^*} + t'_{A'(b^*)}$ for all b maximal back-sliding in A'. By construction of $\boldsymbol{t}'$, we have $t'_i + t'_{i^*} + t'_{A'(i^*)} \leq t'_j + t'_{j^*} + t'_{A'(j^*)}$. Consequently, $\boldsymbol{t}' \in \mathcal{T}_3(A')$, meaning that $A' \in \mathcal{A}(\boldsymbol{t}')$.

We have proven that $|\mathcal{A}(\boldsymbol{t})| \geq |\mathcal{A}(\boldsymbol{t}')|$. By symmetry, both quantities are equal. The theorem results from the fact that the graph of maximal cones of $\mathsf{O}_n^\circ \setminus \bigcup_{H \in \mathcal{H}_n} H$ is connected. $\square$

Corollary 3.37. *If $\boldsymbol{t} \in \mathsf{O}_n^\circ \setminus \bigcup_{H \in \mathcal{H}_n} H$, then $|\mathcal{A}(\boldsymbol{t})| = \binom{n}{2} - 1$.*

Proof. By Theorem 3.36 it is enough to show $|\mathcal{A}(\boldsymbol{t})| = \binom{n}{2} - 1$ for some $\boldsymbol{t} \in \mathsf{O}_n^\circ \setminus \bigcup_{H \in \mathcal{H}_n} H$. Let $\boldsymbol{t}_{lex} = (2, 2^2, \ldots, 2^n)$. Recall that for $m \in \mathbb{N}$, $2^m - 1 = \sum_{i=0}^{m-1} 2^i$ and thus $2^m > \sum_{i=0}^{m-1} 2^i$. Thus for any triples $(i_1, i_2, i_3), (j_1, j_2, j_3)$ with $1 \leq i_1 < i_2 < i_3 \leq n$ and $1 \leq j_1 < j_2 < j_3 \leq n$, $2^{i_1} + 2^{i_2} + 2^{i_3} = 2^{j_1} + 2^{j_2} + 2^{j_3}$ if and only if $(i_1, i_2, i_3) = (j_1, j_2, j_3)$. This implies that $\boldsymbol{t}_{lex} \in \mathsf{O}_n^\circ \setminus \bigcup_{H \in \mathcal{H}_n} H$.

Let $A : [n] \to [n]$ be a non-crossing arborescence with $\mu(A) \leq 3$. If A is universal, then $A \in \mathcal{A}(\boldsymbol{t}_{lex})$. There are two such arborescences of intrinsic degree 2 by Corollary 3.18 and $n - 3$ of intrinsic degree 3 by Corollary 3.19(i). Otherwise, $\mathbb{L}(A) = \mathbb{L}^\circ(A) = \{\ell\} \subseteq [2, n-2]$, see Corollary 3.19. By Lemma 3.26, if j is backward-sliding and maximal, then $j < \ell$, $j^* > \ell$ and $A(j^*) = j^* + 1$. If i is forward-sliding and minimal, then $i \leq \ell$ and $A(i^*) - 1 = i^* > \ell$. Consequently, $A \in \mathcal{A}(\boldsymbol{t}_{lex})$ if and only if for all i minimal forward and j maximal backward:

$$2^j + 2^{j^*} + 2^{A(j^*)} < 2^i + 2^{i^*} + 2^{A(i^*)}. \tag{12}$$

As A is non-crossing, if $j < i$, then $j^* \geq i^*$. Assume $j^* \geq i^* + 1$. Then $2^{A(j^*)} > 2^i + 2^{i^*} + 2^{A(i^*)}$ and thus $A \notin \mathcal{A}(\boldsymbol{t}_{lex})$. Otherwise, if $i^* = j^*$, then Equation (12) holds. If $j \geq i$, then $A(i^*) > j^*$ and thus Equation (12) also holds. Consequently, if $A \in \mathcal{A}(\boldsymbol{t}_{lex})$, then there is $p \in [\ell]$ and $k \in [\ell + 1, n - 1]$ such that for all $i < p$, $A(i) = k$ and for all $p \leq i < \ell$, $A(i) = k + 1$. If $k > \ell + 1$, then ℓ is forward-sliding and minimal and $2^{k+1} > 2^\ell + 2^{\ell+1} + 2^{\ell+2}$. Thus $A \in \mathcal{A}(\boldsymbol{t}_{lex})$ if and only if $k = \ell + 1$. The number of those arborescences is

$$\sum_{\ell=2}^{n-2} \sum_{p=1}^{\ell} 1 = \sum_{\ell=2}^{n-2} \ell = \binom{n-1}{2} - 1.$$

Hence in total

$$|\mathcal{A}(\boldsymbol{t}_{lex})| = \binom{n-1}{2} - 1 + (n-1) = \binom{n}{2} - 1.$$

$\square$

Example 3.38. We consider the subdivision of O_n induced by $\mathcal{H}_n$, and then merge the maximal cones C and C' such that $\mathcal{A}(\boldsymbol{t}) = \mathcal{A}(\boldsymbol{t}')$ for $\boldsymbol{t} \in \mathsf{C}$ and $\boldsymbol{t}' \in \mathsf{C}'$. Said differently, we consider the subdivision $\boldsymbol{S}_n$ of O_n induced by the family of polytopes $\mathcal{T}_3(A)$ for A with $\mu(A) \leq 3$.

For $n = 6$, thanks to a computer analysis, we can show that $\boldsymbol{S}_6$ is composed of 12 cones, separated by the 5 hyperplanes of $\mathcal{H}_6$. Corollary 3.37 ensures that to each cone C correspond a 14-gon whose vertices are in bijection with the 14 non-crossing arborescences A with $\mu(A, \boldsymbol{t}) \leq 3$ for all $\boldsymbol{t} \in \mathsf{C}$. In Figure 43 is pictured the dual graph of $\boldsymbol{S}_6$: each maximal cone is represented by its 14-gon whose vertices are labelled by the corresponding non-crossing arborescence. The edges of the dual graph are colored according to the hyperplane of $\mathcal{H}_6$ they correspond to. Moreover, the vertices of the 14-gons are colored with the same colors: hence, following the cyan edge amounts to performing a diagonal switch on the label of the cyan vertex (and similarly for the other colors). As before, green vertices correspond to non-crossing arborescences A with $\mu(A) = 2$, and blue vertices to universal ones with $\mu(A) = 3$.

3.2.3 Pivot polytopes of cyclic polytopes of dimension 2 and 3

The cyclic polytope $\mathsf{Cyc}_d(\boldsymbol{t})$ has a complete graph for $d \geq 4$, but for $d = 2$ and $d = 3$, one can also define the max-slope pivot polytope $\Pi(\mathsf{Cyc}_d(\boldsymbol{t}), \boldsymbol{e}_1)$, even though it will not be the projection of an associahedron. Hence, its vertices will not be associated (in general) to non-crossing arborescences, but just to arborescences on n nodes. For the sake of completeness, we present here the study of the cases $d = 2$ and $d = 3$. To this end, we make use of the method developed in the previous sections (and the proof will be exposed in a more concise way).

Dimension 2 In dimension 2, whatever the chosen $\boldsymbol{t} \in \mathbb{R}^n$, $n \geq 4$, the cyclic polytope $\mathsf{Cyc}_2(\boldsymbol{t})$ is not neighborly: as it is a polygon, its graph is not complete. Thus, its max-slope pivot polytope is not in general the projection of an associahedron. In this section, we describe the max-slope pivot polytope of $\mathsf{Cyc}_2(\boldsymbol{t})$ for the objective function $\boldsymbol{e}_1$.

Theorem 3.39. *For all $\boldsymbol{t} \in \mathsf{O}_n^\circ$, the 1-dimensional polytope $\Pi(\mathsf{Cyc}_2(\boldsymbol{t}), \boldsymbol{e}_1)$ has two vertices, one corresponding to the arborescence $A_m^{(2)}$ defined by $A_m^{(2)}(i) = i+1$ for all $i \in [n]$, and one corresponding to $A_M^{(2)}$ defined by $A_M^{(2)}(1) = n$ and $A_M^{(2)}(i) = i+1$ for $i \neq 1$, see Figure 44.*

Proof. As $\mathsf{Cyc}_2(\boldsymbol{t})$ is 2-dimensional, $\Pi(\mathsf{Cyc}_2(\boldsymbol{t}), \boldsymbol{e}_1)$ is 1-dimensional and has precisely two vertices.

Fix $\boldsymbol{t} \in \mathsf{O}_n^\circ$. When orienting $\mathsf{Cyc}_2(\boldsymbol{t})$ along $\boldsymbol{e}_1$, the only improving neighbor of $\boldsymbol{\gamma}_d(t_i)$ is $\boldsymbol{\gamma}_d(t_{i+1})$ for $i \neq 1$; on the other hand, $\boldsymbol{\gamma}_d(t_1)$ has two improving neighbors: $\boldsymbol{\gamma}_d(t_2)$ and $\boldsymbol{\gamma}_d(t_n)$. Thus, there are exactly two possible arborescences on $\mathsf{Cyc}_2(\boldsymbol{t})$: $A_m^{(2)}$ and $A_M^{(2)}$. Consequently $A_m^{(2)}$ and $A_M^{(2)}$ correspond to the two vertices of $\Pi(\mathsf{Cyc}_2(\boldsymbol{t}), \boldsymbol{e}_1)$. ☐

Remark 3.40. The arborescences $A_m^{(2)}$ and $A_M^{(2)}$ are universal in the sense that they appear as vertices of $\Pi(\mathsf{Cyc}_2(\boldsymbol{t}), \boldsymbol{e}_1)$ for all $\boldsymbol{t} \in \mathsf{O}_n^\circ$. Note that $A_m^{(2)}$ is captured on all $\boldsymbol{t} \in \mathsf{O}_n^\circ$ by any polynomial $P(t) = a_2t^2 + a_1t + a_0$ with $a_2 < 0$ whereas $A_M^{(2)}$ is captured when $a_2 > 0$.

Last but not least, although $A_m^{(2)} = A_m$ is a non-crossing arborescence of degree 2 in the sense of the previous section, this is not the case of $A_M^{(2)} \neq A_M$.

Dimension 3 In dimension $d = 3$, whatever the chosen $\boldsymbol{t} \in \mathsf{O}_n$, $n \geq 5$, the cyclic polytope $\mathsf{Cyc}_3(\boldsymbol{t})$ is not neighborly: its graph is not complete. Thus, its max-slope pivot polytope is not in general the projection of an associahedron. In this section, we will explore the max-slope pivot polytope of the cyclic polytope $\mathsf{Cyc}_3(\boldsymbol{t})$. First, we need to understand the edges of $\mathsf{Cyc}_3(\boldsymbol{t})$.

Lemma 3.41. *When orienting $\mathsf{Cyc}_3(\boldsymbol{t})$ along $\boldsymbol{e}_1$, the vertex $\boldsymbol{\gamma}_d(t_i)$ has at most 2 improving neighbors for $i \neq 1$: $\boldsymbol{\gamma}_d(t_j)$ with $j \in \{i+1, n\}$. Moreover, every $\boldsymbol{\gamma}_d(t_j)$ for $1 < j \leq n$ is an improving neighbor of $\boldsymbol{\gamma}_d(t_1)$.*

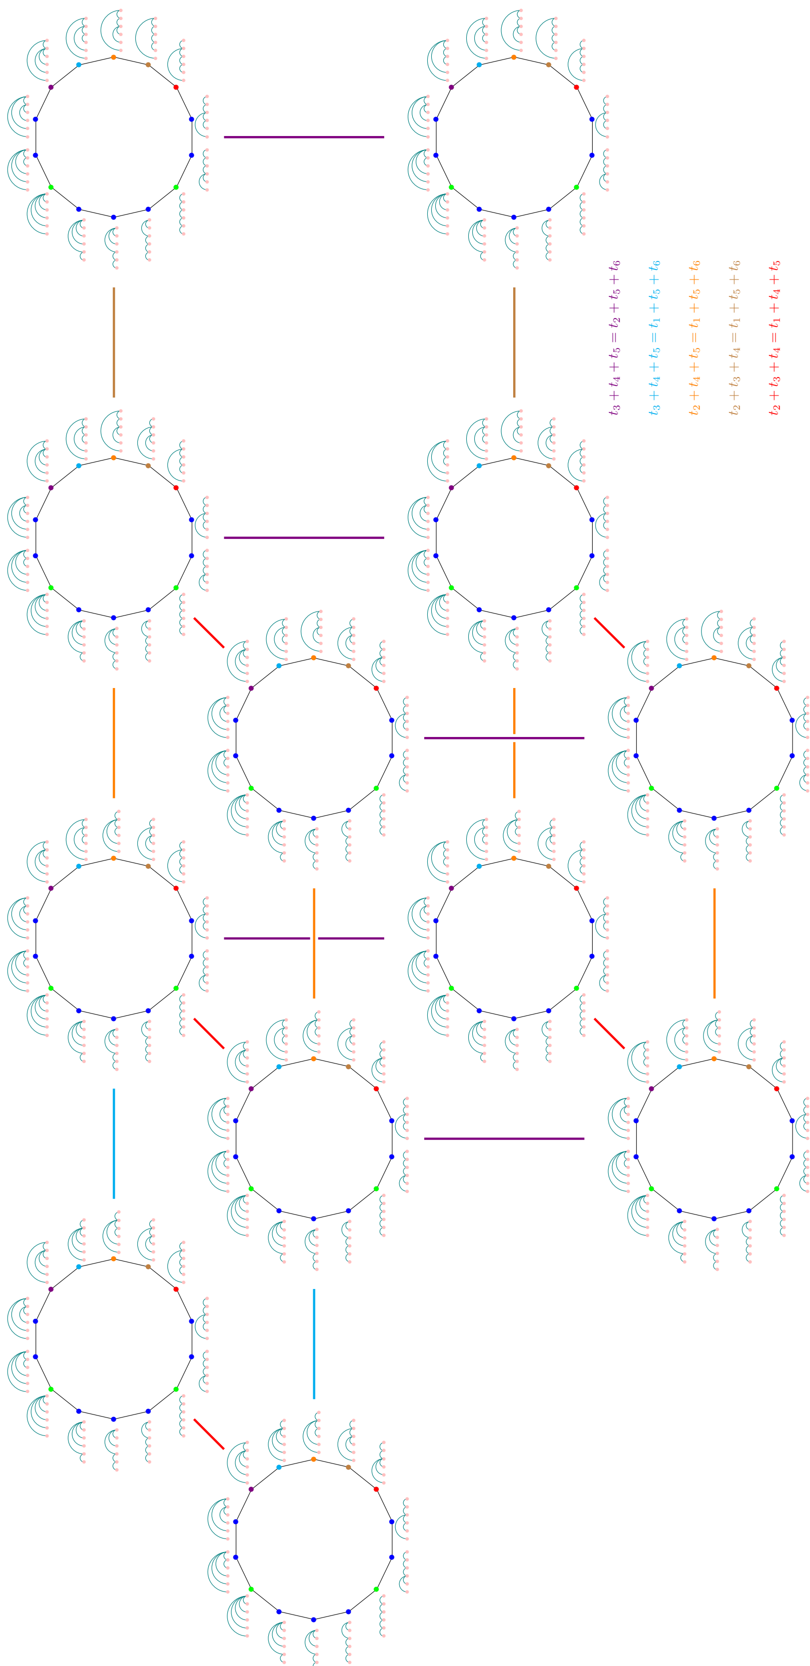

Figure 43: Dual graph of the subdivision of O_6 induced by $\mathcal{T}_3(A)$ for A non-crossing with $\mu(A) \leq 3$.