FPSAC 2009, Hagenberg, Austria DMTCS proc. **AK**, 2009, 647–660

Triangulations of root polytopes and reduced forms (Extended abstract)

Karola Mészáros[1]

[1]*Department of Mathematics, Massachusetts Institute of Technology, Cambridge MA 02139 USA*

Abstract. The type A_n root polytope $\mathcal{P}(A_n^+)$ is the convex hull in $\mathbb{R}^{n+1}$ of the origin and the points $e_i - e_j$ for $1 \leq i < j \leq n+1$. Given a tree T on vertex set $[n+1]$, the associated root polytope $\mathcal{P}(T)$ is the intersection of $\mathcal{P}(A_n^+)$ with the cone generated by the vectors $e_i - e_j$, where $(i,j) \in E(T)$, $i < j$. The reduced forms of a certain monomial $m[T]$ in commuting variables x_{ij} under the reduction $x_{ij}x_{jk} \rightarrow x_{ik}x_{ij} + x_{jk}x_{ik} + \beta x_{ik}$, can be interpreted as triangulations of $\mathcal{P}(T)$. If we allow variables x_{ij} and x_{kl} to commute only when i, j, k, l are distinct, then the reduced form of $m[T]$ is unique and yields a canonical triangulation of $\mathcal{P}(T)$ in which each simplex corresponds to a noncrossing alternating forest.

Résumé. Le polytope des racines $\mathcal{P}(A_n^+)$ de type A_n est l'enveloppe convexe dans $\mathbb{R}^{n+1}$ de l'origine et des points $e_i - e_j$ pour $1 \leq i < j \leq n+1$. Étant donné un arbre T sur l'ensemble des sommets $[n+1]$, le polytope des racines associé, $\mathcal{P}(T)$, est l'intersection de $\mathcal{P}(A_n^+)$ avec le cône engendré par les vecteurs $e_i - e_j$, où $(i,j) \in E(T)$, $i < j$. Les formes réduites d'un certain monôme $m[T]$ en les variables commutatives x_{ij} sous la reduction $x_{ij}x_{jk} \rightarrow x_{ik}x_{ij} + x_{jk}x_{ik} + \beta x_{ik}$ peuvent être interprétées comme des triangulations de $\mathcal{P}(T)$. Si on impose la restriction que les variables x_{ij} et x_{kl} commutent seulement lorsque les indices i, j, k, l sont distincts, alors la forme réduite de $m[T]$ est unique et produit une triangulation canonique de $\mathcal{P}(T)$ dans laquelle chaque simplexe correspond à une forêt alternée non croisée.

Keywords: root polytope, triangulation, volume, Ehrhart polynomial, reduced form, noncrossing alternating tree

1 Introduction

This work was initially inspired by an exercise in Stanley's Catalan Addendum (S, Exercise 6.C6), which calls on us to consider the monomial $w = x_{12}x_{23}\cdots x_{n,n+1}$ in commuting variables x_{ij}. Starting with $p_0 = w$, produce a sequence of polynomials $p_0, p_1, \ldots, p_m$ as follows. To obtain p_{r+1} from p_r, choose a term of p_r which is divisible by $x_{ij}x_{jk}$, for some i, j, k, and replace the factor $x_{ij}x_{jk}$ in this term with $x_{ik}(x_{ij} + x_{jk})$. Note that p_{r+1} has one more term than p_r. Continue this process until a polynomial p_m is obtained, in which no term is divisible by $x_{ij}x_{jk}$, for any i, j, k. Such a polynomial p_m is a **reduced form** of w. Remarkably, while the reduced form is not unique, it turns out that the number of terms in a reduced form is always the Catalan number $C_n = \frac{1}{n+1}\binom{2n}{n}$.

Certain generalizations of this problem reach far beyond its setting in the world of polynomials. On one hand, the reductions can be interpreted in terms of root polytopes and their subdivisions, yielding a geometric, and subsequently a combinatorial, interpretation of reduced forms. On the other hand, using the combinatorial results obtained about the reduced forms, we obtain a method for calculating the volumes and Ehrhart polynomials of certain root polytopes. This "two way traffic" between the combinatorics of reduced forms and geometry is a satisfying outcome.

Root polytopes were defined by Postnikov in (P). The full root polytope $\mathcal{P}(A_n^+)$, which is the convex hull in $\mathbb{R}^{n+1}$ of the origin and points $e_i - e_j$ for $1 \leq i < j \leq n+1$, already made an appearance in the work of Gelfand, Graev and Postnikov (GGP), who gave a canonical triangulation of it in terms of noncrossing alternating trees on $[n+1]$. We obtain canonical triangulations for all acyclic root polytopes, of which $\mathcal{P}(A_n^+)$ is a special case.

We define **acyclic root polytopes** $\mathcal{P}(T)$ for a tree T on vertex set $[n+1]$ as the intersection of $\mathcal{P}(A_n^+)$ with a cone generated by the vectors $e_i - e_j$, where $(i,j) \in E(T)$, $i<j$. Let

$$\bar{G} = ([n+1], \{(i,j) \mid \text{there exist edges } (i,i_1)\ldots,(i_k,j) \text{ in } G \text{ such that } i<i_1<\ldots<i_k<j\}),$$

denote the **transitive closure** of the graph G. Recall that a graph G on vertex set $[n+1]$ is said to be **noncrossing** if there are no vertices $i<j<k<l$ such that (i,k) and (j,l) are edges in G. A graph G on vertex set $[n+1]$ is said to be **alternating** if there are no vertices $i<j<k$ and (i,j) and (j,k) are edges in G. Alternating trees were introduced in (GGP). Gelfand, Graev and Postnikov (GGP) showed that the number of noncrossing alternating trees on $[n+1]$ is counted by the Catalan number C_n.

Theorem 1 *If T is a noncrossing tree on vertex set $[n+1]$ and $T_1,\ldots,T_k$ are the noncrossing alternating spanning trees of $\bar{T}$, then the root polytopes $\mathcal{P}(T_1),\ldots,\mathcal{P}(T_k)$ are n-dimensional simplices with disjoint interiors whose union is $\mathcal{P}(T)$. Furthermore,*

$$\mathrm{vol}\mathcal{P}(T) = f_{T,n}\frac{1}{n!},$$

where $f_{T,n}$ denotes the number of noncrossing alternating spanning trees of $\bar{T}$.

We can describe the intersections of the top dimensional simplices $\mathcal{P}(T_1),\ldots,\mathcal{P}(T_k)$ in Theorem 1 in terms of certain noncrossing alternating spanning forests of $\bar{T}$ and using this calculate the Ehrhart polynomial of the polytope $\mathcal{P}(T)$. Theorem 1 can also be generalized for any tree T, but we omit these details in this abstract.

The proof of Theorem 1 relies on relating the triangulations of a root polytope $\mathcal{P}(T)$ to reduced forms of a monomial $m[T]$ in variables x_{ij}, which we now define. Let A and B be two associative algebras over the polynomial ring $\mathbb{Q}[\beta]$, where β is a fixed constant, generated by the set of elements $\{x_{ij} \mid 1 \leq i < j \leq n+1\}$ modulo the relation $x_{ij}x_{jk} = x_{ik}x_{ij} + x_{jk}x_{ik} + \beta x_{ik}$. Algebra A is commutative, i.e. it has additional relations $x_{ij}x_{kl} = x_{kl}x_{ij}$ for all i,j,k,l, while B is noncommutative and has additional relations $x_{ij}x_{kl} = x_{kl}x_{ij}$ for i,j,k,l distinct only.

We treat the first relation as a **reduction rule**:

$$x_{ij}x_{jk} \to x_{ik}x_{ij} + x_{jk}x_{ik} + \beta x_{ik}. \tag{1}$$

A **reduced form** of the monomial m in algebra A (algebra B) is a polynomial P_n^A (polynomial P_n^B) obtained by successive applications of reduction (1) until no further reduction is possible, where we allow commuting any two variables (commuting any two variables x_{ij} and x_{kl} where i, j, k, l are distinct) between reductions. Note that the reduced forms are not necessarily unique.

A possible sequence of reductions in algebra A yielding a reduced form of $x_{12}x_{23}x_{34}$ is given by

$$\begin{aligned} x_{12}\boldsymbol{x_{23}x_{34}} &\to \boldsymbol{x_{12}}x_{24}\boldsymbol{x_{23}} + \boldsymbol{x_{12}}x_{34}\boldsymbol{x_{24}} + \beta\boldsymbol{x_{12}x_{24}} \\ &\to \boldsymbol{x_{24}}x_{13}\boldsymbol{x_{12}} + x_{24}x_{23}x_{13} + \beta x_{24}x_{13} + x_{34}x_{14}x_{12} + x_{34}x_{24}x_{14} \\ &\quad +\beta x_{34}x_{14} + \beta x_{14}x_{12} + \beta x_{24}x_{14} + \beta^2 x_{14} \\ &\to x_{13}x_{14}x_{12} + x_{13}x_{24}x_{14} + \beta x_{13}x_{14} + x_{24}x_{23}x_{13} + \beta x_{24}x_{13} \\ &\quad +x_{34}x_{14}x_{12} + x_{34}x_{24}x_{14} + \beta x_{34}x_{14} + \beta x_{14}x_{12} + \beta x_{24}x_{14} \\ &\quad +\beta^2 x_{14} \end{aligned} \tag{2}$$

where the pair of variables on which the reductions are performed is in boldface. The reductions are performed on each monomial separately.

Some of the reductions performed above are not allowed in the noncommutative algebra B. The following is an example of how to reduce $x_{12}x_{23}x_{34}$ in the noncommutative case.

$$\begin{aligned} x_{12}\boldsymbol{x_{23}x_{34}} &\to \boldsymbol{x_{12}x_{24}}x_{23} + \underline{x_{12}x_{34}}x_{24} + \beta\boldsymbol{x_{12}x_{24}} \\ &\to x_{14}\boldsymbol{x_{12}x_{23}} + x_{24}x_{14}x_{23} + \beta x_{14}x_{23} + x_{34}\boldsymbol{x_{12}x_{24}} + \beta x_{14}x_{12} \\ &\quad +\beta x_{24}x_{14} + \beta^2 x_{14} \\ &\to x_{14}x_{13}x_{12} + x_{14}x_{23}x_{13} + \beta x_{14}x_{13} + x_{24}x_{14}x_{23} + \beta x_{14}x_{23} \\ &\quad +x_{34}x_{14}x_{12} + x_{34}x_{24}x_{14} + \beta x_{34}x_{14} + \beta x_{14}x_{12} + \beta x_{24}x_{14} \\ &\quad +\beta^2 x_{14} \end{aligned} \tag{3}$$

In the example above the pair of variables on which the reductions are performed is in boldface, and the variables which we commute are underlined.

The "reason" for allowing x_{ij} and x_{kl} to commute only when i, j, k, l are distinct might not be apparent at first, but as we will prove in section 5 it insures that, unlike in the commutative case, there are unique reduced forms for a natural set of monomials. Kirillov (K) observed previously that the monomial $w = x_{12}x_{23}\cdots x_{n,n+1}$ has a unique reduced form in algebra B and asked for a nice combinatorial proof of this fact. We provide such a proof, as the uniqueness of the reduced form of w is a special case of our results.

Before we can state a simplified version of our main result on reduced forms, we need one more piece of notation. Given a graph G on vertex set $[n+1]$ we associate to it the monomial $m^A[G] = \prod_{(i,j)\in E(G)} x_{ij}$; if G is edge-labeled with labels $1, \ldots, k$, we can also associate to it the noncommutative monomial $m^B[G] = \prod_{a=1}^{k} x_{i_a,j_a}$, where $E(G) = \{(i_a, j_a)_a \mid a \in [k]\}$.

Theorem 2 *Let T be a noncrossing tree on vertex set $[n+1]$, and P_n^A a reduced form of $m^A[T]$. Then,*

$$P_n^A(x_{ij} = 1, \beta = 0) = f_{T,n},$$

where $f_{T,n}$ denotes the number of noncrossing alternating spanning trees of $\bar{T}$.

If we label the edges of T so that it becomes a good tree (to be defined in section 4), then the reduced form P_n^B of the monomial $m^B[T]$ is

$$P_n^B(x_{ij}, \beta = 0) = \sum_{T_0} x^{T_0},$$

where the sum runs over all noncrossing alternating spanning trees T_0 of $\bar{T}$ with lexicographic edge-labels (to be defined in section 5) and x^{T_0} is defined to be the noncommutative monomial $\prod_{l=1}^{n} x_{i_l, j_l}$ if T_0 contains the edges $(i_1, j_1)_1, \ldots, (i_n, j_n)_n$.

We can generalize Theorem 2 for any β; see sections 2 and 5. Theorem 2 can also be generalized for any tree T, but we omit these details in this abstract.

Our extended abstract is organized as follows. In section 2 we reformulate the reduction process in terms of graphs and elaborate further on Theorem 2. In section 3 we discuss acyclic root polytopes and relate them to reduced forms. In section 3 we also outline some of the steps for proving Theorems 1 and 2. The lemmas in section 4 indicate the significance of considering reduced forms in the noncommutative algebra B and prepares the ground for proving Theorem 2. In section 5 we conclude the proofs of Theorems 1 and 2.

2 Reductions in terms of graphs

We can phrase the reduction process described in section 1 in terms of graphs. This view will be useful throughout the abstract. Think of a monomial $m \in A$ as a directed graph G on the vertex set $[n+1]$ with an edge directed from i to j for each appearance of x_{ij} in m. Let $G^A[m]$ denote this graph. If, however, we are in the noncommutative version of the problem, and $m = \prod_{l=1}^{p} x_{i_l, j_l}$, then we can think of m as a directed graph G on the vertex set $[n+1]$ with p edge labels $1, \ldots, p$, such that the edge labeled l is directed from vertex i_l to j_l. Let $G^B[m]$ denote the edge-labeled graph just described. Let $(i,j)_a$ be the notation for an edge (i,j) labeled a. It is straightforward to reformulate the reduction rule (1) in terms of reductions on graphs. If $m \in A$, then it reads as follows.

The **reduction rule for graphs:** Given a graph G_0 on vertex set $[n+1]$ and $(i,j), (j,k) \in E(G_0)$ for some $i < j < k$, let G_1, G_2, G_3 be graphs on vertex set $[n+1]$ with edge sets

$$\begin{aligned} E(G_1) &= E(G_0) \backslash \{(j,k)\} \cup \{(i,k)\}, \\ E(G_2) &= E(G_0) \backslash \{(i,j)\} \cup \{(i,k)\}, \\ E(G_3) &= E(G_0) \backslash \{(i,j)\} \backslash \{(j,k)\} \cup \{(i,k)\}. \end{aligned} \tag{4}$$

We say that G_0 **reduces** to G_1, G_2, G_3 under the reduction rules defined by equations (4).

The reduction rule for graphs $G^B[m]$ with $m \in B$ is explained in section 4.

An A**-reduction tree** $\mathcal{T}_R^A$ for a monomial m_0, or equivalently, the graph $G^A[m_0]$, is constructed as follows. The root of $\mathcal{T}_R^A$ is labeled by $G^A[m_0]$. Each node $G^A[m]$ in $\mathcal{T}_R^A$ has three children, which depend on the choice of the edges of $G^A[m]$ on which we perform the reduction. Namely, if the reduction is performed on edges $(i,j), (j,k) \in E(G^A[m])$, then the three children of the node $G_0 = G^A[m]$ are labeled by the graphs G_1, G_2, G_3 as described by equation (4). For an example of an A-reduction tree, see Figure 1 (disregard the edge-labels).

Summing the monomials to which the graphs labeling the leaves of the reduction tree $\mathcal{T}_R^A$ correspond multiplied by suitable powers of β, we obtain a reduced form of m_0.

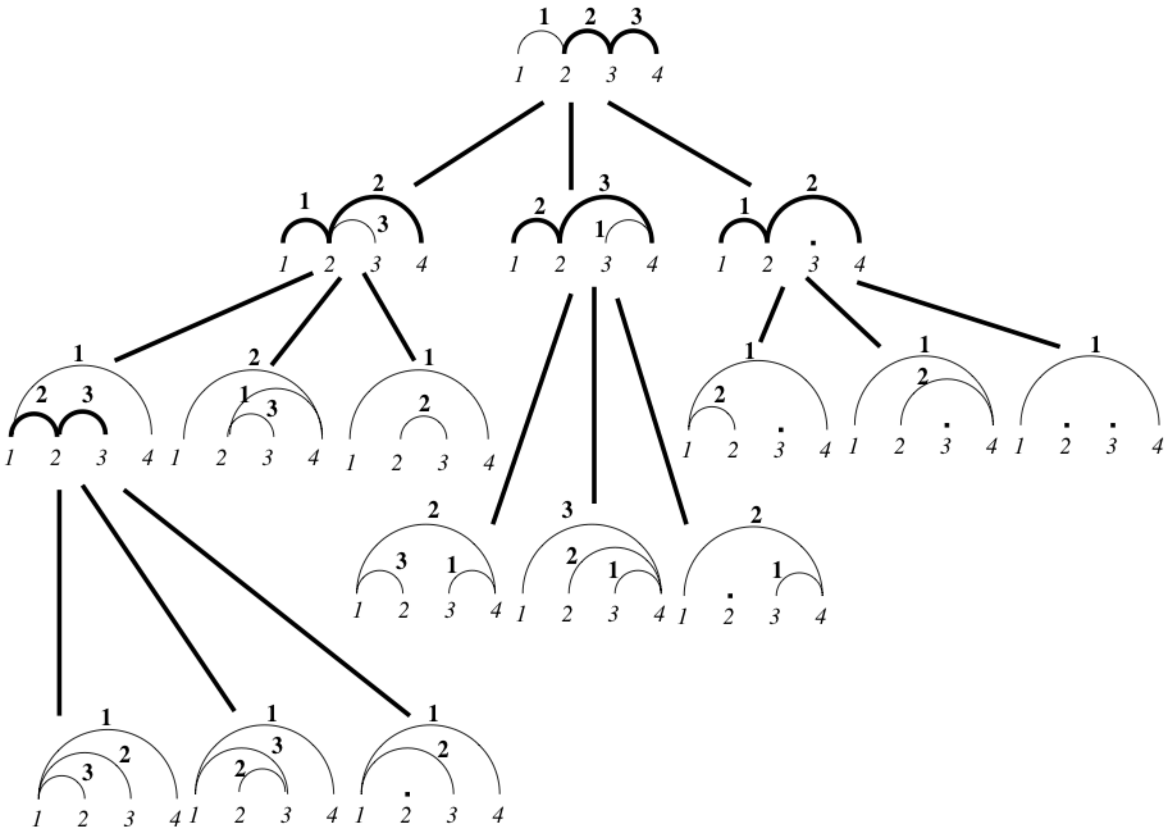

Fig. 1: This is an A-reduction tree with root $m^A[x_{12}x_{23}x_{34}]$, when the edge-labels are disregarded. The boldface edges indicate where the reduction is performed. We can read off the following reduced form of $x_{12}x_{23}x_{34}$ from the set of leaves: $x_{14}x_{13}x_{12} + x_{14}x_{23}x_{13} + \beta x_{14}x_{13} + x_{24}x_{14}x_{23} + \beta x_{14}x_{23} + x_{34}x_{14}x_{12} + x_{34}x_{24}x_{14} + \beta x_{34}x_{14} + \beta x_{14}x_{12} + \beta x_{24}x_{14} + \beta^2 x_{14}$. When the edge-labels are taken into account, this is the B-reduction tree corresponding to equation (3). Note that in the second child of the root we commuted edge-labels 1 and 2.

Let T be a noncrossing tree on vertex set $[n+1]$. In terms of reduction trees, Theorem 2 states that the number of leaves labeled by graphs with exactly n edges of an A-reduction tree with root labeled T is independent of the particular A-reduction tree. The generalization of Theorem 2 states that the number of leaves labeled by graphs with exactly k edges of an A-reduction tree with root labeled T, is independent of the particular A-reduction tree for any k. In terms of reduced forms we can write this as follows. If P_n^A is the reduced form of a monomial $m^A[T]$ for a noncrossing tree T, then

$$P_n^A(x_{ij}=1)=\sum_{m=0}^{n-1} f_{T,n-m}\beta^m,$$

where $f_{T,k}$ denotes the number of noncrossing alternating spanning forests of $\bar{T}$ with k edges and additional technical requirements. Also, if P_n^B is the reduced form of a monomial $m^B[T]$ for a noncrossing good tree T (defined in section 4), then

$$P_n^B(x_{ij})=\sum_F x^F,$$

where the sum runs over all noncrossing alternating spanning forests F of $\bar{T}$ with lexicographic edge-labels (defined in section 5) and additional technical requirements.

If we consider the reduced forms of the path monomial $w=\prod_{i=1}^n x_{i,i+1}$, then $T=P=([n+1],\{(i,i+1)\mid i\in[n]\})$, and $f_{P,k}$ is simply the number of noncrossing alternating spanning forests on $[n+1]$ with k edges containing edge $(1,n+1)$. Furthermore, $P_n^B(x_{ij})=\sum_F x^F$, where the sum runs over all noncrossing alternating spanning forests F on $[n+1]$ with lexicographic edge-labels and containing edge $(1,n+1)$. Let $s_{n,k}$ denote the number of ways to draw k diagonals of a convex $(n+2)$-gon that do not intersect in their interiors.

Lemma 3 *With the notation above,* $f_{P,k+1}=s_{n,k}$.

Cayley (C) in 1890 showed that $s_{n,k}=\dfrac{1}{n+1}\dbinom{n+k+1}{n}\dbinom{n-1}{k}$. Using Lemma 3 and the expression by Cayley, we obtain $P_n^A(x_{ij}=1)=\sum_{m=0}^{n-1} s_{n,n-m-1}\beta^m$.

3 Acyclic root polytopes

In the terminology of (P), a root polytope of type A_n is the convex hull of the origin and some of the points e_i-e_j for $1\le i<j\le n+1$, where e_i denotes the i^{th} coordinate vector in $\mathbb{R}^{n+1}$. A very special root polytope is the full root polytope

$$\mathcal{P}(A_n^+)=\text{ConvHull}(0,e_{ij}^-\mid 1\le i<j\le n+1),$$

where $e_{ij}^-=e_i-e_j$. We study a class of root polytopes including $\mathcal{P}(A_n^+)$, which we now discuss.

Let G be a connected simple graph on vertex set $[n+1]$. Define

$$\mathcal{V}_G=\{e_{ij}^-\mid (i,j)\in E(G), i<j\},\ \text{a set of vectors associated to } G;$$

$$\mathcal{C}_G=\langle\mathcal{V}_G\rangle:=\{\sum_{e_{ij}^-\in\mathcal{V}_G} c_{ij}e_{ij}^-\mid c_{ij}\ge 0\},\ \text{the cone associated to } G;\text{ and}$$

$$\bar{\mathcal{V}}_G = \Phi^+ \cap \mathcal{C}_G, \text{ all the positive roots of type } A_n \text{ contained in the cone associated to } G,$$

where $\Phi^+ = \{e_{ij}^- \mid 1 \leq i < j \leq n+1\}$ is the set of positive roots of type A_n. The idea to consider the positive roots of a root system inside a cone appeared earlier in Reiner's work (R1), (R2) on signed posets.

The root polytope $\mathcal{P}(G)$ associated to graph G is

$$\mathcal{P}(G) = \text{ConvHull}(0, e_{ij}^- \mid e_{ij}^- \in \bar{\mathcal{V}}_G).$$

Note that $\mathcal{P}(A_n^+) = \mathcal{P}(P)$ for the path graph $P = ([n], \{(i, i+1) \mid i \in [n]\})$. While the choice of G such that $\mathcal{P}(A_n^+) = \mathcal{P}(G)$ is not unique, it becomes unique if we require that for no edge $(i,j) \in E(G)$ can the corresponding vector e_{ij}^- be written as a nonnegative linear combination of the vectors corresponding to the edges $E(G) \backslash \{e\}$. Graph P satisfies this requirement.

Define

$$\mathcal{L}_n = \{G = ([n+1], E(G)) \mid \ G \text{ is an acyclic graph}\},$$

and

$$\mathcal{L}(A_n^+) = \{\mathcal{P}(G) \mid G \in \mathcal{L}_n\}, \text{the set of } \textbf{acyclic root polytopes}.$$

Note that the condition that G is an acyclic graph is equivalent to $\mathcal{V}_G$ being a set of linearly independent vectors. To avoid too much detail, in this extended abstract we restrict out attention to the acyclic root polytopes arising from noncrossing trees, however, our methods yield analogous results for all acyclic root polytopes.

The full root polytope $\mathcal{P}(A_n^+) \in \mathcal{L}(A_n^+)$, since the path graph P is acyclic. We show below how to obtain central triangulations for all polytopes $\mathcal{P} \in \mathcal{L}(A_n^+)$. A **central triangulation** of a d-dimensional root polytope $\mathcal{P}$ is a collection of d-dimensional simplices with disjoint interiors whose union is $\mathcal{P}$, the vertices of which are vertices of $\mathcal{P}$ and the origin is a vertex of all of them. Depending on the context we also take the intersections of these maximal simplices to be part of the triangulation.

Lemma 4 *For an acyclic root polytope* $\mathcal{P}(G) = \mathcal{C}_G \cap \mathcal{P}(A_n^+)$.

We now state the crucial lemma which relates root polytopes and algebras A and B defined in section 2.

Lemma 5 ***(Reduction Lemma)*** *Given a graph* $G_0 \in \mathcal{L}_n$ *with* d *edges let* $(i,j), (j,k) \in E(G_0)$ *for some* $i < j < k$ *and* G_1, G_2, G_3 *as described by equation (4). Then* $G_1, G_2, G_3 \in \mathcal{L}_n$*,*

$$\mathcal{P}(G_0) = \mathcal{P}(G_1) \cup \mathcal{P}(G_2)$$

where all polytopes $\mathcal{P}(G_0), \mathcal{P}(G_1), \mathcal{P}(G_2)$ *are* d*-dimensional and*

$$\mathcal{P}(G_3) = \mathcal{P}(G_1) \cap \mathcal{P}(G_2) \textit{ is } (d-1)\textit{-dimensional.}$$

What the Reduction Lemma really says is that performing a reduction on graph $G_0 \in \mathcal{L}_n$ is the same as "cutting" the d-dimensional polytope $\mathcal{P}(G_0)$ into two d-dimensional polytopes $\mathcal{P}(G_1)$ and $\mathcal{P}(G_2)$, whose vertex sets are subsets of the vertex set of $\mathcal{P}(G_0)$, whose interiors are disjoint and whose union is $\mathcal{P}(G_0)$.

Proof Idea of Theorems 1 and 2: Let T be any tree on vertex set $[n+1]$ and consider any A-reduction tree $\mathcal{T}_R^A$ with root T. For simplicity only consider the nodes labeled by graphs with n edges, which corresponds to setting $\beta = 0$. Let $T_1, \ldots, T_k$ be the trees with n edges labeling leaves of $\mathcal{T}_R^A$. Then, by applying the Reduction Lemma multiple times, we obtain that $\mathcal{P}(T_1), \ldots, \mathcal{P}(T_k)$ are d-dimensional polytopes, whose vertex sets are subsets of the vertex set of $\mathcal{P}(T)$, whose interiors are disjoint and whose union is $\mathcal{P}(T)$. Clearly then

$$\text{vol}\mathcal{P}(T) = \text{vol}\mathcal{P}(T_1) + \cdots + \text{vol}\mathcal{P}(T_k).$$

Since $T_1, \ldots, T_k$ are leaves of $\mathcal{T}_R^A$, they must be alternating. Postnikov showed (P, Lemma 13.2) that for an alternating tree T_i on vertex set $[n+1]$, the polytope $\mathcal{P}(T_i)$ a simplex with $\text{vol}\mathcal{P}(T_i) = \frac{1}{n!}$. His definitions in (P) differ from ours, but for this alternating tree case his Lemma 13.2 can be reformulated suitably. Thus, we obtain that $\text{vol}\mathcal{P}(T) = \frac{k}{n!}$, where k is the number of leaves of $\mathcal{T}_R^A$ labeled by trees. Note, this shows that the number of leaves of an A-reduction tree labeled by trees is independent of which A-reduction tree we consider. As explained in section 2, summing the monomials to which the graphs labeling the leaves of a reduction tree $\mathcal{T}_R^A$ correspond multiplied by suitable powers of β, we obtain a reduced form of the monomial corresponding to the root of $\mathcal{T}_R^A$. Thus, we just showed that if P_n^A is this reduced form, then $P_n^A(x_{ij} = 1, \beta = 0) = k$, which is part of the statement of Theorem 2. What k exactly is, how to obtain the canonical triangulation described in Theorem 1 and how to express explicitly the reduced form of $m^B[T]$ as stated in Theorem 2 is all outlined in section 5. $\square$

4 Reductions in the noncommutative case

In this section we state two crucial lemmas about reduction (1) in the noncommutative case necessary for proving Theorem 2. While in the commutative case reductions on $G^A[m]$ could result in crossing graphs, we prove that in the noncommutative case exactly those reductions from the commutative case are allowed which result in no crossing graphs, provided that $m = m^B[T]$ for a noncrossing tree T with suitable edge labels specified below. Furthermore, we also show that if there are any two edges (i, j) and (j, k) with $i < j < k$ in a successor of $G^B[m]$, then after suitably many commutations it is possible to apply reduction (1). Thus, once the reduction process terminates, the set of graphs obtained as leaves of the reduction tree are alternating forests. Now, unlike in the commutative case, they are also noncrossing. In fact, each noncrossing alternating spanning forest of $\bar{T}$ satisfying certain additional technical conditions occurs among the leaves of the reduction tree exactly once, yielding a complete combinatorial description of the reduced form of $m^B[T]$.

In terms of graphs the partial commutativity means that if G contains two edges $(i, j)_a$ and $(k, l)_{a+1}$ with i, j, k, l distinct, then we can replace these edges by $(i, j)_{a+1}$ and $(k, l)_a$, and vice versa. Reduction rule (1) on the other hand means that if there are two edges $(i, j)_a$ and $(j, k)_{a+1}$ in G_0, then we replace G_0 with three graphs G_1, G_2, G_3 on vertex set $[n+1]$ and edge sets

$$\begin{aligned}
E(G_1) &= E(G_0)\backslash\{(i,j)_a\}\backslash\{(j,k)_{a+1}\} \cup \{(i,k)_a\} \cup \{(i,j)_{a+1}\} \\
E(G_2) &= E(G_0)\backslash\{(i,j)_a\}\backslash\{(j,k)_{a+1}\} \cup \{(j,k)_a\} \cup \{(i,k)_{a+1}\} \\
E(G_3) &= (E(G_0)\backslash\{(i,j)_a\}\backslash\{(j,k)_{a+1}\})^a \cup \{(i,k)_a\},
\end{aligned} \tag{5}$$

where $(E(G_0)\backslash\{(i,j)_a\}\backslash\{(j,k)_{a+1}\})^a$ denotes the edges obtained from the edges $E(G_0)\backslash\{(i,j)_a\}\backslash\{(j,k)_{a+1}\}$ by reducing the label of each edge which has label greater than a by 1.

A B**-reduction tree** $\mathcal{T}_R^B$ is defined analogously to an A-reduction tree, except we use equation (5) to describe the children. See Figure 1 for an example. A graph H is called a B**-successor** of G if it is obtained by a series of reductions from G. For convenience, we refer to commutativity of x_{ij} and x_{kl} for distinct i, j, k, l as **reduction (2)**, by which we mean the rule $x_{ij}x_{kl} \leftrightarrow x_{kl}x_{ij}$, for i, j, k, l distinct, or, in the language of graphs, exchanging edges $(i, j)_a$ and $(k, l)_{a+1}$ with $(i, j)_{a+1}$ and $(k, l)_a$ for i, j, k, l distinct.

A forest H on vertex set $[n + 1]$ and k edges labeled $1, \ldots, k$ is **good** if it satisfies the following conditions:

(i) If edges $(i, j)_a$ and $(j, k)_b$ are in H, $i < j < k$, then $a < b$.

(ii) If edges $(i, j)_a$ and $(i, k)_b$ in H are such that $j < k$, then $a > b$.

(iii) If edges $(i, j)_a$ and $(k, j)_b$ in H are such that $i < k$, then $a > b$.

(iv) H is noncrossing.

No graph H with a cycle could satisfy all of $(i), (ii), (iii), (iv)$ simultaneously, which is why we only define good forests. Note, however, that any forest H has an edge-labeling that makes it a good forest.

Lemma 6 *If the root of a B-reduction tree is labeled by a good forest F, then all nodes of it are also labeled by good forests.*

A reduction applied to a noncrossing graph G is **noncrossing** if the graphs resulting from the reduction are also noncrossing.

The following is then a trivial corollary of Lemma 6.

Corollary 7 *If G is a good forest, then all reductions that can be applied to G and its B-successors are noncrossing.*

Lemma 8 *Let G be a good forest. Let $(i, j)_a$ and $(j, k)_b$ with $i < j < k$ be edges of G such that no edge of G crosses (i, k). Then after finitely many applications of reduction (2) we can apply reduction (1) to edges (i, j) and (j, k).*

Corollary 9 *If F labels a leaf of a B-reduction tree whose root is labeled by a good forest, then F is a noncrossing alternating forest.*

5 Proof Idea of Theorems 1 and 2

This section is devoted to understanding how to conclude the proofs of Theorems 1 and 2 started at the end of section 3. We first finish the sketch of the proof of Theorem 2, and then conclude with Theorem 1.

To prove Theorem 2 we describe the reduced form of $m^B[T]$ for a good graph T, which, unlike in the commutative case, is unique. For simplicity we lay out the exact details for the monomial $w_B = \prod_{i=1}^{n} x_{i,i+1}$. We index w by B to indicate that we are in the noncommutative algebra B.

Given a noncrossing alternating forest F on vertex set $[n+1]$ with k edges, the **lexicographic order** on its edges is as follows. Edge (i_1, j_1) is less than edge (i_2, j_2) in lexicographic order if $j_1 > j_2$, or $j_1 = j_2$ and $i_1 > i_2$. The forest F is said to have **lexicographic edge-labels** if its edges are labeled with integers $1, \ldots, k$ such that if edge (i_1, j_1) is less than edge (i_2, j_2) in lexicographic order, then the label of (i_1, j_1) is less than the label of (i_2, j_2) in the usual order on the integers. Clearly, given any graph G there is a unique edge-labeling of it which is lexicographic.

Lemma 10 *If a noncrossing alternating forest F is a B-successor of a good forest T, then upon some number of reductions (2) it is possible to obtain a noncrossing alternating forest F' with lexicographic edge-labels.*

Proposition 11 *By choosing the series of reductions suitably, the set of leaves of a B-reduction tree with root $G^B[w_B]$ can be all noncrossing alternating forests F on vertex set $[n+1]$ containing edge $(1, n+1)$ with lexicographic edge-labels.*

Idea of Proof: By Corollary 9 all leaves of a B-reduction tree are noncrossing alternating forests on vertex set $[n+1]$. It is easily seen that they all contain edge $(1, n+1)$. By the correspondence between the leaves of a B-reduction tree and simplices in a triangulation of $\mathcal{P}(G^B[w_B])$, it follows that no forest appears more than once among the leaves. Thus, it suffices to prove that any noncrossing alternating forest F on vertex set $[n+1]$ containing edge $(1, n+1)$ appears among the leaves of a B-reduction tree and that all these forests have lexicographic edge-labels. One can construct such a B-reduction tree inductively. $\square$

Theorem 12 *The set of leaves of a B-reduction tree with root $G^B[w_B]$ is, up to applications of reduction (2), the set of all noncrossing alternating forests with lexicographic edge-labels on the vertex set $[n+1]$ containing edge $(1, n+1)$.*

Idea of Proof: By Proposition 11 there exists a B-reduction tree which satisfies the conditions above. Since the roots of type A_n are unimodular, it can be shown that the number of k-dimensional simplices in a central triangulation of a type A_n root polytope is fixed for any k. Thus, the number of forests on vertex set $[n+1]$ and k edges among the leaves of an B-reduction tree is fixed. Also, no vertex-labeled forest, with edge-labels disregarded, can appear twice among the leaves of a B-reduction tree. Together with Lemma 10 these imply the statement of Theorem 12. $\square$

Using Theorem 12 we obtain the following characterziation of reduced forms of the noncommutative monomial w_B.

Theorem 13 *If the polynomial $P_n^B(x_{ij})$ is a reduced form of w_B, then*

$$P_n^B(x_{ij}) = \sum_F \beta^{n-|E(F)|} x^F,$$

where the sum runs over all noncrossing alternating forests F with lexicographic edge-labels on the vertex set $[n+1]$ containing edge $(1, n+1)$, and x^F is defined to be the noncommutative monomial $\prod_{l=1}^k x_{i_l, j_l}$ if F contains the edges $(i_1, j_1)_1, \ldots, (i_k, j_k)_k$.

As a corollary to Theorem 13 we obtain the other part of Theorem 2 for the commutative monomial $w_A = \prod_{i=1}^n x_{i,i+1}$.

Theorem 14 *If the polynomial $P_n^A(x_{ij})$ is a reduced form of w_A, then*

$$P_n^A(x_{ij} = 1) = \sum_{m=0}^{n-1} s_{n,n-m-1} \beta^m,$$

where $s_{n,k}$ is the number of noncrossing alternating forests on vertex set $[n+1]$ with $k+1$ edges, containing edge $(1, n+1)$.

Idea of Proof: While w_A may have many reduced forms, any reduced form arises from an A-reduction tree, which in turn gives a triangulation of $\mathcal{P}(A_n^+)$. A triangulation of $\mathcal{P}(A_n^+)$ can be shown to have a fixed number of simplices of a certain dimension, using that the positive roots of type A_n are unimodular. Using this it can be shown that there is a fixed number of leaves with k edges in any A-reduction tree. Using Theorem 13 we obtain that there are $s_{n,k}$ leaves with with $k+1$ edges in any A-reduction tree. □

Observe that the above theorems imply that the poset of all noncrossing alternating spanning forests on the vertex set $[n+1]$ containing the edge $(1, n+1)$ equals the face poset of the triangulation of the full type A_n root polytope $\mathcal{P}(A_n^+)$ obtained from the noncommutative process as described in Theorem 12. By face poset we mean the poset whose elements are the top dimensional simplices in the triangulation of $\mathcal{P}(A_n^+)$ and all their nonempty intersections and the order is given by inclusion.

The **Schröder numbers** s_n count the number of ways to draw any number of diagonals of a convex $(n+2)$-gon that do not intersect in their interiors. Recall that in section 2 $s_{n,k}$ denoted the number of ways to draw k diagonals of a convex $(n+2)$-gon that do not intersect in their interiors. Cayley (C) in 1890 showed that $s_{n,k} = \dfrac{1}{n+1}\dbinom{n+k+1}{n}\dbinom{n-1}{k}$. As indicated in Lemma 3, it is not by coincidence that we used $s_{n,k}$ to also denote the number of noncrossing alternating forests on vertex set $[n+1]$ and $k+1$ edges, containing edge $(1, n+1)$.

Theorems 13 and 14 imply Theorem 2 for the special case $T = P = ([n+1], \{(i, i+1) \mid i \in [n]\})$. We can generalize Theorems 12, 13 and 14 to monomials $m^B[T]$, where T is a good tree. Theorem 2 stated in the Introduction is a weaker version of these generalizations, but is easier to state. In the most general statements of Theorems 12, 13 and 14 we need to replace the condition "all noncrossing alternating forests on $[n+1]$ containing edge $(1, n+1)$" with "all noncrossing alternating forests on $[n+1]$ containing edge $(1, n+1)$ and certain technical requirements," the details of which we omit here. The proofs of the analogous statements use the statements of Theorems 12, 13 and 14 as base cases. If the polynomial $P_n^B(x_{ij})$ is a reduced form of $m^B[T]$ for a good tree T, then

$$P_n^B(x_{ij}) = \sum_F \beta^{n-|E(F)|} x^F,$$

where the sum runs over all noncrossing alternating spanning forests F of $\bar{T}$ with lexicographic edge-labels on the vertex set $[n+1]$ containing edge $(1, n+1)$ and satisfying some technical requirement. Also,

$$P_n^A(x_{ij} = 1) = \sum_{m=0}^{n-1} f_{T,n-m} \beta^m,$$

where $f_{T,n-m}$ is the number of noncrossing alternating forests on vertex set $[n+1]$ with $n-m$ edges, containing edge $(1, n+1)$ and satisfying some technical requirement.

We are now ready conclude the proof of Theorem 1. Recall that at the end of section 3 we proved that if $T_1, \ldots, T_k$ are the trees labeling leaves of $\mathcal{T}_R^A$ with root T, then $\mathcal{P}(T_1), \ldots, \mathcal{P}(T_k)$ form a central triangulation of $\mathcal{P}(T)$. Note that the set of leaves of a B-reduction tree $\mathcal{T}_R^B$ can also be obtained as a set of leaves of some A-reduction tree $\mathcal{T}_R^A$, by simply disregarding the edge-labels of the graphs corresponding to the nodes of $\mathcal{T}_R^B$. The generalization of Theorem 12 implies that that the set of leaves of a B-reduction tree with root $m^B[T]$ which are trees are all noncrossing alternating spanning trees of $\bar{T}$ with lexicographic ordering. Thus, there is an A-reduction tree with root $m^A[T]$ whose leaves that are trees are all noncrossing alternating spanning trees of $\bar{T}$. Therefore, if $T_1, \ldots, T_k$ are all noncrossing alternating spanning trees of $\bar{T}$, then the root polytopes $\mathcal{P}(T_1), \ldots, \mathcal{P}(T_k)$ are n-dimensional simplices with disjoint interiors whose union is $\mathcal{P}(T)$, yielding the canonical triangulation described in Theorem 1. Also, from this triangulation it follows that $\text{vol}\mathcal{P}(T) = f_{T,n}\frac{1}{n!}$, where $f_{T,n}$ denotes the number of noncrossing alternating spanning trees of $\bar{T}$, since as noted at the end of section 3 each $\mathcal{P}(T_i)$ has volume $\frac{1}{n!}$. This concludes the proof of Theorem 1.

Theorem 1 can be generalized so that we not only describe the n-dimensional simplices in the triangulation of $\mathcal{P}(T)$, but also describe their intersections in terms of noncrossing alternating spanning forests in $\bar{T}$. Using the special property of Φ^+ that the vectors in it are unimodular, we can also calculate the Ehrhart polynomial of $\mathcal{P}(T)$ for any tree T. We now define Ehrhart polynomials for integer polytopes, and state our main result pertaining to them. For further background on the theory of Ehrhart polynomials see (BR).

Given a polytope $\mathcal{P} \subset \mathbb{R}^{n+1}$, the t^{th} **dilate** of $\mathcal{P}$ is

$$t\mathcal{P} = \{(tx_1, \ldots, tx_{n+1}) | (x_1, \ldots, x_{n+1}) \in \mathcal{P}\}.$$

The **Ehrhart polynomial** of an integer polytope $\mathcal{P} \subset \mathbb{R}^{n+1}$ is then defined to be

$$L_{\mathcal{P}}(t) = \#(t\mathcal{P} \cap \mathbb{Z}^{n+1}).$$

Theorem 15 *The Ehrhart polynomial of the polytope $\mathcal{P}(T)$, where T is a noncrossing tree on vertex set $[n+1]$, is*

$$L_{\mathcal{P}(T)}(t) = (-1)^n \sum_{i=0}^{n} (-1)^i f_{T,i} \binom{t+i}{i},$$

where $f_{T,i}$ is the number of noncrossing alternating forests on vertex set $[n+1]$ with i edges, containing edge $(1, n+1)$ and satisfying some technical requirement.

For $T = P = ([n], \{(i, i+1) \mid i \in [n]\})$ Theorem 15 specializes to the Ehrhart polynomial of $\mathcal{P}(P) = \mathcal{P}(A_n^+)$ with $f_{P,i} = s_{n,i-1}$. The Ehrhart polynomial of $\mathcal{P}(A_n^+)$ was previously calculated by Fong (F) by different methods.

Acknowledgements

I am grateful to my advisor Richard Stanley for suggesting this problem and for many helpful suggestions. I would like to thank Alex Postnikov for sharing his insight into root polytopes. I would also like to thank Anatol Kirillov for drawing my attention to the noncommutative side of the problem.

References

[GGP] I. M. GELFAND, M. I. GRAEV, A. POSTNIKOV, Combinatorics of hypergeometric functions associated with positive roots, in *Arnold-Gelfand Mathematical Seminars: Geometry and Singularity Theory*, Birkhäuser, Boston, 1996, 205–221.

[BR] M. BECK, S. ROBINS, *Computing the continuous discretely*, Springer Science + Business Media, LLCC, 2007.

[C] A. CAYLEY, On the partitions of a polygon, *Proc. Lond. Math. Soc. 22* (1890), 237-262.

[F] W. FONG, *Triangulations and Combinatorial Properties of Convex Polytopes*, Ph.D. Thesis, 2000.

[K] A. KIRILLOV, personal communication, 2007.

[P] A. POSTNIKOV, Permutohedra, associahedra, and beyond, `http://arxiv.org/abs/math.CO/0507163`.

[R1] V. REINER, *Quotients of Coxeter complexes and P-Partitions*, Ph.D. Thesis, 1990.

[R2] V. REINER, Signed posets, *J. Combin. Theory Ser. A 62* (1993), 324-360.

[S] R. STANLEY, Catalan addendum (version of 20 September 2007), `http://www-math.mit.edu/~rstan/ec/catadd.pdf`.