



ELSEVIER

Contents lists available at ScienceDirect

Journal of Combinatorial Theory, Series A

www.elsevier.com/locate/jcta



Combinatorial formulas for \sqcup -coordinates in a totally nonnegative Grassmannian

Kelli Talaska¹

Department of Mathematics, University of Michigan, Ann Arbor, MI 48109, USA

ARTICLE INFO

Article history:

Received 21 May 2009

Available online 3 November 2009

Keywords:

Positroid

Totally nonnegative Grassmannian

Le-diagram

ABSTRACT

Postnikov constructed a decomposition of a totally nonnegative Grassmannian $(Gr_{kn})_{\geq 0}$ into positroid cells. We provide combinatorial formulas that allow one to decide which cell a given point in $(Gr_{kn})_{\geq 0}$ belongs to and to determine affine coordinates of the point within this cell. This simplifies Postnikov's description of the inverse boundary measurement map and generalizes formulas for the top cell given by Speyer and Williams. In addition, we identify a particular subset of Plücker coordinates as a totally positive base for the set of non-vanishing Plücker coordinates for a given positroid cell.

© 2009 Elsevier Inc. All rights reserved.

Postnikov [4] has described a cell decomposition of a totally nonnegative Grassmannian into positroid cells, which are indexed by \sqcup -diagrams; this decomposition is analogous to the matroid stratification of a real Grassmannian given by Gelfand, Goresky, MacPherson, and Serganova [2]. Postnikov also introduced a parametrization of each positroid cell using a collection of parameters which we call \sqcup -coordinates.

In this paper, we give an explicit criterion for determining which positroid cell contains a given point in a totally nonnegative Grassmannian and explicit combinatorial formulas for the \sqcup -coordinates of a point. This generalizes the formulas of Speyer and Williams given for the top dimensional positroid cell [5], and provides a simpler description of Postnikov's inverse boundary measurement map, which was given recursively in [4]. For a fixed positroid cell, our formulas are written in terms of a minimal set of Plücker coordinates, and this minimal set forms a totally positive base (in the sense of Fomin and Zelevinsky [1]) for the set of Plücker coordinates which do not vanish on the specified cell.

E-mail address: kellcar@umich.edu.

¹ The author was supported by NSF Grants DMS-0502170 and DMS-0555880.

| | | | | | | |
|----------|----------|----------|----------|----------|---|----------|
| T_{17} | 0 | 0 | 0 | 0 | 0 | 0 |
| 0 | 0 | 0 | T_{24} | 0 | 0 | 0 |
| 0 | T_{36} | 0 | T_{34} | T_{33} | 0 | T_{31} |
| 0 | 0 | T_{45} | T_{44} | T_{43} | 0 | |
| T_{57} | T_{56} | T_{55} | T_{54} | | | |




Fig. 1. The \mathcal{J} -tableau T and Γ -network N_T for a point in $(Gr_{5,12})_{\geq 0}$. We have $\lambda = (7, 7, 7, 6, 4)$ and $I = \{1, 2, 3, 5, 8\}$.

1. Positroid stratification and the boundary measurement map

In this section, we review Postnikov’s positroid stratification of a totally nonnegative Grassmannian and boundary measurement map.

Let Gr_{kn} denote the Grassmannian of k -dimensional subspaces of \mathbb{R}^n . A point $x \in Gr_{kn}$ can be described by a collection of (projective) Plücker coordinates $(P_J(x))$, indexed by k -element subsets $J \subset [n]$. The *totally nonnegative Grassmannian* $(Gr_{kn})_{\geq 0}$ is the subset of points $x \in Gr_{kn}$ such that all Plücker coordinates $P_J(x)$ can be chosen to be simultaneously nonnegative.

In [2], the authors gave a decomposition of the Grassmannian Gr_{kn} into *matroid strata*. Each stratum satisfies the property that certain Plücker coordinates are zero for all points in the stratum, and the remaining Plücker coordinates are all nonzero. More precisely, for a matroid \mathcal{M} whose bases are k -element subsets of $[n]$, let $S_{\mathcal{M}}$ denote the stratum consisting of precisely the points $x \in Gr_{kn}$ such that $P_J(x) \neq 0$ if and only if $J \in \mathcal{M}$. In particular, each possible vanishing pattern of Plücker coordinates is given by a unique (realizable) matroid \mathcal{M} . In [4], Postnikov studies a natural analogue of the matroid stratification for the totally nonnegative Grassmannian, a decomposition into disjoint *positroid cells* of the form $(S_{\mathcal{M}})_{\geq 0} = S_{\mathcal{M}} \cap (Gr_{kn})_{\geq 0}$.

Definition 1.1. A \mathcal{J} -*diagram* is a partition λ together with a filling of the boxes of the Young diagram of λ with entries 0 and + satisfying the \mathcal{J} -property: there is no 0 which has both a + above it (in the same column) and a + to its left (in the same row).

Replacing the boxes labeled + in a \mathcal{J} -diagram with positive real numbers, called \mathcal{J} -*coordinates*, we obtain a \mathcal{J} -*tableau*. Let \mathbf{T}_L denote the set of \mathcal{J} -tableaux whose vanishing pattern is given by the \mathcal{J} -diagram L . Note that \mathbf{T}_L is an affine space whose dimension is equal to the number of “+” entries in L , which we denote by $|L|$.

For a box B in λ , we let L_B and T_B denote the labels of the box B in the \mathcal{J} -diagram L and the \mathcal{J} -tableau T , respectively.

In the positroid cell decomposition of $(Gr_{kn})_{\geq 0}$ given in [4], the positroid cells are indexed by \mathcal{J} -diagrams L which fit inside a $k \times (n - k)$ rectangle. Further, the positroid cell corresponding to a fixed \mathcal{J} -diagram L is parametrized by the \mathcal{J} -tableaux $T \in \mathbf{T}_L$, i.e., those whose vanishing pattern is given by L .

The parametrization described below is a special case of Postnikov’s boundary measurement map. To give a formula for this parametrization, we need to introduce certain planar networks, called Γ -*networks*, which are in bijection with \mathcal{J} -tableaux.

For each \mathcal{J} -diagram L of shape λ which fits inside a $k \times (n - k)$ rectangle, we will construct a Γ -graph G_L corresponding to L . For each \mathcal{J} -tableau $T \in \mathbf{T}_L$, we will then assign weights to the faces of G_L to obtain a Γ -network N_T .

We begin by establishing the boundary of the planar network. First, we draw a disk whose boundary consists of the north and west edges of the $k \times (n - k)$ box and the path determining the southeast boundary of λ , all shifted slightly northwest. Place a vertex, called a *boundary source*, at the right end of each row (including empty rows) of λ , and a vertex, called a *boundary sink*, at the

bottom of each column of λ (including empty columns). Label these in sequence with the integers $\{1, 2, \dots, n\}$, following the path from the northeast corner to the southwest corner which determines λ . Let $I = \{i_1 < i_2 < \dots < i_k\} \subset [n]$ be the set of boundary sources, so that the complement of I , $[n] \setminus I = \{j_1 < j_2 < \dots < j_{n-k}\}$, is the set of boundary sinks.

Whenever $L_B = +$, we draw the B hook, i.e., the hook whose corner is the northwest corner of the box $B = (r, c)$ (in the r th row from the top and the c th column from the right) and which has a horizontal path directed from the boundary source i_r to the corner and a vertical path directed from the corner to the boundary sink j_c . This process yields a Γ -graph G_L .

To obtain the Γ -network N_T from G_L , we must assign weights to each of the faces. Note that there is exactly one face for each box B in λ satisfying $L_B = +$ (and this face has a portion of the B hook as its northwest boundary), and in addition, there is one face whose northwest boundary is the boundary of the disk. For each box B with $L_B = +$, we assign to the corresponding face the positive real weight T_B . To the face whose northwest boundary is the boundary of the disk, we assign the weight $\prod \frac{1}{T_B}$, taking the product over boxes satisfying $L_B = +$, so that the product of all face weights in N is exactly 1.

Although the complete description of Postnikov’s map given in [4] is far more complicated (see [6] for explicit combinatorial formulas), in the special case of Γ -networks, it can be viewed as an instance of the classical formula of Lindström [3]. This formula is usually given in terms of weights of edges; we apply Postnikov’s transformation from edge weights to face weights [4] to obtain the following restatement of his definition.

Definition 1.2. For each \sqcup -diagram L which fits in a $k \times (n - k)$ rectangle, the *boundary measurement map* $\text{Meas}_L : \mathbf{T}_L \rightarrow (\text{Gr}_{kn})_{\geq 0}$ is defined by

$$P_J(\text{Meas}_L(T)) = \sum_{A \in \mathcal{A}_J(N_T)} \text{wt}(A),$$

where

- N_T is the Γ -network corresponding to the \sqcup -tableau T , and its boundary source set is labeled by I ,
- $\mathcal{A}_J(N_T)$ is the collection of non-intersecting directed path families $A = \{A_i\}_{i \in I}$ in N_T from the boundary sources I to the boundary destinations J ,
- $\text{wt}(A) = \prod_{i \in I} \text{wt}(A_i)$, and
- the weight $\text{wt}(A_i)$ of a path A_i in the family A is the product of the weights of the faces of N_T which lie southeast of A_i .

We note that the destination set J may contain both sources and sinks, so that I and J may overlap, in which case some of the paths in the collection will have zero edges.

Define the set $\mathcal{M}_L \subseteq \binom{[n]}{k}$ by the condition that $J \in \mathcal{M}_L$ if and only if there exists a non-intersecting path collection in the Γ -graph G_L with sources I and destinations J . It can be shown that \mathcal{M}_L has the structure of a matroid, but this is not necessary for our purposes. Further, it is easily verified that for distinct \sqcup -diagrams L and L^* , we have $\mathcal{M}_L \neq \mathcal{M}_{L^*}$.

Theorem 1.3. (See [4].) For each \sqcup -diagram L which fits in a $k \times (n - k)$ rectangle, the map $\text{Meas}_L : \mathbf{T}_L \rightarrow (\text{Gr}_{kn})_{\geq 0}$ is injective, and the image $\text{Meas}_L(\mathbf{T}_L)$ is precisely the positroid cell $(S_{\mathcal{M}_L})_{\geq 0}$.

These positroid cells are pairwise disjoint, and the union $\bigcup_L (S_{\mathcal{M}_L})_{\geq 0}$, taken over all \sqcup -diagrams L which fit inside the $k \times (n - k)$ rectangle, is the entire totally nonnegative Grassmannian $(\text{Gr}_{kn})_{\geq 0}$. Each positroid cell $(S_{\mathcal{M}_L})_{\geq 0}$ is a topological cell; that is, $(S_{\mathcal{M}_L})_{\geq 0}$ is isomorphic to $\mathbb{R}^{|L|}$, where $|L|$ is the number of “+” entries in L . Thus, the positroid cells form a cell decomposition of $(\text{Gr}_{kn})_{\geq 0}$.

In Postnikov’s work [4], this result is proved by giving a recursive algorithm for finding the \sqcup -tableau T corresponding to a given point in $(\text{Gr}_{kn})_{\geq 0}$. In this paper, we obtain explicit combinatorial formulas solving the same problem. This is done in two stages. In Section 2, we give an




Fig. 2. The Γ -graph of an example in $(Gr_{5,12})_{\geq 0}$ and the path families corresponding to $M'((2, 6), \mathcal{M}_L)$ and $M((2, 6), \mathcal{M}_L)$.

explicit rule for determining which positroid cell contains a given point. In Sections 3 and 4, we give two combinatorial formulas for the inverse of each particular map Meas_L (i.e., formulas for the corresponding \sqcup -coordinates) in terms of the relevant Plücker coordinates.

2. Determining the positroid cell of a point in $(Gr_{kn})_{\geq 0}$

In this section, we give an explicit formula for the \sqcup -tableau $L(x)$ that determines which positroid cell $(S_{\mathcal{M}_L})_{\geq 0}$ a given point $x \in (Gr_{kn})_{\geq 0}$ belongs to.

Let $x \in (Gr_{kn})_{\geq 0}$ be given by its Plücker coordinates $(P_J(x) : J \in \binom{[n]}{k})$. Order the k -subsets of $[n]$ lexicographically. That is, a set $A = \{a_1 < a_2 < \dots < a_k\}$ is less than or equal to a set $B = \{b_1 < b_2 < \dots < b_k\}$ if at the smallest index m for which $a_m \neq b_m$, we have $a_m < b_m$.

For $\mathcal{M} \subseteq \binom{[n]}{k}$, let $I = \{i_1 < i_2 < \dots < i_k\}$ be the lexicographically minimum set in \mathcal{M} . Let $[n] \setminus I = \{j_1 < j_2 < \dots < j_{n-k}\}$ be the complement of I . For a box $B = (r, c)$, set $A_{r,c} = [n] \setminus \{i_r + 1, i_r + 2, \dots, j_c - 2, j_c - 1\}$. Set $M(B, \mathcal{M}) = (M'(B, \mathcal{M}) \setminus \{i_r\}) \cup \{j_c\}$, where

$$M'(B, \mathcal{M}) = \text{lexmax}\{J \in \mathcal{M} : J \cap A_{r,c} = I \cap A_{r,c}\}.$$

In plain language, this says that we are taking the maximum over sets J which contain all of the sources outside the open interval from i_r to j_c and none of the sinks, i.e., those sets whose interesting behavior happens *inside* the interval.

Recall that for a \sqcup -diagram L , we have $J \in \mathcal{M}_L$ if and only if there exists a non-intersecting path collection in G_L with source set I and destination set J . Note that the lexicographically minimum set in \mathcal{M}_L labels the sources of the appropriate Γ -graph, so that this set corresponds to the family of $|I|$ zero-edge paths, one from each source to itself.

Lemma 2.1. *Suppose that B is a box in a \sqcup -diagram L of shape $\lambda(L)$. Then*

- (1) $M'(B, \mathcal{M}_L)$ is the destination set of a unique non-intersecting path collection in the Γ -graph G_L , namely the northwest-most path collection whose edges lie strictly southeast of the B hook;
- (2) $M(B, \mathcal{M}_L) \in \mathcal{M}_L$ if and only if $L_B = +$; and
- (3) the vanishing pattern for the Plücker coordinates of $(S_{\mathcal{M}_L})_{\geq 0}$ is uniquely determined by the vanishing pattern of the subset $\{P_{M(B, \mathcal{M}_L)}\}$, ranging over all boxes B in $\lambda(L)$.

Proof. The proof of the first claim is left as a straightforward exercise for the reader; the second and third then follow immediately from the definitions. \square

Example 2.2. On the left in Fig. 2, we have the Γ -graph of the example in Fig. 1. We see that $M'((2, 6), \mathcal{M}_L) = \{1, 2, 7, 9, 10\}$, corresponding to the solid path collection on the right in Fig. 2. Adding in the potential (dotted) hook from $i_2 = 2$ to $j_6 = 11$, we have $M((2, 6), \mathcal{M}_L) = \{1, 7, 9, 10, 11\}$. Since this hook does not occur in the Γ -graph, we must have $P_{M((2,6), \mathcal{M}_L)}(x) = 0$ for this point.

Let $\lambda(\mathcal{M})$ be the partition in the $k \times (n - k)$ rectangle whose southeastern border is given by the path from the northeast corner of the $k \times (n - k)$ rectangle to its southwest corner which has vertical




Fig. 3. Finding the corners of $D_{F(4,4)}$ and $D_{F(1,7)}$.

edges in positions I and horizontal edges in positions $[n] \setminus I$. More precisely, the length of the t th row of λ is the number of elements of $[n] \setminus I$ which are greater than i_t , i.e., $\lambda_t = |\{j_s \in [n] \setminus I : j_s > i_t\}|$.

Theorem 2.3. For $x \in (Gr_{kn})_{\geq 0}$, set $\mathcal{M}(x) = \{J \in \binom{[n]}{k} : P_J(x) \neq 0\}$. Then the filling of $\lambda(\mathcal{M}(x))$ given by

$$L(x)_B = \begin{cases} 0 & \text{if } P_{M(B, \mathcal{M}(x))}(x) = 0, \\ + & \text{if } P_{M(B, \mathcal{M}(x))}(x) \neq 0 \end{cases} \tag{1}$$

is a \lrcorner -diagram, and x lies in the positroid cell $(S_{\mathcal{M}_L(x)})_{\geq 0}$.

Proof. Combining Theorem 1.3 and Lemma 2.1, each point $x \in (Gr_{kn})_{\geq 0}$ lies in a unique positroid cell $(S_{\mathcal{M}_L})_{\geq 0}$ and therefore we must have a unique \lrcorner -diagram L such that $P_{M(B, \mathcal{M}(x))} = P_{M(B, \mathcal{M}_L)}$ for all boxes $B \in \lambda(L) = \lambda(\mathcal{M}(x))$. \square

3. The \lrcorner -tableau associated with a point in $(S_{\mathcal{M}_L})_{\geq 0}$

In Postnikov’s original work, the map from $(Gr_{kn})_{\geq 0}$ to $\bigcup_L \mathbf{T}_L$ is given recursively. In this section, we provide an explicit description of that map. More precisely, given a point $x \in (S_{\mathcal{M}_L})_{\geq 0}$, we give combinatorial formulas for the entries of the parametrizing \lrcorner -tableau, which we call \lrcorner -coordinates for x .

For each box B in λ , let $H(B)$ denote the collection of boxes lying weakly southeast of the B hook. For each box B with $L_B = +$, let $F(B)$ denote the face with northwest corner B , i.e., the collection of boxes which lie in the same face as B in the corresponding Γ -graph G . We may simply write F for $F(B)$ if there is no need to emphasize the northwest corner of F . The \lrcorner -property ensures that the northwest boundary of each face $F = F(B)$ is a portion of a single hook, namely the B hook; we may also refer to this hook as the F hook.

Definition 3.1. In a Γ -graph G , call a collection W of paths a *generalized path* if the paths in W are pairwise disjoint, and no path of W lies southeast of another path in W .

We say that a collection of paths lies (strictly or weakly) southeast of a given generalized path W if each of the edges in the path collection lies (strictly or weakly) southeast of some path of W .

For a generalized path W in a Γ -graph G , let $OC(W)$ denote the set of *outer corners* of W , that is, those boxes B for which the northern and western boundaries of B are both edges of W . We order the outer corners from northeast to southwest. Let $IC(W)$ denote the *inner corners* of W , that is, those boxes B such the northwest boundary of B is formed by portions of the hooks of two consecutive outer corners. Note that an inner corner need not be adjacent to the corresponding outer corners. The \lrcorner -property ensures that each outer or inner corner B satisfies $L_B = +$.

Consider the generalized paths which lie weakly southeast of the F hook and contain the entire southeast border of F ; these generalized paths must all have the same edge set. That is, they are all identical up to addition or removal of paths with zero edges; take D_F to be the unique such generalized path whose paths each consist of at least one edge. Essentially, D_F traces out the southeast boundary of F , but it may be broken into several paths if F borders the boundary of our disk.




Fig. 4. Here we have the poset for the example given in Fig. 1. Dotted lines indicate the boundaries of the original faces, rotated so that the northwest corner is at the lowest point of the poset.

We can see that $\mathcal{OC}(D_F)$ indexes the hooks which determine the southeast boundary of the face F , and $\mathcal{IC}(D_F)$ indexes the hooks which are intersections of two hooks corresponding to adjacent outer corners.

Example 3.2. Consider the Γ -graph in Fig. 3. We find the inner and outer corners of $D_{F(4,4)}$ and of $D_{F(1,7)}$. In each graph, the relevant face is labeled “ F ”, the generalized path D_F is in bold, and outer and inner corners are labeled “oc” and “ic” respectively. Note that $D_{F(4,4)}$ consists of two paths, and it has no inner corners.

We recall that the Möbius function μ_S of a partially ordered set S is given recursively by the rules

$$\begin{aligned} \mu_S(x, x) &= 1, \quad \text{for all } x \in S, \text{ and} \\ \mu_S(x, y) &= - \sum_{x \leq z < y} \mu_S(x, z), \quad \text{for all } x < y \text{ in } S. \end{aligned}$$

For a \sqcup -diagram L , let \mathcal{F}_L denote the set of faces of the Γ -graph G_L which are indexed by $+$ entries in L . We partially order \mathcal{F}_L by the condition that $F_1 \leq_L F_2$ if the F_1 hook lies weakly northwest of the F_2 hook.

Lemma 3.3. Let $\mu_L = \mu_{\mathcal{F}_L}$ denote the Möbius function for \mathcal{F}_L , with the partial order \leq_L . Then for any two faces $F_1 = F(B_1)$ and $F_2 = F(B_2)$ of G_L , we have

$$\mu_L(F_1, F_2) = \begin{cases} 1 & \text{if } F_1 = F_2 \text{ or } B_2 \in \mathcal{IC}(D_{F_1}), \\ -1 & \text{if } B_2 \in \mathcal{OC}(D_{F_1}), \\ 0 & \text{otherwise.} \end{cases}$$

Proof. We see that our Möbius function μ_L has the following interpretation. For a fixed face F_1 , we assign to each face in the collection $H(F_2)$ the quantity $\mu_L(F_1, F_2)$. That is, we count the faces lying southeast of the F_2 hook with signed multiplicity $\mu_L(F_1, F_2)$. By the definition of a Möbius function, this means we want the total count for a face F to be exactly one if $F = F_1$ and zero if $F \neq F_1$. The proof is then completed by a simple inclusion-exclusion argument, which is left to the reader. \square

To avoid unwieldy notation, we will write $M(B)$ and $M'(B)$ in place of $M(B, \mathcal{M}_L)$ and $M'(B, \mathcal{M}_L)$ when the appropriate \sqcup -diagram L is clear from context.

Theorem 3.4. Suppose $x \in (S, \mathcal{M}_L)_{\geq 0}$. Then the \sqcup -coordinates of x are the entries of the \sqcup -tableau $T(x) \in \mathbf{T}_L$ defined below:




Fig. 5. On the left, we have $U_{F(3,6)}$ in bold and $U'_{F(3,6)}$ dotted; on the right, we have $D_{F(3,6)}$ in bold and $D'_{F(3,6)}$ dotted.

$$T(x)_B = \begin{cases} 0 & \text{if } P_{M(B)}(x) = 0, \\ \prod_{L_C=+} \left(\frac{P_{M(C)}(x)}{P_{M'(C)}(x)} \right)^{\mu_L(B,C)} & \text{if } P_{M(B)}(x) \neq 0. \end{cases} \tag{2}$$

That is, $\text{Meas}_L(T(x)) = x$, and $T(x)$ is the unique \sqcup -tableau whose image under Meas_L is x .

Before proving Theorem 3.4, we note that the concrete description of μ_L given in Lemma 3.3 allows us to quickly write the expressions in Eq. (2) by simply inspecting the graph.

Proof. By Theorem 1.3, there exists a unique \sqcup -tableau T satisfying $\text{Meas}_L(T) = x$. Here we show that T must be the \sqcup -tableau $T(x)$ defined above. Suppose that T satisfies $P_J(\text{Meas}_L(T)) = P_J(x)$ for all $J \in \binom{[n]}{k}$. By Theorem 2.3, if $P_{M(B)}(x) = 0$, we must have $L_B = 0$, and therefore $T_B = 0$. Whenever $L_B = +$, we can easily see that the ratio

$$\frac{P_{M(B)}(\text{Meas}_L(T))}{P_{M'(B)}(\text{Meas}_L(T))}$$

is the product of the weights of all faces southeast of the B hook in the Γ -network N_T , each with multiplicity one. By assumption, we have

$$\frac{P_{M(B)}(\text{Meas}_L(T))}{P_{M'(B)}(\text{Meas}_L(T))} = \frac{P_{M(B)}(x)}{P_{M'(B)}(x)}.$$

Since the weight of a hook is simply the product of the weights of faces southeast of the hook, a multiplicative version of Möbius inversion implies that the weight of the face whose northwest corner is B is given by the ratio

$$\prod_{L_C=+} \left(\frac{P_{M(C)}(x)}{P_{M'(C)}(x)} \right)^{\mu_L(B,C)}.$$

Since the positive entries of T are simply the weights of the faces in N_T , the entries of T must be those of $T(x)$ given in Eq. (2). \square

4. \sqcup -coordinates of a positroid cell in terms of a minimal set of Plücker coordinates

By Theorem 1.3, the dimension of a positroid cell $(S_{\mathcal{M}_L})_{\geq 0}$ is $|L|$, the number of “+” entries in the corresponding \sqcup -diagram L . However, finding the \sqcup -coordinates of a point $x \in (S_{\mathcal{M}_L})_{\geq 0}$ via Eq. (2) may require roughly twice this many Plücker variables. In this section, we give a formula for the map from $(S_{\mathcal{M}_L})_{\geq 0}$ to \mathbf{T}_L , using precisely $|L|$ Plücker variables. This formula is, of course, equivalent to our first formula modulo Plücker relations, but we now use exactly the desired number of parameters.

Suppose $x \in (S_{\mathcal{M}_L})_{\geq 0}$ and $\text{Meas}_L(T) = x$. For a box B in λ with $L_B = +$, let $F = F(B)$ be the corresponding face in the Γ -network N_T . We have already defined U_F and D_F . Let U'_F and D'_F be the northwest-most generalized paths lying strictly southeast of U_F and D_F , respectively. See Fig. 5 for an example.




Fig. 6. Finding the weight of a nested path family.

For a generalized path W in a Γ -network N and a box B in λ , set

$$\varepsilon_W(B) = \begin{cases} 1 & \text{if } B \in \mathcal{OC}(W), \\ -1 & \text{if } B \in \mathcal{IC}(W), \\ 0 & \text{otherwise.} \end{cases}$$

Theorem 4.1. Suppose $x \in (S_{\mathcal{M}_L})_{\geq 0}$ and $\text{Meas}_L(T) = x$. Then the \lrcorner -coordinates of x may be written in the alternate form

$$T_B = \begin{cases} 0 & \text{if } P_{M(B)}(x) = 0, \\ \prod_{L_C=+} (P_{M(C)}(x))^{\varepsilon(B,C)} & \text{if } P_{M(B)}(x) \neq 0, \end{cases} \tag{3}$$

where $\varepsilon(B, C) = [\varepsilon_{U_{F(B)}}(C) - \varepsilon_{U'_{F(B)}}(C)] - [\varepsilon_{D_{F(B)}}(C) - \varepsilon_{D'_{F(B)}}(C)]$.

Before proving Theorem 4.1, we first state one nearly immediate corollary using the notion of a totally positive base given in [1].

Corollary 4.2. The set of Plücker coordinates

$$\mathcal{P}_L = \{P_{M(B)} : L_B = +\}$$

forms a totally positive base for the non-vanishing Plücker coordinates $\{P_J : J \in \mathcal{M}_L\}$ of the positroid cell $(S_{\mathcal{M}_L})_{\geq 0}$. That is, every Plücker coordinate P_J with $J \in \mathcal{M}_L$ can be written as a subtraction-free rational expression (i.e., a ratio of two polynomials with nonnegative integer coefficients) in the elements of \mathcal{P}_L , and \mathcal{P}_L is a minimal (with respect to inclusion) set with this property. Further, each P_J with $J \in \mathcal{M}_L$ is a Laurent polynomial in the elements of \mathcal{P}_L , with nonnegative coefficients.

Proof. Suppose $x \in (S_{\mathcal{M}_L})_{\geq 0}$, with $\text{Meas}_L(T) = x$. By Theorem 4.1, every face weight of the Γ -network N_T can be written as a monomial rational expression in the elements of \mathcal{P}_L . Each Plücker coordinate $P_J(x)$ is a sum of products of face weights, by Definition 1.2. It is then clear that each P_J is a Laurent polynomial with nonnegative coefficients in elements of \mathcal{P}_L . Finally, \mathcal{P}_L is minimal, as it is easily verified that the elements of \mathcal{P}_L are algebraically independent. Indeed, the simple form of Eq. (3) shows that we can explicitly construct a network realizing any choice of positive values for the Plücker coordinates in \mathcal{P}_L , starting by choosing appropriate face weights for those faces at the top of the poset in Fig. 4 and working towards the faces at the bottom of the poset. \square

To prove Theorem 4.1, we will need the following technical lemma, which gives the weights of certain nested path families. For a generalized path W , let $\text{Nest}(W)$ denote the northwest-most non-intersecting path family lying weakly southeast of W . That is, $\text{Nest}(W)$ consists of W , the northwest-most generalized path W' which lies strictly southeast of W , the northwest-most generalized path W'' which lies strictly southeast of W' , and so on, until no more paths will fit. (See Fig. 6.)

Lemma 4.3. Suppose T is a \lrcorner -tableau with corresponding Γ -network N_T . Let W be a generalized path in N_T . Then

$$\text{wt}(\text{Nest}(W)) = \prod_{L_C=+} (P_{M(C)}(\text{Meas}_L(T)))^{\varepsilon_W(C)}. \tag{4}$$

Proof. We proceed by induction on the number of outer corners of W . If W has a single outer corner, the result follows from the definition of $M(B)$. Otherwise, assume W has ℓ outer corners and split W as follows: let \overrightarrow{W} be the path determined by the first $\ell - 1$ outer corners of W (ordered from northeast to southwest) and let \widehat{W} be the hook determined by the last outer corner of W . If \overrightarrow{W} and \widehat{W} do not intersect, the result clearly holds. (This can happen when λ is not the full $k \times n$ rectangle.) Otherwise, let \widehat{W} be the hook determined by the inner corner of W which is between the last two outer corners of W .

Now, $\text{Nest}(W)$ is a disjoint union of paths in N_T . Write $\text{Nest}(W)$ as the ordered collection of path families (W_1, W_2, \dots) , where a path Y in $\text{Nest}(W)$ lies in the block W_i if exactly i paths of $\text{Nest}(W)$ lie strictly northwest of Y . (For i large enough, W_i will be empty. Recall that the weight of an empty path collection is 1.) Write $\text{Nest}(\overrightarrow{W})$, $\text{Nest}(\widehat{W})$, and $\text{Nest}(\widehat{W})$ in the same manner.

We claim that for each i , $\text{wt}(\overrightarrow{W}_i) \text{wt}(\widehat{W}_i) = \text{wt}(W_i) \text{wt}(\widehat{W}_i)$. More precisely, let (v_1, \dots, v_m) be the vertices at which \overrightarrow{W}_i and \widehat{W}_i intersect. Then we claim that W_i is the path along edges of \overrightarrow{W}_i or \widehat{W}_i which starts at the source of \overrightarrow{W}_i and takes the northwest-most path between each v_m and v_{m+1} and \widehat{W}_i is the path which starts at the source of \widehat{W}_i and takes the southeast-most path between each v_m and v_{m+1} . This is clearly true for $i = 1$. The remainder, which depends on the \perp -property, is left as an exercise for the reader.

Since the weight of a path family is the product of the weights of the individual paths, we then have

$$\text{wt}(\text{Nest}(W)) = \frac{\prod_B (P_{M(B)}(\text{Meas}_L(T)))^{\varepsilon_{\overrightarrow{W}}(B)} \cdot \prod_B (P_{M(B)}(\text{Meas}_L(T)))^{\varepsilon_{\widehat{W}}(B)}}{\prod_B (P_{M(B)}(\text{Meas}_L(T)))^{\varepsilon_{\widehat{W}}(B)}}$$

which is precisely equation (4), since \overrightarrow{W} has a single outer corner (which is an outer corner of W) and no inner corners, and \widehat{W} has a single outer corner (which is an inner corner of W) and no inner corners. \square

Proof of Theorem 4.1. Suppose W is a generalized path in N_T . Let W' be the northwest-most generalized path lying strictly southeast of W . We can easily see that the ratio $\frac{\text{wt}(\text{Nest}(W))}{\text{wt}(\text{Nest}(W'))}$ is the product of the weights of the faces lying southeast of W , each with multiplicity one, since the weight of each face appearing in this ratio occurs exactly one more time in $\text{wt}(\text{Nest}(W))$ than it does in $\text{wt}(\text{Nest}(W'))$.

Then, since U_F and D_F bound precisely the face $F = F(B)$, the face weight T_B must be given by the ratio

$$\left(\frac{\text{wt}(\text{Nest}(U_F))}{\text{wt}(\text{Nest}(U'_F))} \right) / \left(\frac{\text{wt}(\text{Nest}(D_F))}{\text{wt}(\text{Nest}(D'_F))} \right).$$

Combining this with Eq. (4) then yields Eq. (3), since we require that $P_J(\text{Meas}_L(T)) = P_J(x)$ for all $J \in \binom{[n]}{k}$. \square

Acknowledgments

The author would like to thank Sergey Fomin and Lauren Williams for many helpful conversations and several anonymous referees for useful comments.

References

[1] S. Fomin, A. Zelevinsky, Double Bruhat cells and total positivity, *J. Amer. Math. Soc.* 12 (2) (1999) 335–380.
 [2] I.M. Gelfand, R.M. Goresky, R.D. MacPherson, V.V. Serganova, Combinatorial geometries, convex polyhedra, and Schubert cells, *Adv. Math.* 63 (3) (1987) 301–316.
 [3] B. Lindström, On the vector representations of induced matroids, *Bull. London Math. Soc.* 5 (1973) 85–90.
 [4] A. Postnikov, Total positivity, Grassmannians, and networks, version of October 17, 2007, <http://math.mit.edu/~apost/papers.html>.
 [5] D. Speyer, L. Williams, The tropical totally positive Grassmannian, *J. Algebraic Combin.* 22 (2) (2005) 189–210.
 [6] K. Talaska, A formula for Plücker coordinates associated with a planar network, *Int. Math. Res. Not. IMRN* (2008), article ID rnn081.