MATCHING POLYTOPES, TORIC GEOMETRY, AND THE NON-NEGATIVE PART OF THE GRASSMANNIAN

ALEXANDER POSTNIKOV, DAVID SPEYER, AND LAUREN WILLIAMS

Abstract. In this paper we use toric geometry to investigate the topology of the totally non-negative part of the Grassmannian $(Gr_{kn})_{\geq 0}$. $(Gr_{kn})_{\geq 0}$ is a cell complex whose cells Δ_G can be parameterized in terms of the combinatorics of bicolored planar graphs G. To each cell Δ_G we associate a complete fan $\mathcal{F}_G$ which is normal to a certain polytope $P(G)$. The combinatorial structure of the polytopes $P(G)$ is reminiscent of the well-known Birkhoff polytopes, and we describe their face lattices in terms of matchings and unions of matchings of G. We also demonstrate a close connection between the polytopes $P(G)$ and matroid polytopes. We then use the data of $\mathcal{F}_G$ and $P(G)$ to define an associated toric variety X_G. We use our technology to prove that the cell decomposition of $(Gr_{kn})_{\geq 0}$ is a CW complex, and furthermore, that the Euler characteristic of the closure of each cell of $(Gr_{kn})_{\geq 0}$ is 1.

1. Introduction

The classical theory of total positivity concerns matrices in which all minors are non-negative. While this theory was pioneered by Gantmacher, Krein, and Schoenberg in the 1930's, the past decade has seen a flurry of research in this area initiated by Lusztig [10, 11, 12]. Motivated by surprising connections he discovered between his theory of canonical bases for quantum groups and the theory of total positivity, Lusztig extended this subject by introducing the totally non-negative variety $G_{\geq 0}$ in an arbitrary reductive group G and the totally non-negative part $(G/P)_{\geq 0}$ of a real flag variety (G/P).

Recently Postnikov [14] investigated the combinatorics of the totally non-negative part of a Grassmannian $(Gr_{kn})_{\geq 0}$: he established a relationship between $(Gr_{kn})_{\geq 0}$ and certain planar bicolored graphs, producing a combinatorially explicit cell decomposition of $(Gr_{kn})_{\geq 0}$. To each such graph G he constructed a parameterization $Meas_G$ of a corresponding cell of $(Gr_{kn})_{\geq 0}$ by $(\mathbb{R}_{\geq 0})^{\#\mathrm{Faces}(G)-1}$. In fact this cell decomposition is a special case of a cell decomposition of $(G/P)_{\geq 0}$ which was conjectured by Lusztig and proved by Rietsch [16], although that cell decomposition was described in quite different terms. Other combinatorial aspects of $(Gr_{kn})_{\geq 0}$, and more generally of $(G/P)_{\geq 0}$, were investigated by Rietsch [13, 17] and Williams [21, 22].

2000 *Mathematics Subject Classification.* Primary 05Exx; Secondary 20G20, 14Pxx.
Key words and phrases. Total positivity, Grassmannian, CW complexes, Birkhoff polytope, matching, matroid polytope, cluster algebra.

It is known that $(G/P)_{\geq 0}$ is contractible [10] and it is conjectured that $(G/P)_{\geq 0}$ with its cell decomposition is a regular CW complex homeomorphic to a ball. In [22], Williams proved the combinatorial analogue of this conjecture, proving that the partially ordered set (poset) of cells of $(G/P)_{\geq 0}$ is in fact the poset of cells of a regular CW complex homeomorphic to a ball.

In this paper we give an approach to this conjecture which uses toric geometry to extend $Meas_G$ to a map onto the closure of the corresponding cell of $(Gr_{kn})_{\geq 0}$. Specifically, given a planar bicolored graph G, we construct a toric variety X_G and a rational map $m_G : X_G \to Gr_{kn}$. We show that m_G is well-defined on the totally non-negative part of X_G and that its image is the closure of the corresponding cell of $(Gr_{kn})_{\geq 0}$. The totally non-negative part of X_G is homeomorphic to a certain polytope (the *moment polytope*) which we denote $P(G)$, so we can equally well think of this result as a parameterization of our cell by $P(G)$. The restriction of m_G to the toric interior of the non-negative part of X_G (equivalently, to the interior of $P(G)$) is $Meas_G$.

Our technology proves that the cell decomposition of the totally non-negative part of the Grassmannian is in fact a CW complex. While our map m_G is well-defined on $(X_G)_{\geq 0}$ (which is a closed ball) and is a homeomorphism on the interior, in general m_G is not a homeomorphism on the boundary of $(X_G)_{\geq 0}$; therefore this does not lead directly to a proof of the conjecture. However, we do obtain more evidence that the conjecture is true: using Williams' result [22] that the face poset of $(G/P)_{\geq 0}$ is *Eulerian*, it follows that the Euler characteristic of the closure of each cell of $(Gr_{kn})_{\geq 0}$ is 1.

The most elegant part of our story is how the combinatorics of the planar graph G reflects the structure of X_G. The torus fixed points of X_G correspond to *perfect orientations* of G, equivalently, to planar-perfect matchings of G. The other faces of X_G correspond to certain *elementary subgraphs* of G, that is, to unions of planar-perfect matchings of G. Every face of X_G is of the form $X_{G'}$ for some planar bicolored graph G' obtained by deleting some edges of G, and m_G restricted to $X_{G'}$ is $m_{G'}$. It will follow from this that, for every face Z of X_G, the interior of Z is mapped to the interior of a cell of the totally non-negative Grassmannian with fibers that are simply affine spaces. We hope that this explicit description of the topology of the parameterization will be useful in studying the topology of $(Gr_{kn})_{\geq 0}$.

The structure of this paper is as follows. In Section 2 we review the combinatorics of plabic graphs and perfect orientations, which gives rise to the parameterizations of cells of $(Gr_{kn})_{\geq 0}$. We then show in Section 3 that every perfectly orientable plabic graph has an acyclic perfect orientation. In Section 4 we review toric varieties and their non-negative parts, and prove a lemma which is key to our CW complex result. In Section 5 we prove that the cell decomposition of $(Gr_{kn})_{\geq 0}$ is in fact a CW complex, and we remark on the connection to cluster algebras. In Section 6 we associate to each perfectly orientable plabic graph G a certain fan $\mathcal{F}_G$ of perfect orientations along with a polytope $P(G)$, and we show in Section 7 that $\mathcal{F}_G$ is the dual fan of $P(G)$. In Section 8 we give an inequality description of $P(G)$, and in

Section 9 we give a combinatorial description of the face lattice of $P(G)$ in terms of matchings and unions of matchings of G. In Section 10 we recall the connection between plabic graphs and *positroids*, and describe the relationship between our polytopes $P(G)$ and matroid polytopes. Finally, in Section 11, we give statistics for a few small plabic graphs, including f-vectors, Ehrhart series, volumes, and the degrees of the corresponding toric varieties.

ACKNOWLEDGEMENTS: We are grateful to Vic Reiner for pointing out the similarity between our polytopes $P(G)$ and Birkhoff polytopes, and to Allen Knutson for many helpful conversations.

2. Plabic graphs and perfect orientations

In this section we review material concerning plabic graphs and perfect orientations from [14].

A *planar directed graph* G is a directed graph drawn inside a disk (and considered modulo homotopy). We allow G to have loops and multiple edges. We will assume that G has n *boundary vertices* on the boundary of the disk labeled $b_1, \dots, b_n$ clockwise. The remaining vertices, called the *internal vertices,* are located strictly inside the disk. We will always assume that each boundary vertex b_i is either a *source* or a *sink.* Even if b_i is an *isolated* boundary vertex, *i.e.* a vertex not incident to any edges, we will assign b_i to be a source or a sink.

A *planar directed network* $N = (G, x)$ is a planar directed graph G as above together with *strictly positive* real weights $y_f > 0$ assigned to all faces f of G such that $\prod y_f = 1$. We'll write $U(\text{Faces}(G))$ for the set of functions on $\text{Faces}(G)$ which take positive real values with product one.

For such a network N, the *source set* $I \subset [n]$ and the *sink set* $\bar{I} := [n] \setminus I$ of N are the sets such that b_i, $i \in I$, are the sources of N (among the boundary vertices) and the b_j, $j \in \bar{I}$, are the boundary sinks.

A *planar bicolored graph,* or simply *plabic graph* is a planar (undirected) graph G, defined as above but without orientations of edges, such that each boundary vertex b_i is incident to a single edge, together with a function $col : V \to \{1, -1\}$ on the set V of internal vertices. We will display vertices with $col(v) = 1$ in black, and vertices with $col(v) = -1$ in white.

A *plabic network* $N = (G, y)$ is a plabic graph G together with positive real weights $y_f > 0$ assigned to faces f of G such that $\prod y_f = 1$.

A *perfect orientation* of a plabic graph or network is a choice of orientation of its edges such that each internal vertex v with $col(v) = 1$ ("black") is incident to exactly one edge directed away from v; and each v with $col(v) = -1$ ("white") is incident to exactly one edge directed towards v. A plabic graph or network is called *perfectly orientable* if it has a perfect orientation.

Let us say that a plabic graph or network has *type* (k, n) if it has n boundary vertices and $k - (n - k) = \sum_{v \in V} col(v)\,(\deg(v) - 2)$. In this case, every perfect orientation will have a source set of size k.

Let P be a directed path through G, starting at the boundary vertex b_i and ending at b_j. Let P' be the closed curve in the plane which first travels along

P and then returns from b_j to b_i along the boundary of the disk in which G is embedded, traveling clockwise. For f any face of G, we write $\mathrm{wind}(P, f)$ for the winding number of P' about f, where we consider a path to wind positively about f if it travels *clockwise* around f. We define $\mathrm{weight}(P, y)$ to be the monomial $\prod_{f \in \mathrm{Faces}(G)} y_f^{\mathrm{wind}(P,f)}$. (Since the product of the y variables is 1, we would get the same result if we returned from b_j to b_i in a counterclockwise direction, but it will be convenient to fix a definite choice for $\mathrm{wind}(P, f)$.) Similarly, if C is a cycle in G we set $\mathrm{weight}(C, y) = \prod_{f \in \mathrm{Faces}(G)} y_f^{\mathrm{wind}(C,f)}$.

Let P be a directed graph in G. Perturb P slightly at the vertices of G to make P a smooth path, without adding any new self crossings. Let $\mathrm{wind}(P)$ be the number of times the tangent vector to P winds around, rounded to the nearest integer. We will actually only care about the value of $\mathrm{wind}(P)$ modulo 2.

For any boundary source b_i and boundary sink b_j, we define the boundary measurement to be

$$M_{ij} = \sum_{P : b_i \to b_j} (-1)^{\mathrm{wind}(P)} \mathrm{weight}(P, y).$$

Remark 2.1. It might seem more natural to place weights on the edges of G and let the weight of a path be the product of the weights of the edges of this path. This is in fact the approach taken at first in [14]. Note, however, that in this case the boundary measurements are unchanged if we take an internal vertex v of G and multiply the edges incident on v by α and $1/\alpha$ (according to whether the edge is directed into v or out of v) for some positive scalar α. We call this operation a gauge transformation. The space of positive functions on edges modulo gauge transformations is isomorphic to the space of positive functions on faces with product one. Specifically, if f is a face of G and x is a function on the edges of G, we set $y_f = \prod_{e \subset \partial f} x_e^{\pm 1}$ where the exponent is 1 if e is oriented clockwise in ∂f and -1 if e is oriented counterclockwise. These y coordinates appear to be more fundamental than the x-coordinates, and we will use them exclusively in this paper.

Proposition 2.2. [14, Lemma 4.3] *Each boundary measurement M_{ij} sums to a subtraction-free rational expression, which gives a well-defined function on $(\mathbb{R}_{>0})^{\mathrm{Faces}(G)}$.*

Definition 2.3. Recall that the (real) Grassmannian $Gr_{kn}(\mathbb{R})$ is the space of all k-dimensional subspaces in $\mathbb{R}^n$. One can represent an element of $Gr_{kn}(\mathbb{R})$ as a full-rank $k \times n$ matrix (not in a unique manner). The *totally non-negative part of the Grassmannian* $(Gr_{kn})_{\geq 0}$ is defined to be the subset of the real Grassmannian $Gr_{kn}(\mathbb{R})$ such that all maximal minors are non-negative.

These boundary measurements can be thought of as giving a matrix representation of an element of the Grassmannian. More specifically, for $y \in U(\mathrm{Faces}(G))$, we define a $k \times n$ matrix A by the following conditions: the $k \times k$ submatrix of A whose columns are indexed by I (the sources of G) is the identity matrix and, if i is the r^{th} source (so $1 \leq r \leq k$) and j is a sink then $A_{rj} = \pm M_{ij}$ where the sign is chosen such that the maximal minor with column set $I \cup \{j\} \setminus \{i\}$ has determinant M_{ij}. We define $Meas_G(y)$ to be the point of the Grassmannian corresponding to A.

Proposition 2.4. [14, Proposition 5.3] *Every maximal minor of A is given by a subtraction free rational expression. If the graph G is acyclic, then the minor A_J is given by the sum over all collections of non-crossing paths from $I \setminus J$ to $J \setminus I$ of the product of the weights of the paths.*

This result is essentially the Gesssel-Viennot lemma, but some care is needed to see that the signs introduced into the entries of A interact correctly with the signs coming from the determinant. When G is not acyclic, Postnikov proves a similar but more complex result, which we will not need here.

Let G be a perfectly orientable bi-colored planar graph and let G_1 and G_2 be two perfect orientations of G. Let y be an element of $U(\mathrm{Faces}(G))$. Clearly, the matrix A described above will be different for G_1 and G_2. However, we have:

Proposition 2.5. [14, Theorem 10.1] *The points $Meas_{G_1}(y)$ and $Meas_{G_2}(y)$ of Gr_{kn} are the same.*

Thus, we can think of $Meas_G$ as being indexed by a bicolored planar graph rather than a directed planar graph, and (by a slight abuse of notation) will refer to the map as $Meas_G$ in either situation.

The following proposition summarizes several main results of [14].

Proposition 2.6. [14, Theorem 12.7] *Let G be any perfectly orientable bicolored planar graph. Then the image of $Meas_G$ is a cell of $(Gr_{kn})_{\geq 0}$. For every cell of $(Gr_{kn})_{\geq 0}$, there is a perfectly orientable bicolored planar graph G such that the given cell is the image of $Meas_G$. There is some non-negative integer r, such that $Meas_G$ is a fiber bundle with fiber r-dimensional affine space.*

*For any cell of $(Gr_{kn})_{\geq 0}$, we can always choose a G such that $Meas_G$ is a homeomorphism onto this cell. (*I.e. *such that $r = 0$.)*

The graphs G such that $Meas_G$ is a homeomorphism are called *reduced* [14].

Sometimes it is convenient to transform a plabic graph G into a bipartite plabic graph. We say that a plabic graph G is *bipartite* if any edge in G joins two vertices of different colors. Here we assume that the color of a boundary vertex is different than the color of its (unique) neighbor. Note that we can easily make any plabic graph bipartite by contracting any edges which join vertices of the same color. There is an easy bijection between perfect orientations of the old plabic graph and perfect orientations of the contracted plabic graph, and the boundary measurements are the same. This does not change the number of faces of G.

A *planar-perfect matching* of a bipartite plabic graph G is a subset M of edges such that each internal vertex is incident to exactly one edge in M and the boundary vertices are incident to either one or no edges in M. There is a bijection between perfect orientations of G and planar-perfect matchings of G where, for a perfect orientation $\mathcal{O}$ of G, an edge e is included in the corresponding matching if e is directed from a black vertex to a white vertex in $\mathcal{O}$.

Recall that a *circuit* in a graph is a cycle which does not cross itself. We define a *trail* of a plabic graph G to be either a circuit or a path without self-intersections and with endpoints on the boundary of the disk.

3. Acyclic perfect orientations

In this section we show that every perfectly orientable plabic graph has an acyclic perfect orientation. This will be an important tool for our subsequent results.

According to Proposition 2.6, each plabic graph G corresponds to a cell in $(Gr_{kn})_{\geq 0}$. However different plabic graphs can be associated with the same cell. These equivalence classes of plabic graphs can be described in terms of local transformations (moves). For simplicity, we only deal with *reduced* plabic graphs. (For all plabic graphs one needs to consider additional local transformations (reductions), as described in [14, Theorem 12.1].)

Proposition 3.1. [14, Theorem 12.1] *Two reduced plabic graphs correspond to the same cell in* $(Gr_{kn})_{\geq 0}$ *if and only if they can be obtained from each other by a sequence of moves of three types:* (M1) *square move,* (M2) *uni-colored edge contraction/insertion,* (M3) *middle vertex insertion/removal; see Figure 1.*

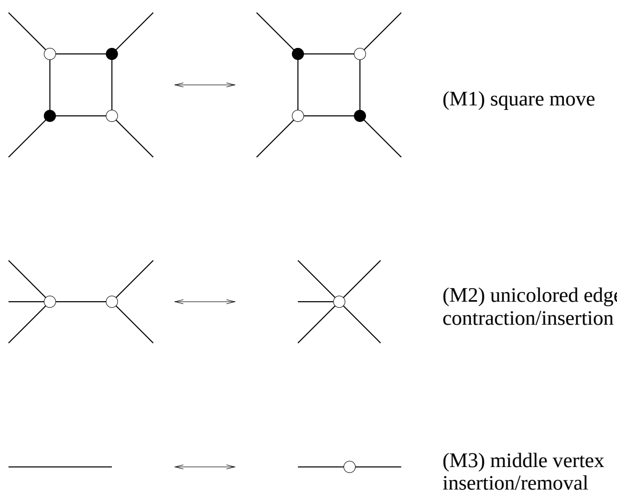

Figure 1. Local transformations of plabic graphs

For each cell in $(Gr_{kn})_{\geq 0}$ there is a unique reduced plabic graph that comes from a Ⅎ-diagram, as described in [14, Section 6]. These special Ⅎ-plabic graphs have very nice structure. By Proposition 3.1 each reduced plabic graph can be transformed by the moves into such special Ⅎ-plabic graph; and Ⅎ-plabic graph are unique representatives of move-equivalence classes of reduced plabic graphs.

We will need the following lemma, which can be proved using techniques of [14].

Lemma 3.2. *Each reduced plabic graph has an acyclic perfect orientation.*

Proof. As we just noted, each reduced plabic graph G can be transformed by the moves into an Ⅎ-plabic graph G'. It is quite clear that an Ⅎ-plabic graph G' has an acyclic perfect orientation. Moreover, one can pick an acyclic orientation of G' such that its set of boundary sources is the lexicographically minimal base of the corresponding matroid.

The fact that a perfect orientation corresponds to the lexicographically minimal base implies that for any directed path in the graph from a boundary vertex b_i to a boundary vertex b_j we have $i > j$. Indeed, if $i < j$, then we can switch directions of edges in the orientation and obtain another orientation that corresponds to a lexicographically smaller base (obtained by replacing j with i).

The original graph G can be obtained from G' by a sequence of moves. Let us show that one can transform an acyclic perfect orientation of G' into an acyclic perfect orientation of G. It is quite clear that each time when we do moves (M2) or (M3) (edge contraction/insertion, vertex insertion/removal) an acyclic perfect orientation is transformed into an acyclic perfect orientation.

We need to be move careful with the square move (M1). Up to symmetries, there are 2 possible perfect acyclic orientations of edges adjacent to the square in a square move, as shown on the left-hand side of Figure 2. These perfect orientation can be transformed as shown on the right-hand side of Figure 2. Let us call these transformations the oriented square move of type I and of type II. Clearly, the oriented square move of type I transforms an acyclic perfect orientation into an acyclic perfect orientation. However, the oriented square move of type II creates a new cycle, so the resulting orientation is no longer acyclic.

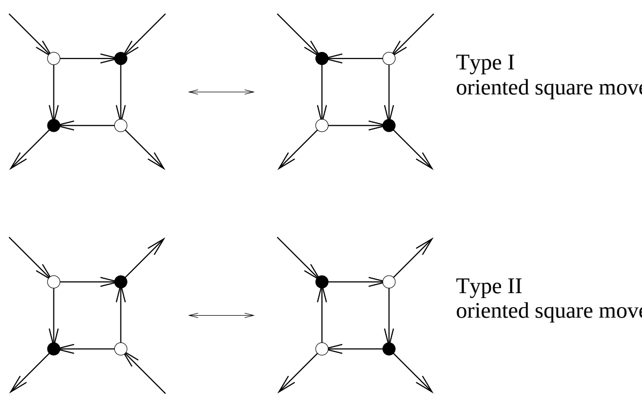

FIGURE 2. Oriented square moves

Let us show that, if we have an acyclic perfect orientation that corresponds to the lexicographically minimal base, then the oriented square move of type II is never possible. In other words, in this case the graph never has a fragment as in the lower left corner of Figure 2. Suppose that we have such a fragment. The 4 outside edges of the square should be connected to the boundary by some directed paths; see Figure 3. These paths cannot cross each other and cannot self-intersect. (Otherwise we get a directed cycle, and the orientation is not acyclic.) It is clear from Figure 3 that we can always find a directed path in the graph from a boundary vertex b_i to a boundary vertex b_j such that $i < j$. Thus the orientation corresponds to a base, which is not lexicographically minimal. This contradicts to our assumption.

It follows that, when we start transforming the graph G' (with acyclic perfect orientation corresponding to the lexicographically minimal base) into the graph

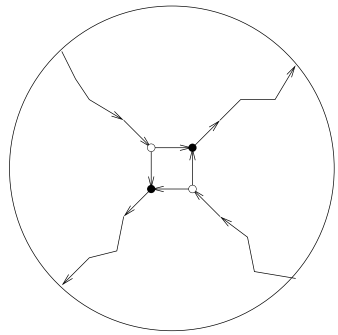

FIGURE 3. Oriented square move of type II

G, we will never perform an oriented square move of type II. Thus at each step we obtain an acyclic perfect orientation. Finally, we obtain an acyclic perfect orientation of the original graph G, as needed. □

4. Toric varieties and their non-negative parts

Here we recall several constructions of a toric variety and define the *non-negative part* of a toric variety. First we give the construction of a toric variety from a fan, following Fulton [6]. Let $N = M$ be dual lattices (free $\mathbb{Z}$-modules of finite rank) with standard basis $e_1, \ldots, e_r$ and $e_1^*, \ldots, e_r^*$. If $\sigma \subset N_{\mathbb{R}}$ is a strongly convex rational polyhedral cone, then $S_\sigma = \sigma^\vee \cap M$ is a finitely generated semigroup, where $\sigma^\vee$ denotes the cone dual to σ. S_σ determines the group ring $\mathbb{C}[S_\sigma]$, which has a basis χ^u, as u varies over S_σ. Multiplication is determined by the addition in S_σ: $\chi^u \cdot \chi^{u'} = \chi^{u+u'}$; the unit 1 is χ^0.

Then $V_\sigma = \mathrm{Spec}(\mathbb{C}[S_\sigma])$ is the affine toric variety associated to σ. In particular, $\mathbb{C}[S_\sigma]$ is the ring of regular functions on V_σ. If Σ is a fan in $N_{\mathbb{R}}$, then X_Σ is the abstract variety constructed from the affine varieties V_σ, $\sigma \in \Sigma$, by gluing V_σ and $V_{\sigma'}$ along their common open subset $V_{\sigma\cap\sigma'}$ for all $\sigma, \sigma' \in \Sigma$. See [6] for more details.

Note that the toric varieties given by the fans in Figures 4 and 5 are $\mathbb{P}^2$ and $\mathbb{P}^1 \times \mathbb{P}^1$, respectively.

For $\sigma \subset N$ any strongly convex rational polyhedral cone, we define the totally non-negative part of V_σ to be the collection of (closed) points on which all of the coordinate functions χ^u, $u \in S_\sigma$, are non-negative real numbers. We will denote the totally non-negative part of V_σ by $(V_\sigma)_{\geq 0}$. If $u_1, \ldots, u_s$ is any set of generators for S_σ, then we can use $\chi^{u_1}, \ldots, \chi^{u_s}$ as coordinate functions to embed V_σ into $\mathbb{C}^s$, and then $(V_\sigma)_{\geq 0}$ is simply the set of points which have non-negative coordinates. If σ and σ' are two cones in a fan Σ, then $(V_\sigma)_{\geq 0} \cap V_{\sigma'} = (V_{\sigma\cap\sigma'})_{\geq 0} = V_\sigma \cap (V_{\sigma'})_{\geq 0}$. Note that if $q = f/g$ is a rational function which is a ratio of two polynomials f and $g \in \mathbb{R}_{\geq 0}[S_\sigma]$, such that the coefficient in g of χ^0 is strictly positive, then f/g is a well defined real valued function on $(V_\sigma)_{\geq 0}$. We denote $\bigcup_{\sigma\in\Sigma}(V_\sigma)_{\geq 0}$ by $(X_\Sigma)_{\geq 0}$, this is a closed subspace of X_Σ.

Now, suppose that Σ is the dual fan to some (bounded) polytope P in $M_{\mathbb{R}}$ with vertices in M. This gives rise to a toric variety X_Σ together with a morphism $\phi : X \to \mathbb{P}^{r-1}$ via the sections χ^u for $u \in P \cap M$ (see Section 3.4 of [6]). There is a map known as the moment map from X_Σ to P. The restriction of the moment map to $(X_\Sigma)_{\geq 0}$ is a homeomorphism. (See [6, Section 4.2], but note that Fulton's sign conventions for normal fans [6, Section 1.5] differ from ours.)

More concretely, we may define a projective toric variety associated to a polytope as follows [3]. Fix a lattice polytope $P \subset \mathbb{R}^n$, and let $\mathbf{m}_i$, $i = 1, \dots, \ell$ be the lattice points of $P \cap \mathbb{Z}^n$. Then consider the map $\phi : (\mathbb{C}^*)^n \to \mathbb{P}^{\ell-1}$ such that $\mathbf{x} = (x_1, \dots, x_n) \mapsto [\mathbf{x}^{\mathbf{m_1}}, \dots, \mathbf{x}^{\mathbf{m}_\ell}]$. If Σ is the dual fan to P, then the toric variety X_Σ (which we also refer to as X_P) is defined to be the closure of the image of this map. The *real part* $X_P(\mathbb{R})$ of X_P is defined to be the intersection of X_P with $\mathbb{RP}^{\ell-1}$; the *positive part* $(X_P)_{>0}$ is defined to be the image of $(\mathbb{R}_{>0})^n$ under ϕ; and the *non-negative part* $(X_P)_{\geq 0}$ is defined to be the closure (in $X_P(\mathbb{R})$) of $(X_P)_{>0}$.

We now prove a simple but very important lemma.

Lemma 4.1. *Suppose we have a map $\Phi : (\mathbb{R}_{>0})^n \to \mathbb{P}^{N-1}$ given by $(t_1, \dots, t_n) \mapsto [h_1(t_1, \dots, t_n), \dots, h_N(t_1, \dots, t_n)]$ where the h_i's are Laurent polynomials with positive coefficients. Let S be the set of all exponent vectors in $\mathbb{Z}^n$ which occur among the (Laurent) monomials of the h_i's, and let P be the convex hull of the points of S. Then the map Φ gives rise to a rational map Φ' from X_P to $\mathbb{P}^{N-1}$ which is well-defined on $(X_P)_{\geq 0}$. In particular, Φ' is a well-defined map from a closed ball to $\mathbb{P}^{N-1}$.*

Proof. Let $S = \{\mathbf{m_1}, \dots, \mathbf{m}_\ell\}$. Clearly the map Φ factors as the composite map $t = (t_1, \dots, t_n) \mapsto [\mathbf{t}^{\mathbf{m_1}}, \dots, \mathbf{t}^{\mathbf{m}_\ell}] \mapsto [h_1(t_1, \dots, t_n), \dots, h_N(t_1, \dots, t_n)]$. Note that the first map is the projective embedding of X_P and the second map Φ' takes a point $[x_1, \dots, x_\ell]$ of X_P to $[g_1(x_1, \dots, x_\ell), \dots, g_N(x_1, \dots, x_\ell)]$, where the g_i's are homogeneous polynomials of degree 1 with positive coefficients. By construction, each x_i occurs in at least one of the g_i's.

Since $(X_P)_{\geq 0}$ is the closure inside X_P of $(X_P)_{>0}$, any point $[x_1, \dots, x_\ell]$ of $(X_P)_{\geq 0}$ has all x_i's non-negative; furthermore, not all of the x_i's are equal to 0. And now since the g_i's have positive coefficients and they involve *all* of the x_i's, the image of any point $[x_1, \dots, x_\ell]$ of $(X_P) \geq 0$ under Φ' is well-defined.

Finally, since $(X_P)_{\geq 0}$ is homeomorphic to the polytope P, the map Φ' is a well-defined map from a closed ball to $\mathbb{P}^{N-1}$. □

In the next section we will use this lemma to prove that $(Gr_{kn})_{\geq 0}$ is a CW complex.

5. $(Gr_{kn})_{\geq 0}$ is a CW complex

In this section we prove that the cell decomposition of $(Gr_{kn})_{\geq 0}$ is a CW complex, and derive some topological consequences regarding the Euler characteristics of closures of cells. We also remark on the connection to cluster algebras.

In a CW complex, a cell is *attached* by gluing a closed i-dimensional ball D^i to the $(i-1)$-skeleton X_{i-1}, i.e. the union of all lower dimensional cells. The gluing is specified by a continuous function f from $\partial D^i = S^{i-1}$ to X_{i-1}. CW complexes are defined inductively as follows. Given X_0 a discrete space (a discrete union of 0-cells), and inductively constructed subspaces X_i obtained from X_{i-1} by attaching some collection of i-cells, the resulting colimit space X is called a *CW complex* provided it is given the weak topology and the *closure-finite* condition is satisfied for its closed cells. Recall that the closure-finite condition requires that every closed cell is covered by a finite union of open cells.

To prove our main result, we will also use the following lemma, which can be found in [14, 17].

Lemma 5.1. [14, 17] *The closure of a cell Δ in $(Gr_{kn})_{\geq 0}$ is the union of Δ together with lower-dimensional cells.*

The following proposition describes a parameterization of a cell of $(Gr_{kn})_{\geq 0}$ using Plucker coordinates, which are expressed as positive Laurent polynomials in terms of the variables on the faces of the plabic graph.

Proposition 5.2. *Each cell Δ_G of $(Gr_{kn})_{\geq 0}$ can be parameterized by a map of the form $(t_1, \dots, t_m) \mapsto [h_1(t_1, \dots, t_m), \dots, h_N(t_1, \dots, t_m)]$ where the h_i's are Laurent polynomials with positive coefficients.*

Proof. By Lemma 3.2, every reduced plabic graph has a perfect orientation which is acyclic. The parameterization that we associate to a perfect orientation only becomes rational (as opposed to Laurent polynomial) if the perfect orientation contains a directed cycle. Therefore an acyclic perfect orientation leads to a parameterization that involves only Laurent polynomials. □

We now define a toric variety which will be an important tool in analyzing the topology of the cell decomposition of $(Gr_{kn})_{\geq 0}$.

Definition 5.3. Define the vector space $V^\vee$ to be the quotient of $\mathbb{R}^{\text{Faces}(G)}$ by $(1, 1, \dots, 1)$. Let S_G be the set of all exponent vectors in $V^\vee$ which occur among the (Laurent) monomials of the h_i's above. Note that these exponent vectors lie in $V^\vee$ because in the parameterizations of the cells of $(Gr_{kn})_{\geq 0}$, the face variables y_f satisfy $\prod y_f = 1$. Let $Q(G)$ be the convex hull of the points of S_G, and define X_G to be the toric variety $X_{Q(G)}$ defined by the polytope $Q(G)$.

Theorem 5.4. *The cell decomposition of $(Gr_{kn})_{\geq 0}$ is a CW complex.*

Proof. All of these cell complexes contain only finitely many cells; therefore the closure-finite condition in the definition of a CW complex is automatically satisfied. What we need to do is define the attaching maps for the cells: we need to prove that for each i-dimensional cell there is a continuous map f from D^i to X_i which maps $\partial D^i = S^{i-1}$ to X_{i-1} and extends the parameterization of the cell (a map from the interior of D^i to X_i).

By Proposition 5.2, if we are given a perfectly orientable plabic graph G, the parameterization $Meas_G$ of the cell Δ_G can be described as a map $(t_1, \dots, t_n) \mapsto [h_1(t_1, \dots, t_n), \dots, h_N(t_1, \dots, t_n)]$ where the h_i's are Laurent polynomials with positive coefficients. By Lemma 4.1, the map $Meas_G$ gives rise to a rational map

$m_G : X_G \to Gr_{kn}$ which is well-defined on $(X_G)_{\geq 0}$ (a closed ball). Furthermore, it is clear that the image of m_G on $(X_G)\geq 0$ lies in $(Gr_{kn})_{\geq 0}$.

Since the totally positive part of the toric variety X_G is dense in the non-negative part, and the interior gets mapped to the cell Δ_G, it follows that $(X_G)_{\geq 0}$ gets mapped to the closure of Δ_G.

And now by Lemma 5.1, it follows that the boundary of $(X_G)_{\geq 0}$ gets mapped to the $(i-1)$-skeleton of $(Gr_{kn})_{\geq 0}$. This completes the proof that the cell decomposition of $(Gr_{kn})_{\geq 0}$ is a CW complex. □

At the moment, the toric variety X_G is defined in a very abstract manner: we are to take the subtraction-free rational expressions described in Section 2, clear out denominators, and take the convex hull of the exponent vectors of the resulting polynomials. This defines a polytope, which we will denote $Q(G)$. In the following sections, we will define a very explicit polytope $P(G)$, which we will eventually show is equal to $Q(G)$.

It has been conjectured that the cell decomposition of $(Gr_{kn})_{\geq 0}$ is a regular CW complex which is homeomorphic to a ball. In particular, if a CW complex is regular then it follows that the Euler characteristic of the closure of each cell is 1.

In [22], Williams proved that the poset of cells of $(G/P)_{\geq 0}$ is thin and lexicographically shellable, hence in particular, Eulerian. This together with our result that the cell complex is actually a CW complex implies the following.

Corollary 5.5. *The Euler characteristic of the closure of each cell of* $(Gr_{kn})_{\geq 0}$ *is* 1.

Finally, we remark that there is a strong connection between total positivity and the cluster algebras of Fomin and Zelevinsky [5]. It is known that the coordinate ring of the Grassmannian can be given the structure of a cluster algebra [18], where the Plucker coordinates are the *cluster variables*. In our situation, the variables which index faces of a plabic graph are related to the variables in a particular cluster up to a certain twist map. Thus the polytope $Q(G)$ which defines the toric varieties used above can be thought of as the Minkowski sum of the Newton polytope of the cluster variables, expressed in terms of face variables. Indeed, in the case of the Grassmannian Gr_{2n}, it is possible to relate the combinatorics of perfect orientations of a given plabic graph to the combinatorial rule of Carroll and Price [2] giving explicit (positive) formulas for the cluster variables.

6. Fan of perfect orientations

In the next few sections we will give concrete combinatorial constructions of a certain polytope $P(G)$ associated to a plabic graph and its dual fan $\mathcal{F}_G$. (It will turn out that $P(G)$ is essentially the polytope $Q(G)$.) In this section we describe in detail the structure of $\mathcal{F}_G$.

Let $\text{Faces}(G) = \{f_0, f_1, \dots, f_{|\text{Faces}(G)|-1}\}$ denote the set of faces of G. We will work with face variables z_{f_i}, and define a fan which lives in the quotient $V := \mathbb{R}^{\text{Faces}(G)}/(\sum z_{f_i} = 0)$.

To each perfect orientation $\mathcal{O}$ associated with the graph, we define a cone $\sigma_{\mathcal{O}}$ by inequalities as follows. For an arbitrary path P from b_i to b_j without self-intersections, let $F' \subset \mathrm{Faces}(G)$ be the subset of faces which lie to the right of P. Then define the inequality $H_P := \sum_{f \in F'} z_f \geq 0$ (which we identify with the halfspace that it defines). For a clockwise (respectively, counterclockwise) circuit P, let $F' \subset \mathrm{Faces}(G)$ be the subset of faces which lie inside (respectively, outside) P. Then define the inequality $H_P := \sum_{f \in F'} z_f \geq 0$. We then define $\sigma_{\mathcal{O}}$ to be the image in V of the intersection $\cap H_P$ where P varies over all trails P.

Let $\mathcal{F}_G$ denote the collection of cones $\sigma_{\mathcal{O}}$ for all perfect orientations $\mathcal{O}$ of G together with all faces of such cones; we will show in the next section that $\mathcal{F}_G$ is a polytopal fan.

Lemma 6.1. *Fix a plabic graph G. One can obtain any perfect orientation $\mathcal{O}_2$ of G from another perfect orientation $\mathcal{O}_1$ by switching all directions in a disjoint collection of (directed) trails $S(\mathcal{O}_1, \mathcal{O}_2)$ in $\mathcal{O}_1$.*

Proof. Let E' denote the set of edges of G in which the orientations $\mathcal{O}_1$ and $\mathcal{O}_2$ disagree. Recall that in a perfect orientation, each black vertex v is incident to exactly one edge directed away from v, and each white vertex v is incident to exactly one edge directed towards v. Therefore every edge e in E' incident to some vertex v can be paired uniquely with another edge e' in E' which is also incident to v (note that at each vertex v of G there are either 0 or 2 incident edges which are in E'). This pairing induces a decomposition of E' into a disjoint union of (undirected) cycles and paths. Moreover, each such cycle or path is directed in both $\mathcal{O}_1$ and $\mathcal{O}_2$ (but of course in opposite directions). Let $S(\mathcal{O}_1, \mathcal{O}_2)$ denote this set of cycles and paths. □

7. The Fan is Polytopal

Let G be a plabic graph. The point of this section is to prove that $\mathcal{F}_G$ is a fan and is the normal fan to a certain polytope $P(G)$. It will turn out that $P(G)$ is essentially the polytope $Q(G)$ that we saw earlier; more precisely, $P(G)$ is equal to $-Q(G)$.

We begin with the following observation:

Proposition 7.1. *Let I be a finite set and let $\{C_i\}_{i \in I}$ be a collection of n-dimensional cones in $\mathbb{R}^n$ indexed by I. Let $\{v_i\}_{i \in I}$ be a collection of points in $\mathbb{R}^n$, indexed by $i \in I$, such that $C_{i_0}^{\vee}$ is the cone spanned by $(v_{i_0} - v_i)_{i \in I}$. Then the v_i are the vertices of $\mathrm{Hull}(\{v_i\}_{i \in I})$ and the C_i are the facets of the dual fan.*

Proof. Let $P = \mathrm{Hull}(\{v_i\}_{i \in I})$. For $i_0 \in I$, let F_{i_0} be the minimal face of P containing v_{i_0}. Then the (outer) tangent cone to P at F_{i_0}, which we write $-T_{F_{i_0}}(P)$, is spanned by $(v_{i_0} - v_i)_{i \in I}$. Now, by assumption, the cone dual to the span of $(v_{i_0} - v_i)_{i \in I}$ is C_{i_0} and, in particular, is n-dimensional. But C_{i_0} is n-dimensional if and only if the lineality space of $C_{i_0}^{\vee}$ is 0-dimensional. For any polytope P and any face F of P, the lineality space of $T_F(P)$ has dimension $\dim F$. Thus, we conclude that $\dim F_{i_0} = 0$, so v_{i_0} is a vertex of P.

Now that we know that the v_i are the vertices of P, then the facets of the dual fan to P are the cones dual to $\mathrm{Span}_{\mathbb{R}_{\geq 0}}(v_{i_0} - v_i)_{i \in I}$, as i_0 ranges through I. We have assumed that the dual cone to $\mathrm{Span}_{\mathbb{R}_{\geq 0}}(v_{i_0} - v_i)_{i \in I}$ is C_{i_0}. □

Note that this observation allows us to avoid checking ahead of time that the C_i form the facets of a complete fan, or that the v_i are the vertices of a polytope. Recall that the vector space $V^\vee$ to be the quotient of $\mathbb{R}^{\mathrm{Faces}(G)}$ by $(1, 1, \dots, 1)$. Let P be any path or cycle in G; we always consider paths and cycles to be given with an orientation, although not necessarily related to a perfect orientation of G. Define ρ_P to be the image in $V^\vee$ of the sum of e_f with f running over the faces on the right of P. Then, for any perfect orientation $\mathcal{O}$, the cone $\sigma_{\mathcal{O}}^\vee$ is spanned by the ρ_P.

We will describe how to associate, to every perfect orientation $\mathcal{O}$ of G, a point $v_{\mathcal{O}} \in V^\vee$. We will then need to check the following two claims:

Proposition 7.2. *Let $\mathcal{O}_0$ be any perfect orientation of G. Then the cone spanned by $v_{\mathcal{O}_0} - v_{\mathcal{O}}$, as $\mathcal{O}$ ranges over all perfect orientations, is the same as the cone spanned by ρ_P where P runs over all trails in G which agree with the perfect orientation $\mathcal{O}_0$.*

Proposition 7.3. *If $\mathcal{O}_0$ is any perfect orientation of G then the cone spanned by ρ_P, where P runs over all trails in G which agree with the perfect orientation $\mathcal{O}_0$, has dual of dimension* $\dim V^\vee = \#\mathrm{Faces}(G) - 1$.

Once we have checked this, we will know by Proposition 7.1 that the $\sigma_{\mathcal{O}}$ are the facets of the dual fan of $\mathrm{Hull}(v_{\mathcal{O}})$, where $\mathcal{O}$ runs over all perfect orientations of G. In particular, we will know that the fan $\mathcal{F}_G$ is polytopal.

For any two perfect orientations $\mathcal{O}_1$ and $\mathcal{O}_2$ of G, we define a subgraph $\delta_{\mathcal{O}_1\mathcal{O}_2}$ of G to be the subgraph consisting of those edges which are oriented differently in $\mathcal{O}_1$ and $\mathcal{O}_2$. We orient $\delta_{\mathcal{O}_1\mathcal{O}_2}$ so that each edge is oriented as in $\mathcal{O}_2$. Then by Lemma 6.1, $\delta_{\mathcal{O}_1\mathcal{O}_2}$ has net flow 0 at each internal vertex of G. There is a function $v_{\mathcal{O}_1,\mathcal{O}_2}$ on the faces of G such that, for two faces f_1 and f_2 separated by an edge e of G, we have $v_{\mathcal{O}_1,\mathcal{O}_2}(f_1) - v_{\mathcal{O}_1,\mathcal{O}_2}(f_2) = 0$ if $e \notin \delta_{\mathcal{O}_1\mathcal{O}_2}$; and we have $v_{\mathcal{O}_1,\mathcal{O}_2}(f_1) - v_{\mathcal{O}_1,\mathcal{O}_2}(f_2) = 1$ if the edge e is oriented to the left in $\delta_{\mathcal{O}_1\mathcal{O}_2}$ as we look from f_1 (up) to f_2 and we have $v_{\mathcal{O}_1\mathcal{O}_2}(f_1) - v_{\mathcal{O}_1\mathcal{O}_2}(f_2) = -1$ if the edge e is oriented to the right in $\delta_{\mathcal{O}_1\mathcal{O}_2}$ as we look from f_1 to f_2. The vector $v_{\mathcal{O}_1,\mathcal{O}_2}$ is unique up to adding the same constant to every face of G, so $v_{\mathcal{O}_1,\mathcal{O}_2}$ is well defined as an element of $V^\vee$.

Remark 7.4. In order to construct $v_{\mathcal{O}_1,\mathcal{O}_2}$ semiexplicitly, fix some face f_0 of G and normalize $v_{\mathcal{O}_1,\mathcal{O}_2}(f_0) = 0$. Then we can compute $v_{\mathcal{O}_1,\mathcal{O}_2}$ on every other face f of G by traveling along a path γ from f_0 to f and figuring out how $v_{\mathcal{O}_1\mathcal{O}_2}$ must change by looking at how we cross $\delta_{\mathcal{O}_1\mathcal{O}_2}$ as we travel along γ.

Note that we have

$$v_{\mathcal{O}_1,\mathcal{O}_3} = v_{\mathcal{O}_1,\mathcal{O}_2} + v_{\mathcal{O}_2,\mathcal{O}_3}. \tag{1}$$

This can be checked by, for each edge e of G, looking at the 8 possible orientations of e in $\mathcal{O}_1$, $\mathcal{O}_2$ and $\mathcal{O}_3$.

Finally we can define the polytopes that we are interested in.

Definition 7.5. Choose a perfect orientation $\mathcal{O}_{\text{base}}$. We set $v_{\mathcal{O}} := v_{\mathcal{O},\mathcal{O}_{\text{base}}}$ for any other perfect orientation $\mathcal{O}$. We define the polytope $P_{\mathcal{O}_{\text{base}}}(G)$ to be the convex hull in $V^{\vee}$ of the $v_{\mathcal{O}}$.

For this definition of $v_{\mathcal{O}}$, we will check Propositions 7.2 and 7.3.

Proof of Proposition 7.2. Let P be any oriented path or cycle in $\mathcal{O}_0$. Let $\mathcal{O}_1$ be the perfect orientation of G obtained by reversing P. Then

$$v_{\mathcal{O}_0} - v_{\mathcal{O}_1} = v_{\mathcal{O}_0,\mathcal{O}_1} = \rho_P$$

where the first equality is by equation (1) and the second is easy to check by hand. Thus, we see that ρ_P is in the span of the vectors of the form $v_{\mathcal{O}_0} - v_{\mathcal{O}}$ because it, in fact, is such a vector.

We now must check the reverse, that any vector $v_{\mathcal{O}_0} - v_{\mathcal{O}}$ can be written in the form $\sum \rho_P$ where P runs over some multiset of oriented paths and cycles in $\mathcal{O}_0$. But this is straightforward. By Lemma 6.1, we can decompose $\delta_{\mathcal{O}_0,\mathcal{O}}$ as an edge disjoint union of such paths and cycles; then $v_{\mathcal{O}_0,\mathcal{O}} = v_{\mathcal{O}_0} - v_{\mathcal{O}}$ is equal to the corresponding sum of ρ_P's. □

Proof of Proposition 7.3. Let K be the cone spanned by ρ_P as P runs over all trails that agree with the orientation $\mathcal{O}_0$. It is enough to find a point in the interior of $K^{\vee}$, in other words, it is enough to find a function λ, such that $\sum_{f \in \text{Faces}(G)} \lambda(f) = 0$ and λ has positive inner product with ρ_P for every trail P which agrees with the orientation $\mathcal{O}_0$.

Let e be an edge of G. As we look along e in the direction that e is oriented in $\mathcal{O}$, let f be the face of G to the right of e and let f' be the face of G to the left of e. Define η_e by $\eta_e(f) = 1$, $\eta_e(f') = -1$ and $\eta_f = 0$ otherwise. (If $f = f'$, we just set $\eta_e = 0$. Note that an edge e for which $f = f'$ can not occur in any trail of G.)

For any trail P of G consistent with the orientation $\mathcal{O}_0$, we have $\langle \eta_e, \rho_P \rangle = 1$ if $e \subseteq P$ and $\langle \eta_e, \rho_P \rangle = 0$ otherwise. Now, set $\lambda := \sum_{e \in \text{Edges}(G)} \eta_e$. Then $\langle \lambda, \rho_P \rangle$ is the length of P, for any trail P of G consistent with the orientation $\mathcal{O}_0$. In particular, $\langle \lambda, \rho_P \rangle$ is positive for any such P.

□

Now by combining Proposition 7.1, Proposition 7.2, and Proposition 7.3, we have proved the following.

Theorem 7.6. *The fan $\mathcal{F}_G$ is polytopal and has dimension* $\#\text{Faces}(G) - 1$*. Moreover, $\mathcal{F}_G$ is normal to the polytope $P_{\mathcal{O}_{\text{base}}}(G)$.*

Note that although $P_{\mathcal{O}_{\text{base}}}(G)$ is dependent on the choice of $\mathcal{O}_{\text{base}}$, changing $\mathcal{O}_{\text{base}}$ simply translates the polytope. Thus, when referring to any property of a polytope which is unaltered by translation, we may forget about our choice of perfect orientation. In this context, we will refer to the polytope as $P(G)$.

In Figure 4, we fix a plabic graph G corresponding to the cell of $(Gr_{24})_{\geq 0}$ such that the Plucker coordinates P_{12}, P_{13}, P_{14} are positive and all others are 0. We display the three perfect orientations, the associated fan $\mathcal{F}_G$, and the associated polytope $P(G)$. Note that since $P(G)$ is defined only up to translation, we set $C = 0$.

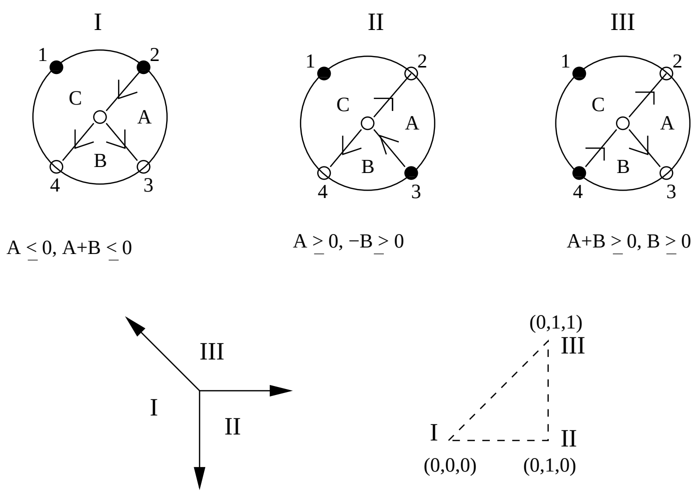

FIGURE 4.

In Figure 5, we fix a plabic graph G corresponding to the cell of $(Gr_{24})_{\geq 0}$ such that the Plucker coordinates $P_{12}, P_{13}, P_{24}, P_{34}$ are positive while P_{14} and P_{23} are 0. We display the four perfect orientations, the associated fan $\mathcal{F}_G$, and the associated polytope $P(G)$.

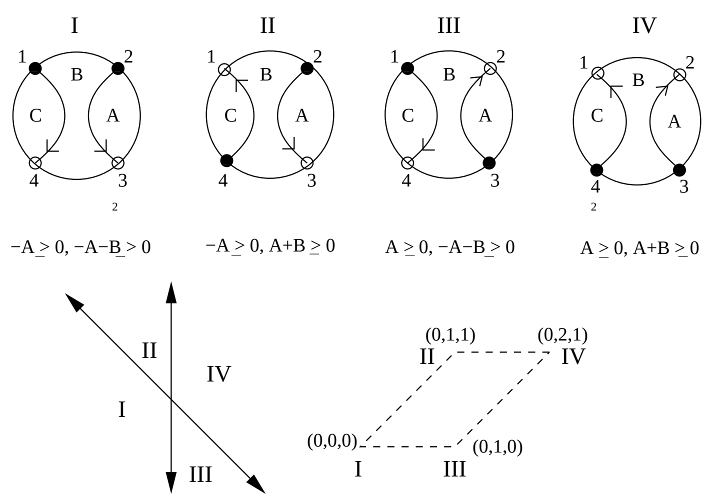

FIGURE 5.

When we constructed our parameterization of our cell of $(Gr_{kn})_{\geq 0}$ by a toric variety, our definition was very ineffective – we cleared out denominators from a collection of rational generating functions. We will know show that the polytope $Q(G)$ defining this toric variety is (essentially) the same as the extremely explicit polytope $P(G)$. If R is a polytope we define $-R$ in the obvious way: $-R$ is the set $\{(-v_1, \ldots, -v_m) \mid (v_1, \ldots, v_m) \in R\}$.

Proposition 7.7. *$Q(G)$ is a translate of $-P(G)$.*

We remark that the negative sign is the result of our desire to use the standard definition of a normal fan (from, for example, [23]) and the standard definition of the toric variety associated to a fan (from, for example, [6]). If we introduced a nonstandard sign in either of these settings, we could have $P(G)$ be a translate of $Q(G)$.

Proof. Fix an acyclic orientation $\mathcal{O}_{\text{base}}$ of G. Let I be the set of sources of G. Every vertex of $Q(G)$ comes from a set S of (directed) pairwise non-intersecting paths from I to another set J of boundary vertices of G. More precisely, the vertex (which we'll refer to as v_S) is the exponent vector of a monomial M in the face variables y_f of G, where the exponent of y_f in M is the number of paths in S that f lies to the right of. If f and f' are adjacent faces of G with no path of S passing between them then y_f and $y_{f'}$ occur with the same exponent in M. On the other hand, if f and f' have a path p of S passing between them with f on the left and f' on the right of p, then the exponent of $y_{f'}$ in M is one more than the exponent of y_f in M.

We now show that the vertex v_S of $Q(G)$ corresponds to a vertex of $P(G)$. Given the set S of non-intersecting paths as above, one can construct another perfect orientation $\mathcal{O}$ of G by taking all paths in S and reversing their orientation. The perfect orientation $\mathcal{O}$ corresponds to the vertex $v_{\mathcal{O}}$ defined earlier, and conversely for every $v_{\mathcal{O}}$ we get a uniquely defined set S of non-intersecting paths and hence a vertex v_S of $Q(G)$. Therefore the vertices of $P(G)$ and $Q(G)$ are in bijection.

The coordinates of $v_{\mathcal{O}}$ are defined (up to translation by $(1, 1, \ldots, 1)$) using the subgraph $\delta_{\mathcal{O},\mathcal{O}_{\text{base}}}$, whose edges are precisely those in the set of paths S. As before, if f and f' are adjacent faces of G with no path of S passing between them then the coordinate $v_{\mathcal{O},\mathcal{O}_{\text{base}}}(f) = v_{\mathcal{O},\mathcal{O}_{\text{base}}}(f')$. And if f and f' have a path p of S passing between them with f on the left and f' on the right then $v_{\mathcal{O},\mathcal{O}_{\text{base}}}(f') = v_{\mathcal{O},\mathcal{O}_{\text{base}}}(f) - 1$. Comparing this rule with the rule for the exponents of the y_f's, we see that $P(G)$ is (a translate of) $-Q(G)$.

□

We can use this description of our parameterization to give an alternate means of computing the M_{ij}. Choose any perfect orientation $\mathcal{O}$ of G, with source set I and let $i \in I$ and $j \notin I$. Then $M_{ij} = \pm p_{I \setminus \{i\} \cup \{j\}}/p_I$. We denote $I \setminus \{i\} \cup \{j\}$ by I'. Now, consider some other, acyclic, perfect orientation $\mathcal{O}_0$ of G with source set I_0. Then p_I and $p_{I'}$ can be computed as the generating functions for non-crossing directed path in $(G, \mathcal{O}_0)$ from $I_0 \setminus I$ to $I \setminus I_0$ and from $I_0 \setminus I'$ to $I' \setminus I_0$ respectively. It is a nice combinatorial problem to give a direct proof that the ratio of these (polynomial) generating functions is the rational generating function counting paths in $(G, \mathcal{O})$ from i to j, twisted by winding number.

Although we have given a global description of the map $X_G \to (Gr_{kn})$ in projective coordinates, we note that the local formulas given by the M_{ij} are quite convenient for analyzing the local behavior of the map. Let $\mathcal{O}$ be a perfect orientation of G, so $V_{\sigma_\mathcal{O}}$ is an affine open set of X_G. Let $z_\mathcal{O}$ be the torus fixed point of $V_{\sigma_\mathcal{O}}$. Then the formal generating function computing M_{ij} for the orientation $\mathcal{O}$ is a sum of monomials which are defined on $V_{\sigma_\mathcal{O}}$ and this generating function even converges in a neighborhood of $z_\mathcal{O}$.

8. The inequality description of $P(G)$

In this section we give an alternate description of the polytope $P_{\mathcal{O}_{\text{base}}}(G)$, through a set of defining inequalities. We note that the following set of inequalities contains redundancy; as we will show afterwards, half of the following inequalities cut out facets of $P_{\mathcal{O}_{\text{base}}}(G)$.

In what follows, we write points of $V^\vee$ as $\sum y_f e_f$, where the variables y_f are defined up to adding the same constant to all of them.

Proposition 8.1. *$P_{\mathcal{O}_{\text{base}}}(G)$ is defined by the following inequalities (two inequalities for each edge e of G): if f and f' are the two faces of G separated by e and if, in $\mathcal{O}_{\text{base}}$, the edge e is oriented to the left as we look from f into f', then $0 \leq y_f - y_{f'} \leq 1$. It follows that the $v_\mathcal{O}$ are precisely the lattice points in $P_{\mathcal{O}_{\text{base}}}(G)$; there are no interior lattice points.*

Proof. Let Q be the intersection of the required inequalities. Then it is clear that the $v_\mathcal{O}$ are precisely the lattice points in Q. We claim that all of the vertices of Q are integral, and thus that the $v_\mathcal{O}$ are the vertices of Q, implying that $Q = P_{\mathcal{O}_{\text{base}}}(G)$.

Now, the inequalities defining Q can be described as $A \cdot \boldsymbol{y} \leq \boldsymbol{b}$, where $\boldsymbol{y}$ is the vector of the y_f's. The matrix A is the transpose of the adjacency matrix of the dual graph of G, and is hence totally unimodular. If A is any totally unimodular matrix, and $\boldsymbol{b}$ any integer vector, then the polytope $A \cdot \boldsymbol{y} \leq \boldsymbol{b}$ has only integral vertices. (See [19, Theorem 19.1].) □

Proposition 8.2. *The facets of $P_{\mathcal{O}_{\text{base}}}(G)$ are in one-to-one correspondence with those edges e of G which separate two distinct faces of G. The bijection is as follows: Fix e and let (f, f') be the ordered pair of faces of G separated by e such that e is oriented to the left (in $\mathcal{O}_{\text{base}}$) as we look from f into f'. Then if e is directed from the black vertex to the white vertex, let $w_e := y_f - y_{f'} - 1$. Otherwise, let $w_e := y_{f'} - y_f$. The edge e corresponds to the facet inequality $w_e \leq 0$.*

Remark: If G is *reduced*, in the terminology of [14], then every edge of G separates two distinct faces of G.

Proof. If e does not separate two distinct faces of G, but rather separates the face f from itself, then the inequality associated to e is $y_f - y_f \geq 0$, which is trivially an equality on all of $P(G)$ and hence doesn't define a facet of $P(G)$. We call such an edge a *non-separating edge*.

We will show that all the inequalities from Proposition 8.1 other than those listed above are a consequence of the inequalities listed above and of the trivial inequalities associated to non-separating edges. Consider a white vertex v of G. Let $e_1, e_2, \ldots, e_m$ denote the edges incident to v (labeled cyclically), with e_1 being the unique edge which is directed towards v in $\mathcal{O}$. Let $f_1, f_2, \ldots, f_m$ denote the faces incident to v, labeled such that e_i separates the faces f_{i-1} and f_i (modulo m). Then the set S of inequalities given by Proposition 8.2 corresponding to the e_i's are $y_{f_1} - y_{f_m} \leq 1$ and also $y_{f_{i+1}} - y_{f_i} \leq 0$ for $1 \leq i \leq m-1$. Consider the complementary set S' of inequalities from Proposition 8.1 corresponding to the e_i's: $y_{f_m} - y_{f_1} \leq 0$ and also $y_{f_i} - y_{f_{i+1}} \leq 1$ for $1 \leq i \leq m-1$. It is now easy to see that by adding $m-1$ of the m inequalities in S at a time, we obtain the inequalities in S'.

It now remains to show that each of the inequalities in Proposition 8.2 actually defines a facet of the polytope. Fix an edge e and let Q be the convex hull of the set of vertices of $P_{\mathcal{O}_{\text{base}}}(G)$ which lie on the hyperplane defined by the inequality associated to e. We will show that Q has dimension #Faces$(G) - 2$. Let (f, f') be two faces of G separated by e as above. If e is directed from the black vertex to the white vertex in $\mathcal{O}_{\text{base}}$, then the vertices of $P_{\mathcal{O}_{\text{base}}}(G)$ on which $y_f - y_{f'} = 1$ correspond exactly to the perfect orientations of G such that e is directed in the opposite direction to $\mathcal{O}_{\text{base}}$, namely from the white vertex to the black vertex. Similarly, if e is directed from the white vertex to the black vertex in $\mathcal{O}_{\text{base}}$, then the vertices of $P_{\mathcal{O}_{\text{base}}}(G)$ on which $y_{f'} - y_f = 0$ also correspond exactly to the perfect orientations of G such that e is directed from the white vertex to the black vertex.

Now note that any perfect orientation of G in which the edge e is directed from the white vertex to the black vertex restricts to a perfect orientation of $G \setminus e$ (and vice-versa). Therefore the polytope Q is equal to the polytope $P(G \setminus e)$ (up to translation). Since we assumed that e separates two distinct faces of G, $G \setminus e$ has one fewer faces than G does. So the dimension of $P(G \setminus e)$ is #Faces$(G) - 2$, and we are done. ☐

Example 8.3. *In Figure 6 we give a plabic graph G corresponding to the top-dimensional cell of $(Gr_{24})_{\geq 0}$, along with its seven perfect orientations and the corresponding vertices of $P(G)$. (The coordinates of the vertices are drawn inside the faces of G.) If we fix $\mathcal{O}_{\text{base}}$ to be the perfect orientation of G listed as "I: Base" in the figure, then the facet inequalities of $P(G)$ are: $A+1 \geq B, D+1 \geq B, B \geq C, B \geq E, E \geq A, C \geq A, E \geq D$, and $C \geq D$. Since the polytope is defined only up to translation, we have set $E = 0$.*

Since we know that $\mathcal{F}_G$ is the normal fan of $P(G)$, we now get a description of the rays of $\mathcal{F}_G$.

Corollary 8.4. *The rays of $\mathcal{F}_G$ are in one-to-one correspondence with the edges e of G, as follows. Fix e and let (f, f') be the ordered pair of faces of G separated by e such that e is oriented to the left (in $\mathcal{O}_{\text{base}}$) as we look from f into f'. Then if e is directed from the black vertex to the white vertex, e corresponds to the ray r_e which is the image of $\mathbb{R}_{\geq 0}(e_f - e_{f'})$ in V. Otherwise, e corresponds to the ray r_e which is the image of $\mathbb{R}_{\geq 0}(e_{f'} - e_f)$ in V.*

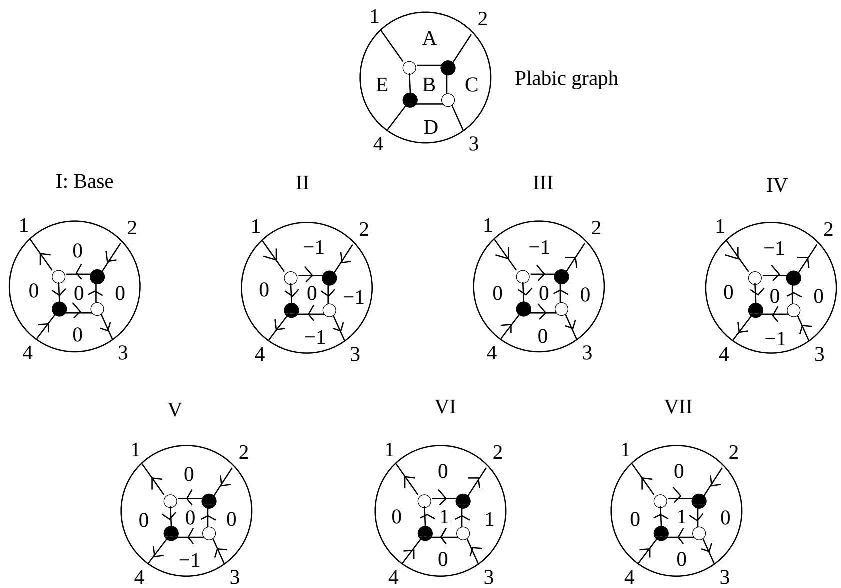

FIGURE 6.

Remark: Propp [15] has considered the general problem of, given a graph H (which need not be planar) and an element $\omega \in H^1(H, \mathbb{Z})$, determining all orientations of H such that the flow around every oriented cycle γ in H is $\omega(\gamma)$. To such a flow, Propp assigns a "height function" on the vertices of H. We can relate Propp's theory to ours in the following manner: let G be a plabic graph and let H be the dual graph to G. (We do not include a vertex dual to the outside of the disc in which G is embedded.) To any orientation $\mathcal{O}$ of G, we assign an orientation of H by rotating each edge 90° clockwise. Then $\mathcal{O}$ is a perfect orientation if and only if the corresponding orientation of G defines a certain class in $H^1(H, \mathbb{Z})$. Propp's height functions, in this setting, become the components of our $v_{\mathcal{O}}$. All the theory of this section can be written in the greater generality which Propp considers.

9. THE FACE LATTICE OF $P(G)$

We now consider the lattice of faces of $P(G)$, and give a description in terms of unions of matchings of G. This description is very similar to the description of the face lattice of the Birkhoff polytopes, as described by Billera and Sarangarajan [1].

In this section we assume that G is bipartite, which as noted in Section 2 is no loss of generality. As before, we fix a perfect orientation $\mathcal{O}_{\text{base}}$. Recall that a *planar-perfect matching* of a planar graph G is a collection of edges M of G in which every internal vertex is incident to precisely one edge in M (but boundary vertices are not necessarily incident to an edge in M.) Note that perfect orientations $\mathcal{O}$ of G are then in bijection with the planar-perfect matchings $M(\mathcal{O})$ of G.

If N is a face of $P(G)$, define $H(N)$ to be the union of $M(\mathcal{O})$ over all perfect orientations $\mathcal{O}$ indexing vertices of N. $H(N)$ then has the property that it has the same number of vertices as G and every edge of $H(N)$ is in some planar-perfect matching of G. Following [9], we call a subgraph H of G *elementary* if it contains every vertex of G and if every edge of H is used in some planar-perfect matching of H. Equivalently, the edges of H are obtained by taking a union of several planar-perfect matchings of G.

Conversely, if $H \subset G$ is an elementary graph, define

$$N(H) := \mathrm{conv}\{v_{\mathcal{O}} \mid M(\mathcal{O}) \subset H\}.$$

Let E be the set of edges in $G \setminus H$ and recall the definition of w_e from Proposition 8.2.

Proposition 9.1. *$N(H)$ is equal to $P_{\mathcal{O}_{\mathrm{base}}}(G) \cap \{w_e = 0 | e \in K\}$.*

Proof. Let P' denote $P_{\mathcal{O}_{\mathrm{base}}}(G) \cap \{w_e = 0 | e \in K\}$. Then P' is the intersection of $P(G)$ with some of its facets – i.e. P' is a face of $P(G)$ – and hence is the convex hull of some $v_{\mathcal{O}}$.

Consider any $v_{\mathcal{O}} \in P'$, and some $e \in K$. Let f and f' be the faces of G incident to e such that e is oriented to the left in $\mathcal{O}_{\mathrm{base}}$ as we look from f into f'. If, in $\mathcal{O}_{\mathrm{base}}$, the edge e is oriented from the white vertex to the black vertex then the fact that the coordinates of $v_{\mathcal{O}}$ satisfy $w_e := y_{f'} - y_f = 0$ implies that e is not in $\delta_{\mathcal{O},\mathcal{O}_{\mathrm{base}}}$ and hence e is oriented from white to black in $\mathcal{O}$. Therefore e is not an edge in the matching $M(\mathcal{O})$. On the other hand, if in $\mathcal{O}_{\mathrm{base}}$ the edge e is oriented from the black vertex to the white vertex then the fact that the coordinates of $v_{\mathcal{O}}$ satisfy $w_e := y_f - y_{f'} - 1 = 0$ implies that e *is* in $\delta_{\mathcal{O},\mathcal{O}_{\mathrm{base}}}$ and hence e is oriented from white to black in $\mathcal{O}$. Again, e is not an edge in the matching $M(\mathcal{O})$. Since this argument works for any $e \in K = G \setminus H$, we conclude that $M(\mathcal{O}) \subset H$ and hence $v_{\mathcal{O}} \in N(H)$.

Conversely, it is obvious that the coordinates of any $v_{\mathcal{O}}$ in $N(H)$ satisfies the equations $w_e = 0$ for $e \in K$. □

Since $N(H)$ is an intersection of facets of $P(G)$, it is a face of $P(G)$. Clearly, $H_1 \subseteq H_2$ implies $N(H_1) \subseteq N(H_2)$ and $N_1 \subseteq N_2$ implies $H(N_1) \subseteq H(N_2)$. We claim that the correspondences $N \mapsto H(N)$ and $H \mapsto N(H)$ are inverse.

Proposition 9.2. *If H is any elementary subgraph of G, then $H(N(H)) = H$. If N is any face of $P(G)$, then $N(H(N)) = N$.*

Proof. Let H be an elementary subgraph of G. From the definitions, we immediately have $H \supseteq H(N(H))$. To establish the converse containment, let e be any edge of H. Since H is elementary, there is a perfect matching M of H which uses the edge e. Let $\mathcal{O}$ be the corresponding perfect orientation of G; the vertex $\mathcal{O}$ is in $N(H)$ so $e \in H(N(H))$.

Now, let N be a face of $P(G)$. From the definitions, we immediately have $N \subseteq N(H(N))$. Suppose (for the sake of contradiction) that $v_{\mathcal{O}}$ is a vertex of $N(H(N))$ which is not in N. Then, by proposition 9.1, there is some edge $e \in K = G \setminus H(N)$ for which $w_e(v_{\mathcal{O}}) \neq 0$. But $w_e(v_{\mathcal{O}}) \neq 0$ means that $e \in M(\mathcal{O})$, so $M(\mathcal{O}) \not\subseteq H(N)$, contradicting our choice of $v_{\mathcal{O}} \in N(H(N))$. □

This yields a description of the face lattice of $P(G)$: it is the lattice of elementary subgraphs of G, ordered by inclusion of edges. (Plus a formal element $\hat{0}$, corresponding to the minimal element of the face lattice.) The join operation is defined by the union of the edge sets and the meet of two elementary graphs is defined to be the join of all elementary graphs less than both (which may be empty). Alternatively, we may describe the meet of $G_1 \cap G_2$ as the subgraph of $G_1 \cap G_2$ obtained by deleting all edges of $G_1 \cap G_2$ not used in any planar-perfect matching of $G_1 \cap G_2$. So we have proved the following.

Theorem 9.3. *The face lattice of $P(G)$ is isomorphic to the lattice of all elementary subgraphs of G, ordered by inclusion.*

The minimal nonempty elementary subgraphs of G are the matchings, corresponding to vertices of $P(G)$.

Corollary 9.4. *Consider a cell Δ_G of $(Gr_{kn})_{\geq 0}$ parameterized by a plabic graph G. For any cell Δ_H in the closure of Δ_G, the corresponding polytope $P(H)$ is a face of $P(G)$.*

Proof. By [14, Theorem 18.3], every cell in the closure of Δ_G can be parameterized using a plabic graph H which is obtained by deleting some edges from G. H is perfectly orientable and hence is an elementary subgraph of G. Therefore by Theorem 9.3, the polytope $P(H)$ is a face of $P(G)$. □

In Figure 7, we have drawn the edge graph of the four-dimensional polytope $P(G)$, where G is the plabic graph from Example 8.3. This time we have depicted the vertices with matchings instead of perfect orientations; note that the Roman numerals indexing matchings in Figure 7 agree with the Roman numeral indexing perfect orientations in Figure 6. Additionally, we have labeled each vertex with the source set of its perfect orientation.

Lemma 9.5. *Consider two (planar-perfect) matchings M_1 and M_2 of G. Let $H = M_1 \cup M_2$. If some vertex v in G is incident to two distinct edges e_1 and e_2 such that $e_1 \in M_1$ and $e_2 \in M_2$ then in fact there is a path through v comprised of edges in H (alternating between M_1 and M_2) which either begins and ends on the boundary of the disk or forms a closed loop in the interior of the disk. Furthermore, the paths in H are vertex-disjoint.*

Proof. This is just Lemma 6.1, translated into the language of matchings. □

Proposition 9.6. *Consider two vertices v_1 and v_2 of $P(G)$ corresponding to the two planar-perfect matchings M_1 and M_2. Let $H = M_1 \cup M_2$ and let r be the number of regions into which the edges of H divide the disk. Then the smallest face of $P(G)$ containing v_1 and v_2 is a cube of dimension $r - 1$.*

Proof. Since H divides the disk into r regions, by Lemma 9.5, it must consist of precisely $r - 1$ disjoint closed paths or cycles (plus some isolated edges). There are exactly 2^{r-1} planar-perfect matchings of G which are contained in H because for each closed path or cycle there are two ways of picking a subset of edges to be in

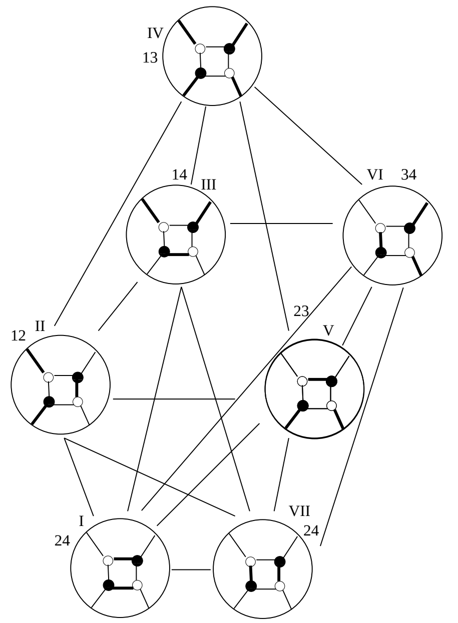

FIGURE 7.

the matching. These 2^{r-1} planar-perfect matchings correspond to the vertices of an $r-1$-dimensional cube. $\square$

Proposition 9.7. *Let N be a face of $P(G)$, and let r be the number of regions into which $H(N)$ divides the disk in which G is embedded. Then* $\dim N = r - 1$.

Proof. Recall that $H(N)$ is the subgraph of G which is the union of the planar-perfect matchings of G corresponding to vertices of N. Since $H(N)$ has a planar-perfect matching, it is perfectly orientable, so we can view it as a perfectly orientable plabic graph in its own right. It is then easy to see that the polytope N is equal to the polytope $P(H(N))$, up to translation, and hence $\dim N = \dim P(H(N)) = \#\text{Faces}(H(N)) - 1$. But this is equal to $r-1$, as desired. $\square$

As a special case of the preceding propositions, we get the following.

Remark 9.8. Let N be a face of $P(G)$ and let r be the number of regions into which the edges of $H(N)$ divide the disk in which G is embedded. Then N is an edge of $P(G)$ if and only if $r = 2$. Equivalently, two vertices $v_{\mathcal{O}_1}$ and $v_{\mathcal{O}_2}$ of $P(G)$ form an edge if and only if $\mathcal{O}_2$ can be obtained from $\mathcal{O}_1$ by switching the orientation along a directed trail in $\mathcal{O}_1$.

Recall that the *Birkhoff polytope* B_n is the convex hull of the $n!$ points in $\mathbb{R}^{n^2}$ $X(\pi) : \pi \in S_n$ where $X(\pi)_{ij}$ is equal to 1 if $\pi(i) = j$ and is equal to 0 otherwise. It is well-known that B_n is an $(n-1)^2$ dimensional polytope, whose face lattice of B_n is isomorphic to the lattice of all elementary subgraphs of the complete bipartite graph $K_{n,n}$ ordered by inclusion [1]. Our polytopes $P(G)$ can be thought of as a kind of analog of the Birkhoff polytopes for planar graphs embedded in a disk.

10. Connection with matroid polytopes

Every perfectly orientable plabic graph encodes a realizable *positroid*, that is, an oriented matroid in which all orientations are positive. The bases of the positroid associated to a plabic graph G of type (k, n) are precisely the k-element subsets $I \subset [n]$ which occur as source sets of perfect orientations of G. This is easy to see, as each perfect orientation of G gives rise to a parametrization of the cell Δ_G of $(Gr_{kn})_{\geq 0}$ in which the Plucker coordinate corresponding to I is 1. Furthermore, if one takes a (directed) path in a perfect orientation $\mathcal{O}$ and switches the orientation of each of its edges, this encodes a basis exchange.

Given this close connection of perfectly orientable plabic graphs to positroids, it is natural to ask whether there is a connection between our polytopes $P(G)$ and matroid polytopes. It turns out that there is a map Ψ from $P(G)$ to the matroid polytope corresponding to the (unoriented) positroid associated to G.

Let M be a matroid of rank k on the ground set $[n]$. Recall that the *matroid polytope* $Q(M)$ is the convex hull of the vectors $\{e(J) \mid J \text{ is a basis of } M\}$ where $e(J)$ is the $0-1$ vector in $\mathbb{R}^n$ whose ith coordinate is 1 if $i \in J$ and is 0 otherwise [8]. The vertices are in one-to-one correspondence with bases of M. This polytope lies in the hyperplane $x_1 + \cdots + x_n = 0$ and has dimension $n - 1$.

In contrast, if G is a plabic graph of type (k, n), one can associate to each vertex $v_{\mathcal{O}}$ of $P(G)$ a basis of the corresponding positroid (the source set corresponding to $\mathcal{O}$), but this is in general a many-to-one correspondence from vertices of $P(G)$ to bases of the positroid. Another way in which $P(G)$ differs from $Q(M_G)$ is that $P(G)$ has the same dimension as the corresponding cell Δ_G of $(Gr_{kn})_{\geq 0}$ (which is also the rank of the associated positroid). This dimension ranges from 0 to $k(n-k)$, and is equal to $\#\text{Faces}(G) - 1$.

Let M_G denote the matroid corresponding to G, i.e. the underlying matroid of the positroid associated to G. We will define a map Ψ from $P(G)$ to $Q(M_G)$.

Fix a perfect orientation $\mathcal{O}_{\text{base}}$ of G. Let S be the set of edges attached to boundary vertices in the plabic graph which are oriented into the disk, i.e. attached to sources. Let $f_1, \ldots, f_n$ denote the n faces bordering the boundary of the disk. Recall that we represent points in $P_{\mathcal{O}_{\text{base}}}(G) \subset V^{\vee}$ in the form $\sum y_f e_f$, where f ranges over the faces of G and the variables y_f are defined up to adding the same constant to all of them.

Definition 10.1. We define $\Psi : P(G) \to Q(M_G)$ by mapping $\sum y_f e_f$ to $(a_1, \ldots a_n) \in \mathbb{R}^n$, where $a_i = y_{f_i} - y_{f_{i+1}}$ if the edge between f_i and f_{i+1} is *not* in S and $a_i = y_{f_i} - y_{f_{i+1}} + 1$ if the edge between f_i and f_{i+1} is in S. (Here indices are regarded modulo n).

Proposition 10.2. *Ψ is well-defined and is a projection of $P(G)$ onto $Q(M_G)$. Furthermore, if $v_{\mathcal{O}}$ is a vertex of $P(G)$ corresponding to the basis $J \subset [n]$, then $\Psi(v_{\mathcal{O}})$ is the vertex of $Q(M_G)$ corresponding to the same basis J.*

Proof. Clearly Ψ is well-defined as a map on $V^{\vee}$. Since there are k boundary vertices of G which are sources with respect to $\mathcal{O}_{\text{base}}$, $|S| = k$ and hence $\sum a_i = k$. Therefore Ψ indeed maps $P(G)$ to $Q(M_G)$. Finally, it is easy to check that if $\mathcal{O}$

is a perfect orientation of G such that $J \subset [n]$ is the set of sources, then $\Psi(v_{\mathcal{O}}) = e(J)$. Therefore Ψ is a surjection which takes vertices to vertices as specified in the proposition. $\square$

Example 10.3. *Consider the plabic graph G from Figure 6. This corresponds to the positroid of rank two on the ground set $[4]$ such that all subsets of size 2 are independent. The edge graph of the four-dimensional polytope $P(G)$ is shown in Figure 7, and each vertex is labeled with the basis it corresponds to. The matroid polytope of this matroid is the (three-dimensional) octahedron with six vertices corresponding to the two-element subsets of $[4]$. Under the map Ψ, the vertex of $P(G)$ corresponding to the two-element subset ij gets mapped to the vertex of the octahedron whose ith and jth coordinates are 1 (and whose other coordinates are 0).*

11. Numerology of the polytopes $P(G)$

In this section we give some statistics about a few of the polytopes $P(G)$. Our computations were made with the help of the software Polymake [7].

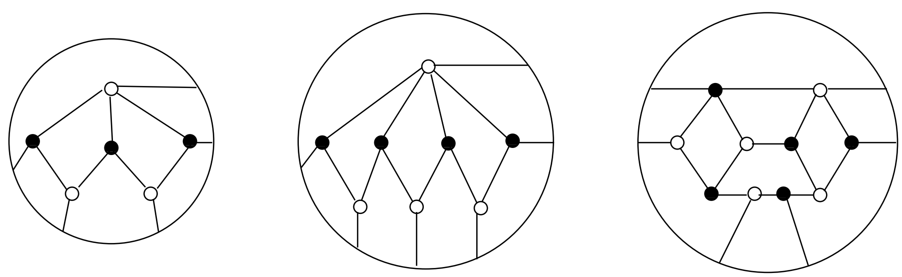

G25 G26 G36

Figure 8.

Let $G24$ denote the plabic graph shown in Figure 6, and let $G25$, $G26$, and $G36$ denote the plabic graphs shown in Figures 8. These four plabic graphs give parameterizations of the top cells of $(Gr_{24})_{\geq 0}, (Gr_{25})_{\geq 0}, (Gr_{26})_{\geq 0}$, and $(Gr_{36})_{\geq 0}$, respectively.

Then the f-vector of the polytope $P(G24)$ is $(7, 17, 18, 8)$.

The f-vector of $P(G25)$ is $(14, 59, 111, 106, 52, 12)$.

The f-vector of $P(G26)$ is $(25, 158, 440, 664, 590, 315, 98, 16)$.

The f-vector of $P(G36)$ is $(42, 353, 1212, 2207, 2368, 1557, 627, 149, 19)$.

The Ehrhart series of $P(G24)$ is $\frac{1+2t+t^2}{(1-t)^5}$.

The Ehrhart series of $P(G25)$ is $\frac{1+7t+12t^2+4t^3}{(1-t)^7}$.

The Ehrhart series of $P(G26)$ is $\frac{1+16t+64t^2+68t^3 15t^4}{(1-t)^9}$.

The volume of $P(G24)$ is $\frac{1}{6} = \frac{4}{4!}$ and so the degree of the corresponding toric variety X_{G24} is 4.

The volume of $P(G25)$ is $\frac{1}{30} = \frac{24}{6!}$ and so the degree of X_{G25} is 24.

The volume of $P(G26)$ is $\frac{41}{10080} = \frac{164}{8!}$ and so the degree of X_{G26} is 164.

The volume of $P(G36)$ is $\frac{781}{181440} = \frac{1562}{9!}$ and so the degree of X_{G36} is 1562.

Note that in general there is more than one plabic graph giving a parameterization of a given cell. But even if two plabic graphs G and G' correspond to the same cell, in general we have $P(G) \neq P(G')$. For example, the plabic graph in Figure 9 gives a parameterization of the top cell of $(Gr_{26})_{\geq 0}$. Let us refer to this graph as $\hat{G}26$. However, $P(\hat{G}26) \neq P(G26)$: the f-vector of $P(\hat{G}26)$ is $(26, 165, 460, 694, 615, 326, 100, 16)$.

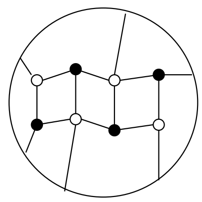

Figure 9.

References

[1] L. Billera, A. Sarangarajan, the Combinatorics of permutation polytopes, DIMACS Series in Discrete Mathematics and Theoretical Computer Science 24, 1996.

[2] G. Carroll, G. Price, Two new combinatorial models for the Ptolemy recurrence, preprint, 2003.

[3] D. Cox, What is a toric variety? Topics in algebraic geometry and geometric modeling, 203–223, Contemp. math., 334, Amer. Math. Soc., Providence, RI, 2003.

[4] S. Fomin, M. Shapiro, Stratified spaces formed by totally positive varieties, Mich. Math. J. 48 (2000), 253–270.

[5] S. Fomin, A. Zelevinsky, Cluster Algebras I: Foundations, J. Amer. Math. Soc. 22 (2002), 497–527.

[6] W. Fulton, Introduction to toric varieties, Annals of Mathematics Studies, 131, Princeton University Press, Princeton, NY, 1993.

[7] E. Gawrilow, M. Joswig, polymake: a framework for analyzing convex polytopes and simplicial complexes. http://www.math.tu-berlin.de/polymake, 1997–2003. Version 2.0, with contributions by Thilo Schroder and Nikolaus Witte.

[8] I. Gelfand, R. Goresky, R. MacPherson, V. Serganova, Combinatorial geometries, convex polyhedra, and Schubert cells, Adv. in Math. 63 (1987), no. 3, 301–316.

[9] L. Lovasz, M. Plummer, Matching Theory, Elsevier Science Pub. Co. New York, 1986.

[10] G. Lusztig, Introduction to total positivity, in Positivity in Lie theory: open problems, ed. J. Hilgert, J.D. Lawson, K.H. Neeb, E.B. Vinberg, de Gruyter Berlin, 1998, 133–145.

[11] G. Lusztig, Total positivity in partial flag manifolds, Representation Theory, 2 (1998) 70–78.

[12] G. Lusztig, Total positivity in reductive groups, in: Lie theory and geometry: in honor of Bertram Kostant, Progress in Mathematics 123, Birkhauser, 1994.

[13] R. Marsh, K. Rietsch, Parametrizations of flag varieties, Representation Theory, 8 (2004).

[14] A. Postnikov, Total positivity, Grassmannians, and networks. http://front.math.ucdavis.edu/math.CO/0609764

[15] J. Propp, Lattice Structure for Orientations of Graphs, preprint. http://front.math.ucdavis.edu/math.CO/0209005

[16] K. Rietsch, Total positivity and real flag varieties, Ph.D. Dissertation, MIT, 1998.

[17] K. Rietsch, Closure relations for totally non-negative cells in G/P, Math. Res. Lett. 13 (2006), no. 5-6, 775–786.

[18] J. Scott, Grassmannians and cluster algebras, Proceedings of the London Math. Soc. (3) 92, 2006, no. 2, 345 – 380.

[19] A. Schrijver, Theory of linear and integer programming, Wiley-Interscience series in discrete mathematics and optimization, John Wiley and Sons, New York, 1986.

[20] F. Sottile, Toric ideals, real toric varieties, and the moment map, in Topics in Algebraic Geometry and Geometric Modeling, Contemp. Math. 334, 2003, 225 – 240.

[21] L. Williams, Enumeration of totally positive Grassmann cells, Adv. Math., 190 (2), 2005, 319–342.

[22] L. Williams, Shelling totally non-negative flag varieties, to appear in J. Reine Angew. Math. (Crelle's Journal). http://front.math.ucdavis.edu/math.RT/0509129

[23] G. Ziegler, Lectures on Polytopes, Graduate Texts in Mathematics, 152, Springer-Verlag, New York, 1995.

Massachusetts Institute of Technology
E-mail address: apost@math.mit.edu

University of Michigan
E-mail address: speyer@umich.edu

Harvard University
E-mail address: lauren@math.harvard.edu