

Weak separation and plabic graphs

Suho Oh, Alexander Postnikov and David E. Speyer

ABSTRACT

Leclerc and Zelevinsky described quasicommuting families of quantum minors in terms of a certain combinatorial condition, called *weak separation*. They conjectured that all inclusion-maximal weakly separated collections of minors have the same cardinality, and that they can be related to each other by a sequence of mutations.

Postnikov studied *total positivity* on the Grassmannian. He described a stratification of the totally non-negative Grassmannian into *positroid strata*, and constructed their parameterization using *plabic graphs*.

In this paper, we link the study of weak separation to plabic graphs. We extend the notion of weak separation to positroids. We generalize the conjectures of Leclerc and Zelevinsky, and related ones of Scott, and prove them. We show that the maximal weakly separated collections in a positroid are in bijective correspondence with the plabic graphs. This correspondence allows us to use the combinatorial techniques of positroids and plabic graphs to prove the (generalized) purity and mutation connectedness conjectures.

1. Introduction

Leclerc and Zelevinsky [11], in their study of quasicommuting families of quantum minors, introduced the following notion of weakly separated sets:

Let I and J be two subsets of $[n] := \{1, 2, \dots, n\}$. Leclerc and Zelevinsky [11] defined I and J to be *weakly separated* if either

- (1) $|I| \leq |J|$ and $I \setminus J$ can be partitioned as $I_1 \sqcup I_2$ such that $I_1 \prec J \setminus I \prec I_2$ or
- (2) $|J| \leq |I|$ and $J \setminus I$ can be partitioned as $J_1 \sqcup J_2$ such that $J_1 \prec I \setminus J \prec J_2$,

where $A \prec B$ indicates that every element of A is less than every element of B .

Leclerc and Zelevinsky proved that, for any collection \mathcal{C} of pairwise weakly separated subsets of $[n]$, one has $|\mathcal{C}| \leq \binom{n}{2} + n + 1$. Moreover, they made the following *Purity Conjecture*.

CONJECTURE 1.1 ([11]). If \mathcal{C} is a collection of subsets of $[n]$, each of which are pairwise weakly separated from each other, and such that \mathcal{C} is not contained in any larger collection with this property, then $|\mathcal{C}| = \binom{n}{2} + n + 1$.

The above notion of weak separation is related to the study of the Plücker coordinates on the flag manifold, see [11]. In the context of Plücker coordinates on the Grassmannian, it is natural to study weak separation of k element subsets of $[n]$, for fixed n and k . Observe that, when $|I| = |J|$, the definition of weak separation becomes invariant under cyclic shifts of $[n]$. Indeed, in this case I and J are weakly separated if and only if after an appropriate cyclic shift $(I \setminus J) \prec (J \setminus I)$.

Received 15 May 2013; revised 3 March 2014; published online 10 February 2015.

2010 *Mathematics Subject Classification* 05E15 (primary), 14M15 (secondary).

Alexander Postnikov was supported in part by NSF CAREER Award DMS-0504629; David E. Speyer was supported in part by a Clay Research Fellowship.

Scott [14] proved that, for any collection $\mathcal{C} \subset \binom{[n]}{k}$ of pairwise weakly separated k element subsets on $[n]$, one has $|\mathcal{C}| \leq k(n-k) + 1$. Here $\binom{[n]}{k}$ denotes the set of k element subsets of $[n]$. Moreover, Scott made the following conjecture.

CONJECTURE 1.2 ([14]). If $\mathcal{C} \subset \binom{[n]}{k}$ is a collection of k element subsets of $[n]$, each of which are pairwise weakly separated from each other, and such that \mathcal{C} is not contained in any larger collection with this property, then $|\mathcal{C}| = k(n-k) + 1$.

We present a stronger statement, which implies both these conjectures.

In [13], a stratification of the totally non-negative Grassmannian $\text{Gr}^{tnn}(k, n)$, called the *positroid stratification*, was investigated. The cells of this stratification can be labeled by many different kinds of combinatorial objects. The three of these objects which we use are

- (1) Grassmann necklaces, which are certain sequences (I_1, I_2, \dots, I_n) in $\binom{[n]}{k}$;
- (2) positroids, which are certain subsets of $\binom{[n]}{k}$;
- (3) decorated permutations, which are permutations of $[n]$ with fixed points colored in two colors.

There is a length function ℓ on decorated permutations, related to the length function on the affine symmetric group.

One of our main results is the following.

THEOREM 1.3. *Let \mathcal{I} , \mathcal{M} and π be a Grassmann necklace, positroid and decorated permutation corresponding to each other. Let $\mathcal{C} \subseteq \binom{[n]}{k}$ be a collection of pairwise weakly separated sets such that $\mathcal{I} \subseteq \mathcal{C} \subseteq \mathcal{M}$, and such that \mathcal{C} is not contained in any larger collection with this property. Then $|\mathcal{C}| = 1 + k(n-k) - \ell(\pi)$.*

We call collections \mathcal{C} satisfying the conditions of this theorem *maximal weakly separated collections inside \mathcal{M}* .

This result implies Conjectures 1.1 and 1.2, and also the w -chamber conjecture of Leclerc and Zelevinsky [11].

From our perspective, Scott's Conjecture 1.2 is especially natural, because it refers to the largest cell in $G^{tnn}(k, n)$. The reader may find it useful to first consider the proofs of the main results in that special case.

We also prove the following result on mutation connectedness, of which special cases were conjectured in [11, 14].

THEOREM 1.4. *Fix a positroid \mathcal{M} . Any two maximal weakly separated collections inside \mathcal{M} can be obtained from each other by a sequence of mutations of the following form: $\mathcal{C} \mapsto (\mathcal{C} \setminus \{Sac\}) \cup \{Sbd\}$, assuming that, for some cyclically ordered elements a, b, c, d in $[n] \setminus S$, \mathcal{C} contains Sab, Sbc, Scd, Sda and Sac . (Here Sab is a shorthand for $S \cup \{a, b\}$, etc.)*

Our main tool is the technology of plabic graphs developed in [13]. Plabic graphs are variants of wiring diagrams of decompositions in the symmetric group. As with wiring diagrams, there is a notion of a reduced plabic graph, and the faces of a reduced plabic graph are labeled with subsets of $\binom{[n]}{k}$.

We show that maximal weakly separated collections and reduced plabic graphs are in a bijective correspondence.

THEOREM 1.5. Fix a positroid \mathcal{M} and the corresponding Grassmann necklace \mathcal{I} . For a reduced plabic graph G associated with \mathcal{M} , let $\mathcal{F}(G) \subset \binom{[n]}{k}$ be the collection of labels of faces of G . Then the map $G \mapsto \mathcal{F}(G)$ is a bijection between reduced plabic graphs for the positroid \mathcal{M} and maximal weakly separated collections \mathcal{C} such that $\mathcal{I} \subseteq \mathcal{C} \subseteq \mathcal{M}$.

We also establish the following consequence for the theory of cluster algebras, in Section 7.

THEOREM 1.6. Let \mathcal{C} be a subset of $\binom{[n]}{k}$. The following are equivalent:

- (i) the set of Plücker coordinates $\{p_I\}_{I \in \mathcal{C}}$ is a cluster in the cluster algebra structure on the coordinate ring of $G(k, n)$;
- (ii) \mathcal{C} is a maximal weakly separated collection;
- (iii) \mathcal{C} is the collection of face labels of a reduced plabic graph for the uniform matroid $\binom{[n]}{k}$.

Readers not interested in cluster algebras may skip this section.

Theorems 1.3 and 1.4 follow from the correspondence in Theorem 1.5 and the properties of plabic graphs from [13].

Our main new tool is a construction we term ‘plabic tilings’. Maximal plabic tilings are dual to plabic graphs, and non-maximal plabic tilings correspond to non-maximal weakly separated collections.

1.1. Parallel work

As we were completing this work, we learned that Danilov, Karzanov and Koshevoy succeeded in proving Leclerc and Zelevinsky’s conjectures. Similar to our proof, their proof relies on certain planar diagrams. Their diagrams, which are called generalized tilings, are closely related to the plabic graphs which play a central role in this paper, see [6]. We see several differences between our work and theirs. We study weakly separated collections of all positroids, which is more general than considering weakly separated collections of w -chamber sets, see Section 12. Moreover, our constructions respect the dihedral symmetry inherent in the definitions of positroids and weak separation, which Danilov, Karzanov and Koshevoy’s constructions do not. Finally, our bicolored surfaces give a natural object to assign to a weakly separated collection which is not maximal.

1.2. Notion of weak separation

From our perspective, it is most natural to work with weak separation as defined by Scott [14]. In Section 3, we review Scott’s definition, and this is the meaning of ‘weak separation’ for the rest of the paper. We refer to Leclerc and Zelevinsky’s definition as *LZ weak separation*. We explain the relation between the notions in Section 12. All of the statements in Section 12 (although not all their proofs) can be understood as soon as the reader has read the definitional material in Sections 3 and 4.

2. Notation

Throughout the paper, we use the following notation. Let $[n]$ denote $\{1, \dots, n\}$ and let $\binom{[n]}{k}$ denote the set of k element subsets of $[n]$. We generally consider $[n]$ as cyclically ordered. We write $<_i$ for the cyclically shifted linear order on $[n]$:

$$i <_i i + 1 <_i i + 2 <_i \dots <_i n <_i 1 <_i \dots <_i i - 1.$$

We say that i_1, i_2, \dots, i_r in $[n]$ are *cyclically ordered* if $i_1 <_i i_2 <_i \dots <_i i_r$ for some $i \in [n]$.

We write (a, b) for the open cyclic interval from a to b . In other words, the set of i such that a, i, b is cyclically ordered. We write $[a, b]$ for the closed cyclic interval, $[a, b] = (a, b) \cup \{a, b\}$, and use similar notations for half open intervals.

If S is a subset of $[n]$ and a an element of $[n]$, then we may abbreviate $S \cup \{a\}$ and $S \setminus \{a\}$ by Sa and $S \setminus a$.

In this paper, we need to deal with three levels of objects: elements of $[n]$, subsets of $[n]$ and collections of subsets of $[n]$. For clarity, we denote these by lower case letters, capital letters and calligraphic letters, respectively.

The use of the notation $I \setminus J$ does not imply $J \subseteq I$.

3. Weakly separated collections

In this section, we define weak separation for collections of k element subsets and discuss the k subset analog of Leclerc and Zelevinsky's conjectures. The relation of this approach to the original definitions and conjectures from [11] be discussed in Section 12.

Let us fix two non-negative integers $k \leq n$.

DEFINITION 3.1. For two k element subsets I and J of $[n]$, we say that I and J are *weakly separated*, if there do *not* exist a, b, a', b' , cyclically ordered, with $a, a' \in I \setminus J$ and $b, b' \in J \setminus I$.

Geometrically, I and J are weakly separated if and only if there exists a chord separating the sets $I \setminus J$ and $J \setminus I$ drawn on a circle.

We write $I \parallel J$ to indicate that I and J are weakly separated.

We call a subset \mathcal{C} of $\binom{[n]}{k}$ a *collection*. We define a *weakly separated collection* to be a collection $\mathcal{C} \subset \binom{[n]}{k}$ such that, for any I and J in \mathcal{C} , the sets I and J are weakly separated.

We define a *maximal weakly separated collection* to be a weakly separated collection which is not contained in any other weakly separated collection.

Following Leclerc and Zelevinsky [11], Scott observed the following claim.

PROPOSITION 3.2 ([14], cf. [11]). Let $S \in \binom{[n]}{k-2}$ and let a, b, c, d be cyclically ordered elements of $[n] \setminus S$. Suppose that a maximal weakly separated collection \mathcal{C} contains Sab, Sbc, Scd, Sda and Sac . Then $\mathcal{C}' := (\mathcal{C} \setminus \{Sac\}) \cup \{Sbd\}$ is also a maximal weakly separated collection.

We define \mathcal{C} and \mathcal{C}' to be *mutations* of each other if they are linked as in Proposition 3.2. We prove the following claim, conjectured by Scott.

THEOREM 3.3. Every maximal weakly separated collection of $\binom{[n]}{k}$ has cardinality $k(n-k)+1$. Any two maximal weakly separated collections are linked by a sequence of mutations.

4. Weakly separated collections in positroids

While our results are purely combinatorial, they are motivated by constructions in algebraic geometry. Specifically, [13] introduced the positroid stratification of the Grassmannian (see also [10]). We prove a version of Theorem 3.3 for every cell in this stratification. Theorem 3.3 itself corresponds to the case of the largest cell.

There are several combinatorial objects which can be used to index the cells of this stratification. We use three of these: *Grassmann necklaces*, *decorated permutations* and *positroids*. See [13] for more details.

DEFINITION 4.1 ([13, Definition 16.1]). A *Grassmann necklace* is a sequence $\mathcal{I} = (I_1, \dots, I_n)$ of k -element subsets of $[n]$ such that, for $i \in [n]$, the set I_{i+1} contains $I_i \setminus \{i\}$. (Here the indices are taken modulo n .) If $i \notin I_i$, then we should have $I_{i+1} = I_i$.

In other words, I_{i+1} is obtained from I_i by deleting i and adding another element, or $I_{i+1} = I_i$. Note that, in the latter case $I_{i+1} = I_i$, either i belongs to all elements I_j of the Grassmann necklace, or i does not belong to all elements I_j of the necklace.

Here is an example of a Grassmann necklace: $I_1 = \{1, 2, 4\}$, $I_2 = \{2, 4, 5\}$, $I_3 = \{3, 4, 5\}$, $I_4 = \{4, 5, 2\}$, $I_5 = \{5, 1, 2\}$. As in this example, we usually sort the elements of I_i according to the order $<_i$ when emphasizing the role of the set I_i as an element of a Grassmann necklace. We often abbreviate examples like this as $(124, 245, 345, 452, 512)$.

Recall the linear order $<_i$ on $[n]$. We extend $<_i$ to k element sets, as follows. For $I = \{i_1, \dots, i_k\}$ and $J = \{j_1, \dots, j_k\}$ with $i_1 <_i i_2 \dots <_i i_k$ and $j_1 <_i j_2 \dots <_i j_k$, define the partial order

$$I \leq_i J \text{ if and only if } i_1 \leq_i j_1, \dots, i_k \leq_i j_k.$$

In other words, \leq_i is the cyclically shifted termwise partial order on $\binom{[n]}{k}$.

DEFINITION 4.2. Given a Grassmann necklace $\mathcal{I} = (I_1, \dots, I_n)$, define the *positroid* $\mathcal{M}_{\mathcal{I}}$ to be

$$\mathcal{M}_{\mathcal{I}} := \{J \in \binom{[n]}{k} \mid I_i \leq_i J \text{ for all } i \in [n]\}.$$

This turns out to be a special kind of matroid.

The term ‘positroid’ is an abbreviation for ‘totally positive matroid’. These are exactly matroids that can be represented by totally positive matrices, see [13].

DEFINITION 4.3. Fix a Grassmann necklace $\mathcal{I} = (I_1, \dots, I_n)$, with corresponding positroid $\mathcal{M}_{\mathcal{I}}$. Then \mathcal{C} is called a *weakly separated collection inside $\mathcal{M}_{\mathcal{I}}$* if \mathcal{C} is a weakly separated collection and $\mathcal{I} \subseteq \mathcal{C} \subseteq \mathcal{M}_{\mathcal{I}}$. We call \mathcal{C} a *maximal weakly separated collection inside $\mathcal{M}_{\mathcal{I}}$* if it is maximal among weakly separated collections inside $\mathcal{M}_{\mathcal{I}}$.

EXAMPLE 4.4. A more naive definition would be simply to define a weakly separated collection inside $\mathcal{M}_{\mathcal{I}}$ to be a weakly separated collection contained in $\mathcal{M}_{\mathcal{I}}$. It is easy to see that this definition would make Theorem 1.5 false: The collection face labels of any reduced plabic graph for \mathcal{I} contains \mathcal{I} . Take $\mathcal{I}_1 := \{13, 23, 34, 14\}$, then $\mathcal{C}_{bad} := \{24, 23, 34, 14\}$ is maximal in $\mathcal{M}_{\mathcal{I}_1}$ in this naive sense, but \mathcal{C}_{bad} does not contain \mathcal{I}_1 . The naive definition would also make Theorem 1.3 false. Take $(k, n) = (3, 6)$ and $\mathcal{I}_2 = \{125, 235, 356, 456, 561, 612\}$. The matroid $\mathcal{M}_{\mathcal{I}_2}$ is the set of those $X \in \binom{[6]}{3}$ so that X contains at most one of $\{3, 4\}$ and at least one of $\{5, 6\}$. We have $\ell(\mathcal{I}_2) = 3$ so $k(n - k) + 1 - \ell(\mathcal{I}_2) = 7$. Consider $\mathcal{C}_{small} := \{136, 146, 145, 126, 156, 456\}$. The collection \mathcal{C}_{small} is maximal in $\mathcal{M}_{\mathcal{I}_2}$ in the naive sense (when checking this by hand, it is helpful to note that the last three elements of \mathcal{C}_{small} are weakly separated from every set in $\binom{[6]}{3}$), yet we have $\#\mathcal{C}_{small} = 6 < 7$.

We have not shown yet that $\mathcal{I} \subseteq \mathcal{M}_{\mathcal{I}}$, or that \mathcal{I} is weakly separated. So there might be no weakly separated collections inside $\mathcal{M}_{\mathcal{I}}$. We remedy this now.

LEMMA 4.5. For any Grassmann necklace \mathcal{I} , we have $\mathcal{I} \subseteq \mathcal{M}_{\mathcal{I}}$, and \mathcal{I} is weakly separated.

Proof. For every i and j in $[n]$, we must show that $\mathcal{I}_i \leq_i \mathcal{I}_j$ and $\mathcal{I}_i \parallel \mathcal{I}_j$.

By the definition, I_{r+1} is either obtained from I_r by deleting r and adding another element, or else $I_{r+1} = I_r$. As we do the changes $I_1 \rightarrow I_2 \rightarrow \dots \rightarrow I_n \rightarrow I_1$, we delete each $r \in [n]$ at most once (in the transformation $I_r \rightarrow I_{r+1}$). This implies that we add each r at most once.

Let us show that $I_j \setminus I_i \subseteq [j, i)$. Suppose that this is not true and there exists $r \in (I_j \setminus I_i) \cap [i, j)$. Note that $I_{r+1} \neq I_r$. Otherwise, r belongs to all elements of the Grassmann necklace, or r does not belong to all elements of the necklace. Consider the sequence of changes $I_i \rightarrow I_{i+1} \rightarrow \dots \rightarrow I_r \rightarrow I_{r+1} \rightarrow \dots \rightarrow I_j$. We should have $r \notin I_i$, $r \in I_r$, $r \notin I_{r+1}$, $r \in I_j$. Thus r should be added twice, as we go from I_i to I_r and as we go from I_{r+1} to I_j . We obtain a contradiction.

Thus $I_j \setminus I_i \subseteq [j, i)$ and, similarly, $I_i \setminus I_j \subseteq [i, j)$. We conclude that $\mathcal{I}_i \leq_i \mathcal{I}_j$ and $\mathcal{I}_i \parallel \mathcal{I}_j$, as desired. \square

Our main theorem says, in part, that all maximal weakly separated collections inside $\mathcal{M}_{\mathcal{I}}$ have the same size. To describe this cardinality, we define decorated permutations.

DEFINITION 4.6 ([13, Definition 13.3]). A decorated permutation $\pi^\cdot = (\pi, \text{col})$ is a permutation $\pi \in S_n$ together with a coloring function col from the set of fixed points $\{i \mid \pi(i) = i\}$ to $\{1, -1\}$.

There is a simple bijection between decorated permutations and Grassmann necklaces. To go from a Grassmann necklace \mathcal{I} to a decorated permutation $\pi^\cdot = (\pi, \text{col})$, we set $\pi(i) = j$ whenever $I_{i+1} = (I_i \setminus \{i\}) \cup \{j\}$ for $i \neq j$. If $i \notin I_i = I_{i+1}$, then $\pi(i) = i$ is a fixed point of color $\text{col}(i) = 1$. Finally, if $i \in I_i = I_{i+1}$, then $\pi(i) = i$ is a fixed point of color $\text{col}(i) = -1$.

To go from a decorated permutation $\pi^\cdot = (\pi, \text{col})$ to a Grassmann necklace \mathcal{I} , we set

$$I_i = \{j \in [n] \mid j <_i \pi^{-1}(j) \text{ or } (\pi(j) = j \text{ and } \text{col}(j) = -1)\}.$$

For example, the decorated permutation $\pi^\cdot = (\pi, \text{col})$ with $\pi = 81425736$ and $\text{col}(5) = 1$ corresponds to the Grassmann necklace (I_1, \dots, I_8) with $I_1 = \{1, 2, 3, 6\}$, $I_2 = \{2, 3, 6, 8\}$, $I_3 = \{3, 6, 8, 1\}$, $I_4 = \{4, 6, 8, 1\}$, $I_5 = \{6, 8, 1, 2\}$, $I_6 = \{6, 8, 1, 2\}$, $I_7 = \{7, 8, 1, 2\}$, $I_8 = \{8, 1, 2, 3\}$. For the fact that this is a bijection between Grassmann necklaces and decorated permutations, see [13, Theorem 17.1].

DEFINITION 4.7 ([13, Section 17]). For $i, j \in [n]$, we say that $\{i, j\}$ forms an *alignment* in π if $i, j, \pi(j), \pi(i)$ are cyclically ordered (and all distinct). The *length* $\ell(\pi^\cdot)$ the number of alignments in π . We define $\ell(\mathcal{I})$ to be $\ell(\pi^\cdot)$, where π^\cdot is the associated decorated permutation of \mathcal{I} .

Geometrically, $\ell(\mathcal{I})$ is the codimension of the positroid variety associated to \mathcal{I} . So $k(n - k) - \ell(\mathcal{I})$ is the dimension of the positroid variety inside $G(k, n)$, and $k(n - k) - \ell(\mathcal{I}) + 1$ is the dimension of the affine cone over it.

We now state our result for an arbitrary Grassmann cell.

THEOREM 4.8. Fix any Grassmann necklace \mathcal{I} . Every maximal weakly separated collection inside $\mathcal{M}_{\mathcal{I}}$ has cardinality $k(n - k) - \ell(\mathcal{I}) + 1$. Any two maximal weakly separated collections inside $\mathcal{M}_{\mathcal{I}}$ are linked by a sequence of mutations.

As a particular case, let $I_i = \{i, i + 1, \dots, i + k - 1\} \subset [n]$. (The entries of I_i are taken modulo n .) Then $\mathcal{M}_{\mathcal{I}} = \binom{[n]}{k}$ and the above theorem becomes Theorem 3.3. In Section 12, we explain how this theorem also incorporates the conjectures of Leclerc and Zelevinsky.

5. *Decomposition into connected components*

Let \mathcal{I} be a Grassmann necklace, let π^\cdot be the corresponding decorated permutation, and let \mathcal{M} be the corresponding positroid $\mathcal{M}_{\mathcal{I}}$.

The connected components of π^\cdot , \mathcal{I} and \mathcal{M} are certain decorated permutations, Grassmann necklaces and positroids, respectively. In each case, the ground set is no longer $[n]$ but rather a subset of $[n]$, which inherits its circular order from $[n]$. We now define these objects.

EXAMPLE 5.1. The arrows in the left-hand side of Figure 10 form a decorated permutation with $(k, n) = (5, 10)$, namely the permutation $(019)(27)(38)(654)$. The connected components are 019, 2378 and 456. The five element set drawn on the faces of this graph form a weakly separated collection for this π . The reader who is curious about the bipartite graph in Figure 10, and about the right-hand side of the figure, finds the explanations in Sections 6 and 9.

DEFINITION 5.2. Let $[n] = S_1 \sqcup S_2 \sqcup \cdots \sqcup S_r$ be a partition of $[n]$ into disjoint subsets. We say that $[n]$ is *non-crossing* if there do not exist (a, b, c, d) , which are circularly ordered, and $i \neq j$ with $\{a, c\} \subseteq S_i$ and $\{b, d\} \subseteq S_j$.

See, for example, [16] for background on non-crossing partitions. The following is obvious from the definition.

PROPOSITION 5.3. *The common refinement of two non-crossing partitions is non-crossing.*

Therefore, the following definition makes sense.

DEFINITION 5.4. Let π^\cdot be a decorated permutation. Let $[n] = \bigsqcup S_i$ be the finest non-crossing partition of $[n]$ such that, if $i \in S_j$, then $\pi(i) \in S_j$.

Let $\pi_{(j)}^\cdot$ be the restriction of π^\cdot to the set S_j , and let $\mathcal{I}_{(j)}$ be the associated Grassmann necklace on the ground set S_j , for $j = 1, \dots, r$.

We call $\pi_{(j)}^\cdot$ the *connected components* of π^\cdot , and $\mathcal{I}_{(j)}$ the *connected components* of \mathcal{I} .

We say that π^\cdot and \mathcal{I} are connected if they have exactly one connected component.

LEMMA 5.5. *The decorated permutation π^\cdot is disconnected if and only if there are two circular intervals $[i, j]$ and $[j, i]$ such that π takes $[i, j]$ and $[j, i]$ to themselves.*

Proof. If such intervals exist, then the pair $[n] = [i, j] \sqcup [j, i]$ is a non-crossing partition preserved by π . So there is a non-trivial non-crossing partition preserved by π and π is not connected. Conversely, any non-trivial non-crossing permutation can be coarsened to a pair of intervals of this form so, if π is disconnected, then there is a pair of intervals of this form. \square

Note that each fixed point of π^\cdot (of either color) forms a connected component.

LEMMA 5.6. *A Grassmann necklace $\mathcal{I} = (I_1, \dots, I_n)$ is connected if and only if the sets I_1, \dots, I_n are all distinct.*

Proof. If π^{\cdot} is disconnected, then let $[i, j)$ and $[j, i)$ be as in Lemma 5.5. As we change from I_i to I_{i+1} to I_{i+2} to \dots to I_j , each element of $[i, j)$ is removed once and is added back in once. So $I_i = I_j$.

Conversely, suppose that $I_i = I_j$. As we change from I_i to I_{i+1} to I_{i+2} and so forth, up to I_j , each element of $[i, j)$ is removed once. To have $I_i = I_j$, each element of $[i, j)$ must be added back in once. So π takes $[i, j)$ to itself. \square

In this section, we explain how to reduce computations about positroids to the connected case.

For the rest of this section, suppose that $I_i = I_j$ for some $i \neq j$. Set $I^1 := [i, j) \cap I_i$ and $I^2 = [j, i) \cap I_i$; set $k_1 = |I^1|$ and $k_2 = |I^2|$. We set $n^1 = |[i, j)|$ and $n^2 = |[j, i)|$.

PROPOSITION 5.7. *For every $J \in \mathcal{M}$, we have $|J \cap [i, j)| = k^1$ and $|(J \cap [j, i))| = k^2$.*

Proof. Since I_i is the \leq_i minimal element of \mathcal{M} , we have

$$|[i, j) \cap J| \leq |[i, j) \cap I_i| = k^1$$

for all $J \in \mathcal{M}$. But also, similarly,

$$|[j, i) \cap J| \leq |[j, i) \cap I_j| = k^2.$$

Adding these inequalities together, we see that

$$|J| = |[i, j) \cap J| + |[j, i) \cap J| \leq k^1 + k^2 = k.$$

But, in fact, $|J| = k$. So we have equality at every step of the process. In particular, $|J \cap [i, j)| = k^1$ and $|J \cap [j, i)| = k^2$. \square

PROPOSITION 5.8. *The matroid \mathcal{M} is a direct sum of two matroids \mathcal{M}^1 and \mathcal{M}^2 , supported on the ground sets $[i, j)$ and $[j, i)$, having ranks k^1 and k^2 . In other words, there are matroids \mathcal{M}^1 and \mathcal{M}^2 such that J is in \mathcal{M} if and only if $J \cap [i, j)$ is in \mathcal{M}^1 and $J \cap [j, i)$ is in \mathcal{M}^2 .*

For every $J \in \mathcal{M}$, write $J^1 := J \cap [i, j)$ and $J^2 := J \cap [j, i)$.

Proof. This is an immediate corollary of Proposition 5.7. [5, Theorem 7.6.4] shows that this rank condition implies that \mathcal{M} is the direct sum of $\mathcal{M}|_{[i, j)}$ and $\mathcal{M}|_{[j, i)}$; [5, Proposition 7.6.1] interprets this in terms of bases of \mathcal{M} . \square

PROPOSITION 5.9. *For $k \in [i, j)$, the set I_k is of the form $J \cup I^2$ for some $J \in \mathcal{M}^1$. For $k \in [j, i)$, the set I_k is of the form $I^1 \cup J$ for some $J \in \mathcal{M}^2$.*

Proof. Consider the case that $k \in [i, j)$, the other case is similar. Recall that I_k is the \leq_k minimal element of \mathcal{M} . Since $\mathcal{M} = \mathcal{M}^1 \oplus \mathcal{M}^2$, we know that $I_k = J^1 \cup J^2$, where J^r is the \leq_k minimal element of \mathcal{M}^r . But, on $[j, i)$, the orders \leq_i and \leq_k coincide, so J^2 is the \leq_i minimal element of \mathcal{M}^2 , namely I^2 . \square

View $[i, j)$ as circularly ordered. Let \mathcal{I}^1 denote the circularly ordered sequence $(I_i^1, I_{i+1}^1, \dots, I_{j-1}^1)$.

PROPOSITION 5.10. *The sequence of sets \mathcal{I}^1 is a Grassmann necklace on the ground set $[i, j]$, and \mathcal{M}^1 is the associated positroid.*

Proof. For $k \in [i, j]$, we know that $I_{k+1} \supseteq I_k \setminus \{k\}$. Since $I_k = I_k^1 \cup I^2$, and $I_{k+1} = I_{k+1}^1 \cup I^2$, we have $I_{k+1}^1 \supseteq I_k^1 \setminus \{k\}$. This is the definition of a Grassmann necklace. (Note that we have used the condition $I_i = I_j$ to cover the boundary case $k = j - 1$.)

Now, we show that \mathcal{M} is the associated positroid. Consider any $J \in \binom{[i, j]}{k^1}$. If $J \in \mathcal{M}^1$, then $J \cup I^2 \in \mathcal{M}$, so $J \cup I^2 \succcurlyeq_k I_k$ for every $k \in [n]$. This immediately implies that $J \succcurlyeq_k I_k^1$ for every $k \in [i, j]$, so J is in the positroid $\mathcal{M}_{\mathcal{I}^1}$.

Conversely, suppose that J is in the positroid $\mathcal{M}_{\mathcal{I}^1}$. We wish to show that $J \cup I^2$ is in \mathcal{M} . Reversing the argument of the previous paragraph shows that $J \cup I^2 \succcurlyeq_k I_k$ for all $k \in [i, j]$. For $k \in [j, i]$, we know that $I^2 \succcurlyeq_k I_k^2$ and, since \succcurlyeq_k and \succcurlyeq_i coincide on $[i, j]$, we know that $J \succcurlyeq_k I^1$. So $J \cup I^2 \succcurlyeq_k I_k$ for k in $[j, i]$ as well. So $J \in \mathcal{M}$. \square

It is easy to check the following proposition.

PROPOSITION 5.11. *With the above definitions, $k(n - k) - \ell(\mathcal{I}) = k_1(n_1 - k_1) - \ell(\mathcal{I}^1) + k_2(n_2 - k_2) - \ell(\mathcal{I}^2)$.*

We now study weakly separated collections in \mathcal{M} .

LEMMA 5.12. *Let $J = J^1 \cup J^2 \in \mathcal{M}$. If J is weakly separated from $I^1 \cup I^2$, then either $I^1 = J^1$ or $I^2 = J^2$.*

Proof. Suppose, to the contrary, that $J^1 \neq I^1$ and $J^2 \neq I^2$. Since $I^1 <_i J^1$, there are a and $b \in [i, j]$, with $i \leq_i a <_i b$, such that $a \in I^1 \setminus J^1$ and $b \in J^1 \setminus I^1$. Similarly, there are c and $d \in [j, i]$, with $j \leq_j c <_j d$, such that $c \in I^2 \setminus J^2$ and $d \in J^2 \setminus I^2$. Then a and c are in $I^1 \cup I^2 \setminus J^1 \cup J^2$, while b and d are in $J^1 \cup J^2 \setminus I^1 \cup I^2$. So $I^1 \cup I^2$ and J are not weakly separated. \square

PROPOSITION 5.13. *If \mathcal{C} is a weakly separated collection in \mathcal{M} , then there are weakly separated collections \mathcal{C}^1 and \mathcal{C}^2 in \mathcal{M}^1 and \mathcal{M}^2 such that*

$$\mathcal{C} = \{J \cup I^2 : J \in \mathcal{C}^1\} \cup \{I^1 \cup J : J \in \mathcal{C}^2\}.$$

Conversely, if \mathcal{C}^1 and \mathcal{C}^2 are weakly separated collections in \mathcal{M}^1 and \mathcal{M}^2 , then the above formula defines a weakly separated collection in \mathcal{M} . The collection \mathcal{C} is maximal if and only if \mathcal{C}^1 and \mathcal{C}^2 are.

Proof. First, suppose that \mathcal{C} is a weakly separated collection in \mathcal{M} . Since $\mathcal{I} \subset \mathcal{C}$, we have $I^1 \cup I^2 \in \mathcal{C}$. By Lemma 5.12, every $J \in \mathcal{C}$ is either of the form $J^1 \cup I^2$, or $I^1 \cup J^2$. Let \mathcal{C}^r be the collection of all sets J^r for which $J^r \cup I^{3-r}$ is in \mathcal{C} . The condition that \mathcal{C} is weakly separated implies that \mathcal{C}^r is; the condition that $\mathcal{I} \subseteq \mathcal{C} \subseteq \mathcal{M}$ implies that $\mathcal{I}^r \subseteq \mathcal{C}^r \subseteq \mathcal{M}^r$. So \mathcal{C}^r is a weakly separated collection in \mathcal{M}^r and it is clear that \mathcal{C} is built from \mathcal{C}^1 and \mathcal{C}^2 in the indicated manner.

Conversely, it is easy to check that, if \mathcal{C}^1 and \mathcal{C}^2 are weakly separated collections in \mathcal{M}^1 and \mathcal{M}^2 , then the above formula gives a weakly separated collection in \mathcal{M} .

Finally, if $\mathcal{C} \subsetneq \mathcal{C}'$ with \mathcal{C}' a weakly separated collection in \mathcal{M} , then either $\mathcal{C}^1 \subsetneq (\mathcal{C}')^1$ or $\mathcal{C}^2 \subsetneq (\mathcal{C}')^2$. So, if \mathcal{C} is not maximal, either \mathcal{C}^1 or \mathcal{C}^2 is not. The converse is similar. \square

In summary, the Grassmann necklace \mathcal{I} can be described in terms of smaller necklaces \mathcal{I}^1 and \mathcal{I}^2 ; the positroid \mathcal{M} can be described in terms of smaller positroids \mathcal{M}^1 and \mathcal{M}^2 ; weakly separated collections \mathcal{C} in \mathcal{M} can be described in terms of weakly separated collections in \mathcal{M}^1 and \mathcal{M}^2 .

REMARK 5.14. The connected components S_i are the connected components of \mathcal{M} in the sense of matroid theory. See [1] or [7].

6. Plabic graphs

Plabic graphs were introduced in [13]. We use these graphs to represent weakly separated collections, and the properties of these graphs are a major key in our proof of the conjecture. For more details on plabic graphs, see [13].

DEFINITION 6.1. A *planar bicolored graph*, or simply a *plabic graph* is a planar undirected graph G drawn inside a disk. The vertices on the boundary are called boundary vertices, and are labeled in clockwise order by $[n]$. All vertices in the graph are colored either white or black.

Let G be a plabic graph in the disc D . We draw n -directed paths, called *strands*, within the disc D , each starting from and ending at a boundary vertex of G .

DEFINITION 6.2. The construction in this definition is depicted in Figure 1; the numeric labels in Figure 1 are explained. The strands are drawn as follows: For each edge of G , draw two strand segments. If the ends of the segment are the same color, then the two strands should be parallel to the edge without crossing, and should run in opposite directions. If the two ends are different colors, then the two strands should cross, with one running toward each endpoint. As we discuss in the below remark, for most purposes, we can reduce to the case that G is bipartite, so the latter case is the important one. Around each vertex, connect up the ends of the strands so that they turn right at each black vertex and left at each white vertex. We have n strands leading from ∂G to itself, plus possibly some loops in the interior of G .

REMARK 6.3. Suppose that u and v are two vertices of G , joined by an edge, which have the same color. Let G' be the graph formed by contracting the edge (u, v) to a single vertex w , and coloring w the same color as u and v are colored. Then the strands of G and G' have the same connectivity. See Figure 2 for a depiction of how contracting an edge leaves strand topology unchanged. For this reason, we can usually reduce any question of interest to the case where G is bipartite. We do not restrict in this paper to bipartite graphs for two reasons: (1) We want to be compatible with the definitions in [13], which does not make this restriction and (2) The description of the square move, (M1) below, would be significantly more complicated.

A plabic graph is called *reduced* [13, Section 13] if the following holds.

- (1) The strands cannot be closed loops in the interior of the graph.
- (2) No strand passes through itself. The only exception is that we allow simple loops that start and end at a boundary vertex i .




FIGURE 1. Labeling the faces of a plabic graph.




FIGURE 2. The effect of replacing G by G' , as in Remark 6.3.

(3) For any two strands α and β , if α and β have two common vertices A and B , then one strand, say α , is directed from A to B , and the other strand β is directed from B to A . (That is the crossings of α and β occur in opposite orders in the two strands.)

REMARK 6.4. Strands, and the reducedness condition, occurs in the physics-inspired literature on quivers. In this literature, strands are called zig-zag paths. Graphs which we call ‘reduced’ are said to ‘obey condition Z ’ in [3, Section 8] and are called ‘marginally geometrically consistent’ in [4, Section 3.4].

The strand which ends at the boundary vertex i is called strand i .

DEFINITION 6.5 ([13, Section 13]). For a reduced plabic graph G , let $\pi_G \in S_n$ be the permutation such that the strand that starts at the boundary vertex i ends at the boundary




FIGURE 3. (M1) Square move.




FIGURE 4. (M2) Unicolored edge contraction.

vertex $\pi_G(i)$. A fixed point $\pi_G(i) = i$ corresponds to simple loop at the boundary vertex i . We color a fixed point i of π_G as follows: $\text{col}(i) = 1$ if the corresponding loop is counter-clockwise; and $\text{col}(i) = -1$ if the loop is clockwise. In this way, we assign the decorated strand permutation $\pi_G^\bullet = (\pi_G, \text{col})$ to each reduced plabic graph G .

We label the faces of a reduced plabic graph with subsets of $[n]$; this construction was first published in [14]. By condition (2), each strand divides the disk into two parts. For each face F we label that F with the set of those $i \in [n]$ such that F lies to the left of strand i . See Figure 1 for an example. So given a plabic graph G , we define $\mathcal{F}(G)$ as the set of labels that occur on each face of that graph.

When we pass from one face F of G to a neighboring one F' , we cross two strands. For one of these strands, F lies on its left and F' on the right; for the other F lie on the right and F' on the left. So every face is labeled by the same number of strands as every other. We define this number to be the rank of the graph; it eventually plays the role of k .

The following claim establishes a correspondence between maximal weakly separated collections in a positroid and reduced plabic graphs. It describes maximal weakly separated collections as labeled sets $\mathcal{F}(G)$ of reduced plabic graphs.

THEOREM 6.6. *For a decorated permutation π^\bullet and the corresponding Grassmann necklace $\mathcal{I} = \mathcal{I}(\pi^\bullet)$, a collection \mathcal{C} is a maximal weakly separated collection inside the positroid $\mathcal{M}_{\mathcal{I}}$ if and only if it has the form $\mathcal{C} = \mathcal{F}(G)$ for a reduced plabic graph with strand permutation π^\bullet .*

In particular, a maximal weakly separated collection \mathcal{C} in $\binom{[n]}{k}$ has the form $\mathcal{C} = \mathcal{F}(G)$ for a reduced plabic graph G with strand permutation $w^\bullet = [k + 1, k + 2, \dots, n, 1, 2, \dots, k]$.

We define the boundary face at i to be the face touching the part of the disk between boundary vertices $i - 1$ and i in clockwise order. Let I_i be the label of that face. At boundary point i , the strand i comes in and the strand $\pi(i)$ leaves. So I_{i+1} is obtained from I_i by deleting i and adding in $\pi(i)$; we deduce that (I_1, \dots, I_n) is the Grassmann necklace $\mathcal{I}(\pi)$.

We now describe how to see mutations in the context of plabic graphs. We have following three moves on plabic graphs.

(M1) Pick a square with vertices alternating in colors, such that all vertices have degree 3. We can switch the colors of all the vertices. See Figure 3.

(M2) For two adjoint vertices of the same color, we can contract them into one vertex. See Figure 4.

(M3) We can insert or remove a vertex inside any edge. See Figure 5.




FIGURE 5. (M3) Vertex removal.

The moves do not change the associated decorated permutation of the plabic graph, and do not change whether or not the graph is reduced. The power of these moves is reflected in Theorem 6.7.

THEOREM 6.7 ([13, Theorem 13.4]). *Let G and G' be two reduced plabic graphs with the same number of boundary vertices. Then the following claims are equivalent:*

- (i) G can be obtained from G' by moves (M1)–(M3);
- (ii) these two graphs have the same decorated strand permutation $\pi_G = \pi_{G'}$.

Moves (M2) and (M3) do not change $\mathcal{F}(G)$, while move (M1) changes $\mathcal{F}(G)$ by a mutation, as shown in Figure 8. So, once we prove Theorem 6.6, we establish that all maximal weakly separated collections within $\mathcal{M}(\mathcal{I})$ are connected to each other by mutations.

Note also that the moves (M1)–(M3) do not change the number of faces of the plabic graph. So all reduced plabic graphs for a given decorated permutation have the same number of faces. This number was calculated in [13].

THEOREM 6.8. *Let G be a reduced plabic graph with decorated permutation π . Then G has $k(n - k) - \ell(\pi) + 1$ faces.*

Proof. By [13, Proposition 17.10], $k(n - k) - \ell(\pi)$ is the dimension of $S_{\mathcal{M}}^{tnn}$, a manifold whose definition we do not need to know. By [13, Theorem 12.7], $S_{\mathcal{M}}^{tnn}$ is isomorphic to $\mathbb{R}^{|F(G)|-1}$. \square

Thus, proving Theorem 6.6 establishes all parts of 4.8, thus proving the conjectures of Scott and of Leclerc and Zelevinsky.

7. Consequences for cluster algebras

This section is not cited in the rest of the paper.

From the beginnings of the theory of cluster algebras, it has been expected that the coordinate ring of the Grassmannian, in its Plücker embedding, would be a cluster algebra, and that the Plücker coordinates would be cluster variables. This was verified by Scott [15]. Moreover, Scott showed that, given any reduced alternating strand diagram, the face labels of that diagram form a cluster in the cluster structure for $G(k, n)$. With the tools of this paper, we can establish the converse statement.

THEOREM 7.1. *Let \mathcal{C} be a subset of $\binom{[n]}{k}$. The following are equivalent:*

- (i) the set of Plücker coordinates $\{p_I\}_{I \in \mathcal{C}}$ is a cluster in the cluster algebra structure on the coordinate ring of $G(k, n)$;
- (ii) \mathcal{C} is a maximal weakly separated collection;
- (iii) \mathcal{C} is the collection of face labels of a reduced plabic graph for the uniform matroid $\binom{[n]}{k}$.

In this section, we use the language of cluster algebras freely.

Proof. The implication (ii) \implies (iii) is the main result of this paper; the implication (iii) \implies (i) is [15, Theorem 2]. We now show (i) \implies (ii).

Let I and $J \in \mathcal{C}$; we show that $I \parallel J$. Let $\mathbb{C}_q(G(k, n))$ be the quantization of the coordinate ring of $G(k, n)$ constructed in [9] (see also [8]). Since p_I and p_J lie in a common cluster of $\mathbb{C}[G(k, n)]$, the corresponding quantum minors q_I and q_J lie in a cluster of $\mathbb{C}_q(G(k, n))$. By [2], this means that q_I and q_J quasi-commute. By [11] (see also [14]), this means that I and J are weakly separated.

So we know that \mathcal{C} is a weakly separated collection. Every cluster has cardinality $1 + k(n - k)$, so \mathcal{C} is a maximal weakly separated collection by the bound of [14]. \square

EXAMPLE 7.2. We do not know whether the analogous statement holds for other positroids. (It has only recently been showed that positroid varieties have cluster structure. See [12].) The following near-counterexample emerged in conversation with Kevin Carde: In $G(3, 6)$, consider the Grassmann necklace (123, 234, 345, 451, 561, 612). This is the Schubert divisor where $p_{456} = 0$. It has a cluster structure of type A_3 ; eight of the nine cluster variables are the Plücker coordinates $\{124, 125, 134, 135, 136, 146, 235, 245\}$ and the ninth is $x := p_{145}p_{236}$. In our opinion, clusters involving x are morally made of Plücker coordinates, as x is the product of a Plücker variable and a frozen variable. Note also that, if the coordinate ring of this variety deforms to a quantum cluster algebra (which we expect would follow from the results of [8]), then q_{145} and $q_{145}q_{236}$ would quasi-commute in this quantum cluster algebra; canceling the non-zero divisor q_{145} from the left of $q_{145}(q_{145}q_{236}) = c(q_{145}q_{236})q_{145}$, we conclude that q_{145} and q_{236} would quasi-commute in this quantum cluster algebra, even though 145 and 236 are not weakly separated.

8. Plabic graphs and weakly separated sets

Our goal in this section is to show that, for any reduced plabic graph G with associated decorated permutation π , the collection $\mathcal{F}(G)$ is a weakly separated collection for $\mathcal{M}_{\mathcal{I}(\pi)}$. This is the easy part of Theorem 6.6. We show that this collection is maximal in Section 9, and that any maximal weakly separated collection is of the form $\mathcal{F}(G)$ in Section 11.

PROPOSITION 8.1. *Let G be a reduced plabic graph. Then $\mathcal{F}(G)$ is a weakly separated collection.*

Proof. Assume $\mathcal{F}(G)$ is not weakly separated. Pick I and $J \in \mathcal{F}(G)$ such that $I \not\parallel J$. Let a, b, c and d be cyclically ordered elements such that $\{a, c\} \subset I \setminus J$ and $\{b, d\} \subset J \setminus I$.

We first consider the case where strands a and c do not cross (see Figure 6). In this case, region I is to the left of strands a and c , and region J is to their right, so strands a and c must be parallel, not antiparallel. In other words, the endpoints $(\pi^{-1}(a), a, b, c, \pi^{-1}(c))$ are circularly ordered. If $\pi^{-1}(b)$ is in $(b, \pi^{-1}(a))$, then strands a and b cannot cross and I is on the left-hand side of b , contrary to our desires. If $\pi^{-1}(b)$ is in $(\pi^{-1}(c), b)$, then strand b cannot cross strand c and we deduce that J is to the right of strand b , again contrary to our desires. But the intervals $(b, \pi^{-1}(a))$ and $(\pi^{-1}(a), c)$ cover all of $[n]$, so this excludes all possible positions for $\pi^{-1}(b)$ and we have a contradiction. This concludes the proof in the case that strands a and c do not cross.

Now, suppose that strands a and c cross (see Figure 7). Let R be the region which is to the left of a and right of c ; let T be to the right of a and left of c , and let S_1, S_2, \dots, S_s be the regions which are both to the left of a and c or both to the right. Our numbering is such that strand a first passes by S_1 and precedes in increasing order, while c starts at S_s and goes in




FIGURE 6. The case where strands a and c do not cross.




FIGURE 7. The case where strands a and c do cross.




FIGURE 8. The effect of (M1) on face labels.

decreasing order. Let face I be in S_i and J be in S_j ; note that $i \neq j$. We discuss the case that $i < j$; the case that $i > j$ is equivalent by relabeling (a, b, c, d) as (c, d, a, b) .

Consider strand d , as it passes through the various regions R, S_1, \dots, S_s and T . (In Figure 7, a portion of strand d is shown as a dotted line.) We claim that it is impossible that strand d

both passes from R to S_e to T and from T to S_f to R . The proof is simple: If it did, either its crossings with strand a or its crossings with c would violate condition (3) in the definition of reducedness. On the other hand, strand d must enter both R and T , as otherwise regions I and J would lie on the same side of strand d . Let m be the index such that strand d travels from R to S_m to T , or else from T to S_m to R . We must have $i \leq m \leq j$, since J and I are opposite sides of d , and d must travel from R to S_m to T , as I is on the right of d and J on the left (Figure 7).

But the boundary point d , where strand d terminates, is in the cyclic interval (c, a) . So d cannot end in T . It also cannot end in R , as it goes from R to S_m to T . So it must end in S_1 or S_s . We discuss the former case, since the latter case is similar. If d ends in S_1 , it must come from T , passing through a . But then the intersections of c and d violate condition (3). \square

Now we want to show that $\mathcal{F}(G)$ is actually a weakly separated collection of $\mathcal{M}_{\mathcal{I}(\pi^\cdot)}$. We need the following proposition.

LEMMA 8.2. *Let G be a reduced plabic graph with decorated permutation π^\cdot such that $\ell(\pi^\cdot) < k(n - k)$. There is some pair of strands (i, j) in π^\cdot which only crosses once, at an edge e of G , and such that, if we delete the edge e , thus uncrossing the strands, the resulting graph G' is still reduced.*

Proof. According to [13, Lemma 18.9], a removable edge has these properties. (See [13] for definitions.) According to [13, Corollary 18.10], the removable edges of G are in bijection with the decorated permutations covered by π^\cdot . Since $\ell(\pi^\cdot) < k(n - k)$, it covers at least one decorated permutation. \square

PROPOSITION 8.3. *If G is a reduced plabic graph with decorated permutation π^\cdot , then the following properties hold:*

- (i) *the boundary cells of G are labeled by $\mathcal{I}(\pi^\cdot)$;*
- (ii) *every face of G receives a separate label in $\mathcal{F}(G)$;*
- (iii) *$\mathcal{F}(G)$ is contained in $\mathcal{M}_{\mathcal{I}(\pi^\cdot)}$.*

Proof. We are going to use induction on $k(n - k) - \ell(\pi^\cdot)$. When this quantity is 0, any reduced plabic graph has only one face due to Theorem 6.8. This is possible only when π is the identity. Then there is only one face, and it has the label $I_1 = I_2 = \dots = I_n$. This covers the base case for our induction argument.

Let i, j, e and G' be as in Lemma 8.2. By induction hypothesis, all the properties hold for G' . We use μ^\cdot to denote the decorated permutation of G' and use $\mathcal{I}' = (I'_1, \dots, I'_n)$ for the Grassmann necklace of G' .

We start with the first property. Swapping i and j going from μ^\cdot to π^\cdot corresponds to changing \mathcal{I}' to \mathcal{I} in a way that, if some I'_k contains only one of i and j , then it gets swapped with the other. When going from G' to G , the label of a boundary face gets changed only when it contains exactly one of i and j , and the change is by swapping one with the other. So the boundary cells of G are labeled by $\mathcal{I}(\pi^\cdot)$.

Now we check the second property. Since i and j only cross once, if two faces are assigned the same label in G , which means they have to be in the same region with respect to strands i and j . But this also means they are assigned the same label in G' , a contradiction.

We prove the third property directly. Assume for the sake of contradiction that we have some $J \in \mathcal{F}(G)$ such that $J \not\leq_i I_i$. For $J = \{j_1, \dots, j_k\}$ and $I_i = \{h_1, \dots, h_k\}$, let t denote the first position such that $j_t <_i h_t$. From this, we obtain $j_t \notin I_i$, and hence $\pi^{-1}(j_t) \leq_i j_t$. And for

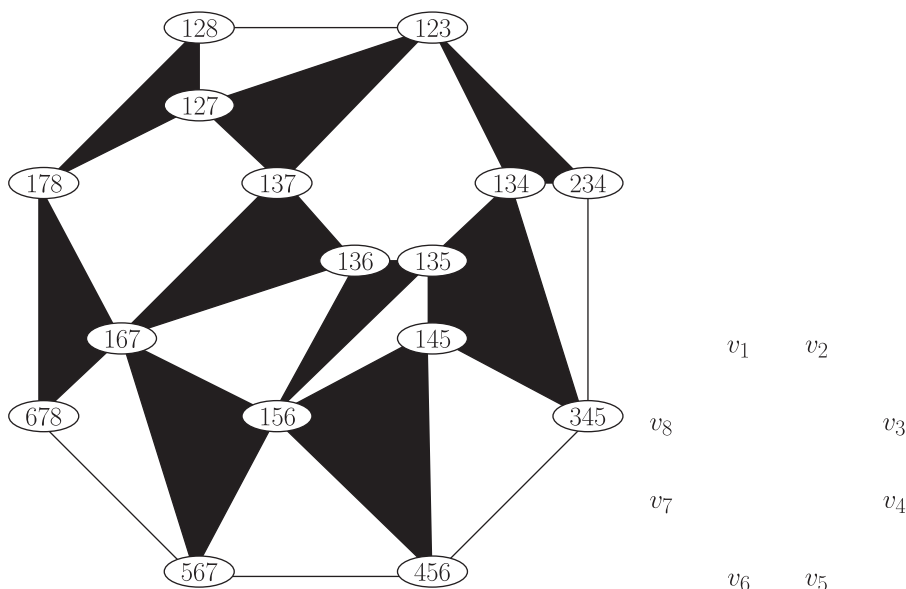


FIGURE 9. $\Sigma(\mathcal{C})$ for \mathcal{C} the face labels of Figure 1, and the points v_i .

any $q \in [t, k]$, we have $\pi^{-1}(h_t) \geq_i h_t$, and it follows from comparing the strands j_t and h_q that J should contain h_q . But this implies that $|J| = k + 1$, a contradiction. \square

Combining the first and third property of Proposition 8.3 with Proposition 8.1, we achieve the following claim.

COROLLARY 8.4. *If G is a reduced plabic graph with boundary \mathcal{I} , then $\mathcal{F}(G)$ is a weakly separated collection in $\mathcal{M}_{\mathcal{I}}$.*

Proof. We have $\mathcal{I} \subseteq \mathcal{F}(G)$ by Proposition 8.3; that \mathcal{I} is weakly separated by Proposition 8.1; and that $\mathcal{F}(G) \subseteq \mathcal{M}_{\mathcal{I}}$ by Proposition 8.3. \square

We have not yet shown that $\mathcal{F}(G)$ is maximal; that is Theorem 9.16.

9. Plabic tilings

Given a maximal weakly separated collection \mathcal{C} of $\binom{[n]}{k}$, we need to construct a plabic graph G such that $\mathcal{F}(G) = \mathcal{C}$. To do so, we define a *plabic tiling*. For a weakly separated collection \mathcal{C} , we construct a two-dimensional CW-complex embedded in \mathbb{R}^2 , and denote this complex by $\Sigma(\mathcal{C})$. Maximal plabic tilings turn out to be dual to reduced bipartite plabic graphs. We associate a plabic tiling to any weakly separated collection, maximal or not. See Example 11.3 for a discussion of $\Sigma(\mathcal{C})$ for a non-maximal \mathcal{C} .

Figure 9 shows $\Sigma(\mathcal{C})$ for \mathcal{C} the collection of face labels in Figure 1. The points v_i , which we introduce soon, are shown in the lower right of Figure 9. They are shown in the same scale as the rest of the figure.

Let us fix \mathcal{C} , a weakly separated collection in $\binom{[n]}{k}$. For I and $J \in \binom{[n]}{k}$, say that I neighbors J if

$$|I \setminus J| = |J \setminus I| = 1.$$

Let K be any $(k - 1)$ element subset of $[n]$. We define the *white clique* $\mathcal{W}(K)$ to be the set of $I \in \mathcal{C}$ such that $K \subset I$. Similarly, for L a $(k + 1)$ element subset of $[n]$, we define the *black clique* $\mathcal{B}(L)$ for the set of $I \in \mathcal{C}$ which are contained in L . We call a clique *non-trivial* if it has at least three elements. Observe that, if \mathcal{X} is a non-trivial clique, then it cannot be both black and white.

REMARK 9.1. Let G be a reduced plabic graph of rank k . The black (respectively, white) cliques of $\mathcal{F}(G)$ correspond to the black (white) vertices of G . More precisely, for each vertex v , the labels of the faces bordering v form a clique, and all cliques are of this form.

Observe that a white clique $\mathcal{W}(K)$ is of the form $\{Ka_1, Ka_2, \dots, Ka_r\}$ for some a_1, a_2, \dots, a_r , which we take to be cyclically ordered. Similarly, $\mathcal{B}(L)$ is of the form $\{L \setminus b_1, L \setminus b_2, \dots, L \setminus b_s\}$, where the indices b_i are circularly ordered. If $\mathcal{W}(K)$ is non-trivial, then we define the boundary of $\mathcal{W}(K)$ to be the cyclic graph

$$(Ka_1) \longrightarrow (Ka_2) \longrightarrow \dots \longrightarrow (Ka_r) \longrightarrow (Ka_1).$$

Similarly, the boundary of a non-trivial $\mathcal{B}(L)$ is

$$(L \setminus b_1) \longrightarrow (L \setminus b_2) \longrightarrow \dots \longrightarrow (L \setminus b_s) \longrightarrow (L \setminus b_1).$$

If (J, J') is a two element clique, then we define its boundary to be the graph with a single edge (J, J') ; we define an one element clique to have empty boundary.

LEMMA 9.2. *Let I neighbor J ; set $K = I \cap J$ and $L = I \cup J$. If $\mathcal{W}(K)$ and $\mathcal{B}(L)$ are non-trivial, then there is an edge between I and J in the boundaries of $\mathcal{W}(K)$ and $\mathcal{B}(L)$.*

Proof. Let i be the lone element of $I \setminus J$ and j the lone element of $J \setminus I$. Let $\mathcal{W}(K) = \{Ka_1, Ka_2, \dots, Ka_r\}$ and $\mathcal{B}(L) = \{L \setminus b_1, L \setminus b_2, \dots, L \setminus b_s\}$. So the i and j are both in $\{a_1, \dots, a_r\}$ and in $\{b_1, \dots, b_s\}$. Consider the four sets:

$$\begin{aligned} S_1 &:= \{a_m : a_m \in (i, j)\}, & S_2 &:= \{a_m : a_m \in (j, i)\}, \\ S_3 &:= \{b_m : b_m \in (i, j)\}, & S_4 &:= \{b_m : b_m \in (j, i)\}. \end{aligned}$$

Our goal is to show that either $S_1 = S_3 = \emptyset$ or $S_2 = S_4 = \emptyset$.

Suppose (for the sake of contradiction) that S_1 and S_4 are both non-empty, with $a \in S_1$ and $b \in S_4$. Set $P = Ka$ and $Q = L \setminus b$. Then $a \in P \setminus Q$, $i \in Q \setminus P$, $b \in P \setminus Q$ and $j \in Q \setminus P$, so P and Q are not weakly separated. We have a contradiction and we deduce that at least one of S_1 and S_4 is empty. Similarly, at least one of S_2 and S_3 is empty. On the other hand, $\mathcal{W}(K)$ is non-trivial, so at least one of S_1 and S_2 is non-empty. Similarly, at least one of S_3 and S_4 is non-empty.

Since at least one of (S_1, S_4) is empty, and at least one of (S_2, S_3) is, we deduce that at least two of (S_1, S_2, S_3, S_4) are empty. Similarly, at most two of them are empty, so precisely two of them are empty. Checking the six possibilities, those consistent with the above restrictions are that either S_1 and S_3 are empty, and S_2 and S_4 are not, or vice versa. \square

We now define a two-dimensional CW-complex $\Sigma(\mathcal{C})$. The vertices of $\Sigma(\mathcal{C})$ are the elements of \mathcal{C} . There is an edge (I, J) if

- (1) $\mathcal{W}(I \cap J)$ is non-trivial and (I, J) appears in the boundary of $\mathcal{W}(I \cap J)$ or
- (2) $\mathcal{B}(I \cup J)$ is non-trivial and (I, J) appears in the boundary of $\mathcal{B}(I \cup J)$ or
- (3) $\mathcal{W}(I \cap J) = \mathcal{B}(I \cup J) = \{I, J\}$ (this third case is only important in certain boundary cases involving disconnected positroids. For most of the results of this paper, it is not important whether or not we include an edge in this case).

There is a (two-dimensional) face of $\Sigma(\mathcal{C})$ for each non-trivial clique \mathcal{X} of \mathcal{C} . The boundary of this face is the boundary of \mathcal{X} . By Lemma 9.2, all edges in the boundary of \mathcal{X} are in $\Sigma(\mathcal{C})$, so this makes sense. We refer to each face of $\Sigma(\mathcal{C})$ as *black* or *white*, according to the color of the corresponding clique. We call a CW-complex of the form $\Sigma(\mathcal{C})$ a *plabic tiling*. So far, $\Sigma(\mathcal{C})$ is an abstract CW-complex. Our next goal is to embed it in a plane.

Fix n points v_1, v_2, \dots, v_n in \mathbb{R}^2 , at the vertices of a convex n -gon in clockwise order. Define a linear map $\pi : \mathbb{R}^n \rightarrow \mathbb{R}^2$ by $e_a \mapsto v_a$. For $I \in \binom{[n]}{t}$, set $e_I = \sum_{a \in I} e_a$. We abbreviate $\pi(e_I)$ by $\pi(I)$.

We need a notion of weak separation for vectors in $\mathbb{R}^{[n]}$. Let e and f be two vectors in $\mathbb{R}^{[n]}$, with $\sum_{a \in [n]} e_a = \sum_{a \in [n]} f_a$. We define e and f to be *weakly separated* if there do *not* exist a, b, a' and b' , cyclically ordered, with $e_a > f_a, e_b < f_b, e_{a'} > f_{a'}$ and $e_{b'} < f_{b'}$. So, for I and $J \in \binom{[n]}{k}$, we have $I \parallel J$ if and only if $e_I \parallel e_J$.

LEMMA 9.3. *Let e and f be two different points in \mathbb{R}^n , with $\sum_{a \in [n]} e_a = \sum_{a \in [n]} f_a$ and $e \parallel f$. Then $\pi(e) \neq \pi(f)$.*

Proof. Since e and f are weakly separated, there are some ℓ and r in $[n]$ such that $e_a - f_a$ is non-negative for a in $[\ell, r)$ and is non-positive for b in $[r, \ell)$. We have

$$\pi(e) - \pi(f) = \sum_{a \in [\ell, r)} (e_a - f_a)v_a - \sum_{b \in [r, \ell)} (f_b - e_b)v_b. \tag{9.1}$$

Since $\sum_{a \in [n]} e_a = \sum_{b \in [n]} f_b$, the right-hand side of the above equation is a positive linear combination of vectors of the form $v_a - v_b$, with $a \in [\ell, r)$ and $b \in [r, \ell)$.

Since the points v_a are the vertex of a convex n -gon, there is a line λ in \mathbb{R}^2 separating $\{v_a : a \in [\ell, r)\}$ from $\{v_b : b \in [r, \ell)\}$. So every vector $v_a - v_b$ as above crosses from the former side of λ to the latter. A positive linear combination of such vectors must cross the line λ , and can therefore not be zero. So the right-hand side of the above equation is non-zero, and we deduce that $\pi(e) \neq \pi(f)$. □

We extend the map π to a map from $\Sigma(\mathcal{C})$ to \mathbb{R}^2 as follows: Each vertex I of $\Sigma(\mathcal{C})$ is sent to $\pi(I)$ and each face of $\Sigma(\mathcal{C})$ is sent to the convex hull of the images of its vertices. We encourage the reader to consult Figure 9 and see that the vector $\pi(S_i) - \pi(S_j)$ is a translation of $v_i - v_j$.

PROPOSITION 9.4. *For any weakly separated collection \mathcal{C} , the map π embeds the CW-complex $\Sigma(\mathcal{C})$ into \mathbb{R}^2 .*

Proof. Suppose, to the contrary, that there are two faces of $\Sigma(\mathcal{C})$ whose interiors have overlapping image. We deal with the case that both of these faces are two-dimensional, with one white and the other black. The other possibilities are similar, and easier. Let the vertices of the two faces be $\{Ka_1, Ka_2, \dots, Ka_r\}$ and $\{L \setminus b_1, L \setminus b_2, \dots, L \setminus b_s\}$. Write $A = \{a_1, \dots, a_r\}$ and $B = \{b_1, \dots, b_s\}$.

Before we analyze the geometry of π , we need to do some combinatorics. Note that, if $K \subset L$, then the faces are convex polygons with a common edge, lying on opposite sides of that edge, and thus have disjoint interiors (the fact that we have to consider this possibility is one of the things that makes the case of two faces of opposite color more difficult). So we may assume that $K \setminus L$ is non-empty. Also, $|L| = k + 1 > |K| = k - 1$, so $L \setminus K$ is non-empty. Let $x \in K \setminus L$ and $y \in L \setminus K$. Suppose that a_i and b_j are in (x, y) . Then weak separation of Ka_i and $L \setminus b_j$ implies that $a_i \leq_x b_j$. More generally, if $a \in (K \cup A) \setminus L$ and $b \in (L \cup B) \setminus K$ are both in (x, y) , then we still have $a \leq_x b$. Similarly, if $a \in (K \cup A) \setminus L$ and $b \in (L \cup B) \setminus K$ are both

in (y, x) , then $a \succeq_x b$. So, there is some $c \in [x, y]$ and $d \in [y, x]$ such that $(K \cup A) \setminus L \subseteq [d, c]$ and $(L \cup B) \setminus K \subseteq [c, d]$.

Now, let ζ be the point which is in the interior of both faces, say

$$\zeta = \sum p_i \pi(Ka_i) = \sum q_i \pi(L \setminus b_i)$$

for some positive scalars p_i and q_i with $\sum p_i = \sum q_i = 1$. Define the vectors u and v by $u = \sum p_i e_{Ka_i}$ and $v = \sum q_i e_{L \setminus b_i}$, so $\pi(u) = \pi(v)$. All the positive entries of $u - v$ are contained in $[d, c]$, and all the negative entries in $[c, d]$. So u and v are weakly separated and, by Lemma 9.3, $\pi(u) \neq \pi(v)$, a contradiction. \square

Looking through the proof, we have proved the more technical result.

LEMMA 9.5. *Let P and Q be different faces of $\Sigma(\mathcal{C})$. Let $V(P)$ and $V(Q)$ be the sets of vertices of the faces P and Q . Let u be a vector of the form $\sum_{I \in V(P)} c_I e_I$, where $c_I > 0$ and $\sum c_I = 1$. Similarly, let v be a vector of the form $\sum_{I \in V(Q)} d_I e_I$, with $d_I > 0$ and $\sum d_I = 1$. Then u and v are weakly separated.*

REMARK 9.6. Let Ka_1, Ka_2, \dots, Ka_r be the vertices of a white face of \mathcal{C} , with the a_i in cyclic order. Then the vertices $\pi(Ka_i)$ appear in clockwise order in the planar embedding. If the vertices of a black face are $L \setminus b_1, L \setminus b_2, \dots, L \setminus b_r$, with the b_i again in cyclic order, then the vertices $\pi(L \setminus b_i)$ again appear in clockwise order. This is because negation is an orientation preserving operation on \mathbb{R}^2 ; see Figure 9.

Now, we study the construction $\Sigma(\mathcal{C})$ when \mathcal{C} is a weakly separated collection inside a particular positroid \mathcal{M} . Propositions 9.7–9.11 are obvious in the case that \mathcal{M} is the largest positroid, $\binom{[n]}{k}$.

Let \mathcal{M} be a positroid and \mathcal{I} the corresponding Grassmann necklace. We first give a lemma to reduce to the case when \mathcal{I} is connected: Suppose that \mathcal{I} is not connected, with $I_i = I_j$. Let \mathcal{C} be a weakly separated collection for \mathcal{M}^r ; and use the notations \mathcal{C}^r, I^r and so forth from Section 5. It is easy to check that the following proposition.

PROPOSITION 9.7. *The complex $\Sigma(\mathcal{C})$ is formed by gluing $\Sigma(\mathcal{C}^1)$ and $\Sigma(\mathcal{C}^2)$ together at the point $I^1 \cup I^2$.*

See Figure 10 for an example of a disconnected plabic graph and the corresponding plabic tiling.

We now, therefore, restrict to the case that \mathcal{I} is a connected Grassmann necklace. We define $\pi(\mathcal{I})$ to be the closed polygonal curve whose vertices are, in order, $\pi(I_1), \pi(I_2), \dots, \pi(I_n), \pi(I_1)$.

PROPOSITION 9.8. *The curve $\pi(\mathcal{I})$ is a simple closed curve.*

REMARK 9.9. When \mathcal{M} is $\binom{[n]}{k}$, the interior of $\pi(\mathcal{I})$ is convex, but this is not true in general.

Proof. Suppose for the sake of contradiction that the curve $\pi(\mathcal{I})$ crosses through itself. First, by assumption, the I_k are all distinct so, by Lemma 9.3, the vertices of $\pi(\mathcal{I})$ are all




FIGURE 10. A disconnected plabic graph and the corresponding plabic tiling.

distinct. We now rule out the case that line segments $(\pi(I_i), \pi(I_{i+1}))$ and $(\pi(I_j), \pi(I_{j+1}))$ cross in their interior. The case that $(\pi(I_i), \pi(I_{i+1}))$ passes through the vertex $\pi(I_j)$ is similar and easier.

The collection \mathcal{I} is weakly separated by Lemma 4.5. By Proposition 9.4, the map π embeds $\Sigma(\mathcal{I})$ into \mathbb{R}^2 . The only way that the line segments $(\pi(I_i), \pi(I_{i+1}))$ and $(\pi(I_j), \pi(I_{j+1}))$ could cross is if I_i, I_j, I_{i+1} and I_{j+1} are vertices of some two-dimensional face of $\Sigma(\mathcal{I})$, arranged in that circular order. (Or the reverse circular order, but then we could switch the indices i and j .) We consider the case that this face is of the form $\mathcal{W}(K)$, the case of $\mathcal{B}(L)$ is similar. Let $(I_i, I_j, I_{i+1}, I_{j+1}) = (Ka, Kb, Kc, Kd)$; the elements a, b, c and d must be circularly ordered.

Now, since i is the unique element of $I_i \setminus I_{i+1}$, we know that $a = i$. Similarly, $b = j$. But then $Kb \leq_{i+1} Kc$, and the inequality is strict because $I_{i+1} \neq I_j$. This contradicts that I_{i+1} , which is Kc , is the \leq_{i+1} minimal element of \mathcal{M} . \square

The aim of the next several propositions is to establish:

PROPOSITION 9.10. *Let \mathcal{I} be a connected Grassmann necklace. Let J be weakly separated from \mathcal{I} , but not an element of \mathcal{I} . Then $\pi(J)$ is in the interior of $\pi(\mathcal{I})$ if and only if J is in \mathcal{M} .*

Define $\phi(t) : \mathbb{R}/\mathbb{Z} \rightarrow \mathbb{R}^2$ so that ϕ sends the interval $[r/n, (r+1)/n]$ linearly to the line segment from $\pi(I_r)$ to $\pi(I_{r+1})$. So ϕ parameterizes the simple closed curve $\pi(\mathcal{I})$. We consider the path $\phi(t) - \pi(J)$ in $\mathbb{R}^2 \setminus (0, 0)$. The point $\pi(J)$ is inside $\pi(\mathcal{I})$ if and only if $\phi(t) - \pi(J)$ has winding number 1; if not, it has winding number 0. We prove Proposition 9.10 by computing this winding number. We first introduce some notation.

By hypothesis, I_r and J are weakly separated. Let $[a_r, b_r]$ be the smallest cyclic interval containing all the elements of $I_r \setminus J$, and $[c_r, d_r]$ the smallest cyclic interval containing all the elements of $J \setminus I_r$. Since $J \neq I_r$, these intervals are non-empty. So a_r, b_r, c_r and d_r are

circularly ordered, with $b_r \neq c_r$ and $d_r \neq a_r$. Now,

$$\pi(I_r) - \pi(J) = \sum_{i \in I_r \setminus J} v_i - \sum_{j \in J \setminus I_r} v_j.$$

Let p_r be the centroid of $\{v_i\}_{i \in I_r \setminus J}$, and q_r the centroid of $\{v_j\}_{j \in J \setminus I_r}$; since I_r and J are weakly separated, $p_r \neq q_r$. So $\pi(I_r) - \pi(J)$ is parallel to $p_r - q_r$.

LEMMA 9.11. *There are circular intervals $[a'_r, b'_r]$, and $[c'_r, d'_r]$, with (a'_r, b'_r, c'_r, d'_r) circularly ordered, so that $[a_r, b_r]$ and $[a_{r+1}, b_{r+1}]$ are contained in $[a'_r, b'_r]$ and $[c_r, d_r]$ and $[c_{r+1}, d_{r+1}]$ are contained in $[c'_r, d'_r]$.*

Proof. We have to consider four cases, depending on whether or not r and $\pi(r)$ are in J . Let us consider the case that neither of them is. In this case, $[c_r, d_r] = [c_{r+1}, d_{r+1}]$ and we can take $[c'_r, d'_r] = [c_r, d_r]$. The intervals $[a_r, b_r]$ and $[a_{r+1}, b_{r+1}]$ are disjoint from $[c_r, d_r]$, and thus live in (d_r, c_r) . It is thus easy to see that there is a minimal subinterval of (d_r, c_r) which contains $[a_r, b_r]$ and $[a_{r+1}, b_{r+1}]$; we take this interval to be $[a'_r, b'_r]$.

The other three cases are similar to this one. □

For t a real number between r/n and $(r + 1)/n$, define $p(t)$ by linear interpolation between p_r and p_{r+1} . Similarly, define $q(t)$ by linear interpolation between q_r and q_{r+1} . For t between r/n and $(r + 1)/n$, the point $p(t)$ is in the convex hull of $\{v_i\}_{i \in [a'_r, b'_r]}$ and $q(t)$ is in the convex hull of $\{v_j\}_{j \in [c'_r, d'_r]}$. Because the intervals $[a'_r, b'_r]$ and $[c'_r, d'_r]$ are disjoint, we know that $p(t) \neq q(t)$. The vector $\phi(t) - \pi(J)$ is a positive scalar multiple of $p(t) - q(t)$.

Let γ be a convex simple closed curve through the points v_i . Let $x(t)$ and $y(t)$ be the points where the line through $p(t)$ and $q(t)$ meets γ , with $x(t)$ closer to $p(t)$. So $x(t)$ lies on the part of γ between $v_{a'_r}$ and $v_{b'_r}$; the point $y(t)$ lies on the part of γ between $v_{c'_r}$ and $v_{d'_r}$. Note that $x(t)$ and $y(t)$ are continuous functions of t .

It is geometrically clear that the paths $x(t)$ and $y(t)$, traveling around γ , have the same winding number as the path $x(t) - y(t)$ around 0. As observed above, the vector $x(t) - y(t)$ is a positive multiple of $\phi(t) - \pi(J)$. Thus, we are reduced to computing the winding number of $x(t)$ around γ .

Proof of Proposition 9.10. First, suppose that $J \in \mathcal{M}$. This means that, for every r , we have $J \geq_r I_r$; since $J \notin \mathcal{I}$, this inequality is in fact strict. The indices (r, a_r, b_r, c_r, d_r) are circularly ordered; the inequality between b_r and c_r are strict, while the others are weak. So the interval $[a'_r, b'_r]$ can only contain r at its left extreme. For t between r/n and $(r + 1)/n$, the point $x(t)$ lies in the interval $[a'_r, b'_r]$, so $x(t)$ has winding number 1, as desired.

Now, suppose that $J \notin \mathcal{M}$. Then, for some $s \in [n]$, we have $J \not\leq_s I_s$; by hypothesis, $J \parallel I_s$. So one of the following two cases must hold:

- (i) the minimal element of $J \setminus I_s$ is less than the minimal element of $I_s \setminus J$, or
- (ii) the maximal element of $I_s \setminus J$ is greater than the maximal element of $J \setminus I_s$.

Here and in the rest of this proof, ‘minimal’, ‘maximal’, ‘less than’ and ‘greater than’ refer to the orders \leq_s .

We discuss the former case; the latter is similar. We know that I_s is minimal among the sets I_r . So, for every r , the minimal element of $J \setminus I_r$ must be less than the minimal element of $I_r \setminus J$, in the \leq_s order.

So, for every r , the interval $[a_r, b_r]$ does not contain s . So $[a'_r, b'_r]$ also does not contain s , and we see that $x(t)$ is not v_s , for any t . So $x(t)$ has winding number 0, as desired. □

We now apply our constructions to the case that G is a plabic graph.

THEOREM 9.12. *Let G be a bipartite plabic graph whose strand permutation is connected and let $\mathcal{C} = \mathcal{F}(G)$. Then $\Sigma(\mathcal{C})$ is isomorphic to the dual complex to the planar graph G . Also, $\pi(\Sigma(\mathcal{F}(G)))$ fills the region surrounded by $\pi(\mathcal{I})$.*

REMARK 9.13. If G is a reduced plabic graph which may not be bipartite, let G' be the graph obtained by contracting edges between vertices of the same color, as discussed in Remark 6.3. Then the faces of G and G' are in bijection, and they have the same labels. Letting \mathcal{C} be this common set of labels, it then follows from Theorem 9.12 that $\Sigma(\mathcal{C})$ is dual to G' .

REMARK 9.14. If G is a bipartite plabic graph whose strand permutation is *not connected*, then $\Sigma(\mathcal{C})$ is a union of $\Sigma(\mathcal{C}')$'s for various smaller collections \mathcal{C}' , by Proposition 9.7. Each of these collections comes from a connected plabic graph G' , to which Theorem 9.12 applies. See Figure 10 for an example.

Proof of Theorem 9.12. By Proposition 8.1, \mathcal{C} is weakly separated, so it makes sense to define $\Sigma(\mathcal{C})$. By definition, the vertices of $\Sigma(\mathcal{C})$ correspond to the faces of G , with a separate vertex for each face of G by Proposition 8.3. By Propositions 9.4 and 9.10, we know that $\Sigma(\mathcal{C})$ embeds into \mathbb{R}^2 , with the boundary vertices corresponding to \mathcal{I} . Let R be the region in \mathbb{R}^2 surrounded by the boundary of $\pi(\Sigma(\mathcal{C}))$.

Let $D(G)$ be the dual to the planar graph G . So $D(G)$ is a two-dimension CW-complex, homeomorphic to a disc, with vertices labeled by \mathcal{C} , and with boundary labeled by \mathcal{I} . We now describe a map ψ of $D(G)$ into \mathbb{R}^2 : For a vertex v of $D(G)$, if J is the label of the dual face of G , then $\psi(v)$ is $\pi(J)$. For each edge e of $D(G)$, connecting u and v , let $\psi(e)$ be a line segment from $\psi(u)$ to $\psi(v)$.

We now discuss faces of $D(G)$. Let F be a face of $D(G)$, dual to a vertex v of G . Suppose that v is black; the case where v is white is similar. Let the edges ending at v cross the strands i_1, i_2, \dots, i_r . So the faces bordering v have labels of the form $L \setminus i_1, L \setminus i_2, \dots, L \setminus i_r$ for some L . In particular, we see that the corresponding vertices of $D(G)$ are in $\mathcal{B}(L)$. The face $\mathcal{B}(L)$ is embedded in \mathbb{R}^2 as a convex polygon. Moreover, the vertices of F occur in cyclic order among those of $\mathcal{B}(L)$. So ψ embeds ∂F as a convex polygon in \mathbb{R}^2 , and we extend ψ to take the interior of F to the interior of that polygon. Note that we have shown that, for every face F of $D(G)$, the image $\psi(F)$ lies within a face of $\Sigma(\mathcal{C})$ of the same color.

The map ψ takes the boundary of $D(G)$ to ∂R . For every face F of $D(G)$, the face $\psi(F)$ is embedded in \mathbb{R}^2 with the correct orientation. Since $D(G)$ is homeomorphic to a disc, these points force ψ to be a homeomorphism onto R .

Every face of $\psi(D(G))$ lies in a face of $\pi(\Sigma(\mathcal{C}))$, so $\pi(\Sigma(\mathcal{C})) \supseteq \psi(D(G)) = R$. By Proposition 9.10, we also have $\pi(\Sigma(\mathcal{C})) \subseteq R$, so we conclude that $\pi(\Sigma(\mathcal{C})) = \psi(D(G)) = R$. We now must show that the two CW-structures on this subset of \mathbb{R}^2 coincide.

Let F be any face of $D(G)$. We consider the case that F is black; the white case is similar. As shown above, there is a black face $\pi(\mathcal{B}(L))$ of $\Sigma(\mathcal{C})$ which contains $\psi(F)$. Suppose, for the sake of contradiction, that some point of $\pi(\mathcal{B}(L))$ is not in $\psi(\Sigma(\mathcal{C}))$. Since $D(G)$ covers all of R , there must be some face F' of $D(G)$, adjacent to F , such that $\psi(F')$ overlaps $\pi(\mathcal{B}(L))$. But F' , being a neighbor of F , is white and thus lies in some $\pi(\mathcal{W}(K))$. So $\pi(\mathcal{B}(L))$ and $\pi(\mathcal{W}(K))$ overlap, a contradiction. We conclude that every face of $\psi(D(G))$ is also a face of $\pi(\Sigma(\mathcal{C}))$. Since ψ is injective, and has image R , we see that every face of $\pi(\Sigma(\mathcal{C}))$ is precisely one face of $\psi(D(G))$.

We now have a bijection between the faces of $\Sigma(\mathcal{C})$ and of $D(G)$. In both $\pi(\Sigma(\mathcal{C}))$ and $\psi(D(G))$, a face is the convex hull of its vertices, so the corresponding faces of $\Sigma(\mathcal{C})$ and $D(G)$ must have the same vertices and the same edges. So $\pi^{-1} \circ \psi$ is a cell-by-cell homeomorphism between the simplicial complexes $\Sigma(\mathcal{C})$ and $D(G)$. \square

REMARK 9.15. The most technical part of this paper is establishing that, when \mathcal{C} is a maximal weakly separated collection which is not yet known to be of the form $\mathcal{F}(G)$, then $\pi(\Sigma(\mathcal{F}(G)))$ fills the interior of the region surrounded by $\pi(\mathcal{I})$. This task is accomplished in Section 11.

We now show that $\mathcal{F}(G)$ is always maximal.

THEOREM 9.16. *Let G be a reduced plabic graph, with boundary \mathcal{I} . Then $\mathcal{F}(G)$ is a maximal weakly separated collection in $\mathcal{M}_{\mathcal{I}}$.*

Proof. We first reduce to the case that \mathcal{I} is connected. If not, let $I_i = I_j$. We resume the notations of Section 5. The face of G labeled I_i touches the boundary of the disc at two points; both between boundary vertices i and $i + 1$ and between boundary vertices j and $j + 1$. So this face disconnects the graph G into two graphs, with vertices in I^1 and the other with vertices in I^2 . Call these plabic graphs G^1 and G^2 . It is an easy exercise that each G^j is a reduced plabic graph, with boundary \mathcal{I}^j . By induction on n , we know that $\mathcal{F}(G^j)$ is a maximal weakly separated collection in $\mathcal{M}_{\mathcal{I}^j}$.

Let J in $\mathcal{M}_{\mathcal{I}}$ be weakly separated from $\mathcal{F}(G)$. Since $\mathcal{F}(G) \cup \{J\}$ is weakly separated, by Proposition 5.13, J must either be of the form $J^1 \cup J^2$ or $I^1 \cup J^2$, where J^r is weakly separated from $\mathcal{F}(G^r)$ and in $\mathcal{M}_{\mathcal{I}^r}$. But, by the maximality of $\mathcal{F}(G^r)$, this implies that J^r is in $\mathcal{F}(G^r)$. Then J is the label of the corresponding face of $\mathcal{F}(G)$. We have shown that any J in $\mathcal{M}_{\mathcal{I}}$ which is weakly separated from $\mathcal{F}(G)$ must lie in $\mathcal{F}(G)$, so $\mathcal{F}(G)$ is maximal in $\mathcal{M}_{\mathcal{I}}$.

We now assume that \mathcal{I} is connected. Let J in $\mathcal{M}_{\mathcal{I}}$ be weakly separated from $\mathcal{F}(G)$ and assume, for the sake of contradiction, that $J \notin \mathcal{F}(G)$. We showed above that $\pi(\Sigma(\mathcal{F}(G)))$ fills the interior of the region surrounded by $\pi(\mathcal{I})$. By Proposition 9.10, $\pi(J)$ is in this region, so $\pi(J)$ lands on some point of $\pi(\Sigma(\mathcal{F}(G)))$. In particular, π is not injective on $\Sigma(\mathcal{F}(G) \cup \{J\})$, contradicting the assumption that J is weakly separated from \mathcal{C} . This contradiction shows that $\mathcal{F}(G)$ is maximal in $\mathcal{M}_{\mathcal{I}}$. \square

10. Proof of a key lemma

Our aim in this section is to prove the following technical lemma. The reader may wish to skip this proof on first reading, and go on to see how it is used in the proof of Theorem 6.6.

LEMMA 10.1. *Let \mathcal{C} be a weakly separated collection of $\binom{[n]}{k}$. Let $E \subset \binom{[n]}{k-1}$ and $x \neq y \in [n]$ be such that $E \cup \{x\}$ and $E \cup \{y\}$ are in \mathcal{C} . If E is not the interval (x, y) , then there is a set of one of the following two types which is weakly separated from \mathcal{C} :*

- (i) $M^a := E \cup \{a\}$ with $a \in (x, y) \setminus E$ or
- (ii) $N^b := E \cup \{x, y\} \setminus \{b\}$ with $b \in (y, x) \cap E$.

Note that this lemma is taken to itself under the symmetry which replaces every subset of $[n]$ by its complement and switches x and y . We call this symmetry *dualization*.

The condition that $E \neq (x, y)$ means that either there exists some $a \in (x, y) \setminus E$ or there exists some $b \in (y, x) \cap E$ (or both).

Consider under what circumstances a set of the form M^a might not be weakly separated from \mathcal{C} . In order for this to happen, there must be some J in \mathcal{C} such that $J \not\parallel M^a$. The next lemma studies the properties of such a J .

LEMMA 10.2. *Let E , x and y be as above, and let $a \in (x, y) \setminus E$. Suppose that $J \parallel Ex$, $J \parallel Ey$ but $J \not\parallel Ea$. Then there are elements ℓ and r such that:*

- (i) ℓ and r are in $J \setminus E$;
- (ii) $x \leq_x \ell <_x a <_x r \leq_x y$;
- (iii) $(r, \ell) \cap (E \setminus J) \neq \emptyset$;
- (iv) $(\ell, r) \cap (J \setminus E) = \emptyset$.

Proof. Note that J is a vertex of $\Sigma(\mathcal{C})$, and (Ex, Ey) is contained in a face of $\Sigma(\mathcal{C})$. (It may or may not be an edge.) Then, by Lemma 9.5, the vectors e_J and $e_E + (1/2)(e_x + e_y)$ are weakly separated. By hypothesis, e_J and $e_E + e_a$ are not weakly separated. The only coordinates to change between $e_J - (e_E + (1/2)(e_x + e_y))$ and $e_J - e_E - e_a$ are in positions x , y and a . Coordinates x and y increase by $1/2$ from half integers to integers. In particular, they cannot change from non-negative to negative or non-positive to positive. So, the failure of weak separation between J and Ea must be attributable to coordinate a . We must have $a \notin J$, and we must have $a \in (\ell, r)$ for some ℓ and $r \in J \setminus E$ such that $(r, \ell) \cap (E \setminus J) \neq \emptyset$. Moreover, we choose to take the interval (ℓ, r) to be of minimal length, subject to the conditions that $a \in (\ell, r)$ and ℓ and $r \in J \setminus E$. This ensures that $(\ell, r) \cap (J \setminus E) = \emptyset$.

We are left to check that ℓ and r are in $[x, y]$. We check that ℓ is; the case of r is similar. Suppose, instead, that $\ell \in (y, x)$. If $x \in J$, then (x, r) would be closer to a than (ℓ, r) , contradicting our minimal choice of (ℓ, r) . So $x \notin J$. Looking at (ℓ, x, a, y) , we see that J and Ey are not weakly separated, a contradiction. \square

We define (J, ℓ, r) to be a *witness against M^a* if $J \in \mathcal{C}$ and the numbered conditions of Lemma 10.2 hold.

LEMMA 10.3. *Suppose that $(x, y) \setminus E$ is non-empty, and no M^a is weakly separated from \mathcal{C} . Then there is a sequence of triples, $(J_1, \ell_1, r_1), \dots, (J_q, \ell_q, r_q)$, each of which obeys*

- (i) $J_m \in \mathcal{C}$;
- (ii) ℓ_m and r_m are in $J_m \setminus E$;
- (iii) $x \leq_x \ell_m <_x r_m \leq_x y$;
- (iv) $(r_m, \ell_m) \cap (E \setminus J_m) \neq \emptyset$;
- (v) $(\ell_m, r_m) \cap (J_m \setminus E) = \emptyset$.

and such that $x = \ell_1 <_x \ell_2 <_x \dots <_x \ell_q$ and $r_1 <_x r_2 <_x \dots <_x r_q = y$ and $\ell_{i+1} <_x r_i$.

Proof. Consider any $a \in (x, y) \setminus E$. Since M^a is not weakly separated from \mathcal{C} , there is some J in M^a which is not weakly separated from M^a . Since Ex and Ey are in \mathcal{C} , we know that J is weakly separated from Ex and Ey . So we can complete J to a witness, (J, ℓ, r) , against M^a . This (J, ℓ, r) obeys the conditions of Lemma 10.2.

Choose a minimal collection $(J_1, \ell_1, r_1), \dots, (J_q, \ell_q, r_q)$, obeying the conditions of Lemma 10.2, so that every $a \in (x, y) \setminus E$ lies in some (ℓ_m, r_m) . (Since there is some $a \in (x, y) \setminus E$, this minimal collection is non-empty.) We cannot have $(\ell_i, r_i) \subseteq (\ell_j, r_j)$ for any

$i \neq j$, as otherwise we could delete (J_i, ℓ_i, r_i) and have a smaller collection. Thus, we may reorder our collection so that $\ell_1 <_x \ell_2 <_x \dots <_x \ell_q$ and $r_1 <_x r_2 <_x \dots <_x r_q = y$.

If $\ell_1 \in (x, y)$, then ℓ_1 is an element of $(x, y) \setminus E$ and ℓ_1 does not lie in any (ℓ_m, r_m) , a contradiction. Similarly, $r_q \notin (x, y)$. So $\ell_1 = x$ and $r_q = y$. Finally, suppose that $\ell_{i+1} \geq_x r_i$. Then r_i is an element of $(x, y) \setminus E$ which is not in any (ℓ_m, r_m) , a contradiction. \square

We now assume that $(x, y) \setminus E$ is non-empty and that no M^a is weakly separated from \mathcal{C} . For clarity, we continue to state these assumptions explicitly in all lemmas that rely on them. We fix, once and for all, a sequence of triples (J_m, ℓ_m, r_m) as in Lemma 10.3.

We introduce a total order \sqsubset on $\binom{[n]}{k}$: Given I and J in $\binom{[n]}{k}$, first compare $|J \cap (x, y)|$ and $|I \cap (x, y)|$. If $|I \cap (x, y)| < |J \cap (x, y)|$, set $I \sqsubset J$. If this comparison is inconclusive and $|I \cap (y, x)| > |J \cap (y, x)|$, set $I \sqsubset J$. If both of these comparisons are inconclusive, and $I \cap (x, y)$ is lexicographically before $J \cap (x, y)$, set $I \sqsubset J$; if that too is inconclusive, and $I \cap (y, x)$ is lexicographically after $J \cap (y, x)$, set $I \sqsubset J$. (Here we use the orderings $<_x$ and $<_y$ on (x, y) and (y, x) , respectively.) Finally, if all of these are inconclusive (which decision we make in this case is completely immaterial to the proof; we just made a decision so that the order would be total), then either $I = J$ or $\{I, J\} = \{Sx, Sy\}$ for some S ; set $Sx \sqsubset Sy$.

This order is constructed to obey the dual lemmas.

LEMMA 10.4. *Suppose that I and J are weakly separated. Suppose that there are $a <_a b <_a c$ in $[x, y]$ such that a and $c \in I \setminus J$ and $b \in J \setminus I$. Then $I \sqsubset J$.*

LEMMA 10.5. *Suppose that I and J are weakly separated. Suppose that there are $a <_a b <_a c$ in $[y, x]$ such that a and $c \in J \setminus I$ and $b \in I \setminus J$. Then $I \sqsubset J$.*

We prove Lemma 10.4 and leave it to the reader to dualize our argument to prove Lemma 10.5.

Proof of Lemma 10.4. Since I and J are weakly separated, there can be no element of $[y, x]$ which is in $J \setminus I$. If there is any element of $I \setminus J$ in $[y, x]$, then $|I \cap (x, y)| < |J \cap (x, y)|$, so we have $I \sqsubset J$ as desired.

If there is no element of $I \setminus J$ in $[y, x]$, then $|I \cap (x, y)| = |J \cap (x, y)|$ and $|I \cap (y, x)| = |J \cap (y, x)|$, so we move to the third and fourth prongs of the test. We have $I \cap (y, x) = J \cap (y, x)$. We know that $a \in (I \cap (x, y)) \setminus (J \cap (x, y))$, and, because I and J are weakly separated, there can be no element of $(J \cap (x, y)) \setminus (I \cap (x, y))$ which is \leq_x -prior to a . So $I \sqsubset J$ in this case as well. \square

We make the following definition: For $1 \leq i \leq m \leq j \leq q$, we say that J_m is an (i, j) -snake if

- (1) all of the integers $\ell_i, \ell_{i+1}, \dots, \ell_m, r_m, \dots, r_{j-1}, r_j$ are in J_m and
- (2) $(r_j, \ell_i) \cap (E \setminus J_m)$ is non-empty.

Note that J_m is always an (m, m) -snake. The following lemma allows us to construct snakes.

LEMMA 10.6. *Assume that $(x, y) \setminus E$ is non-empty and that no M^a is weakly separated from \mathcal{C} .*

For any indices i and j with $1 \leq i \leq j \leq q$, let J_m be the \sqsubset -minimal witness among J_i, J_{i+1}, \dots, J_j . Then J_m is an (i, j) -snake. In particular, for any (i, j) , there is an (i, j) -snake.

Proof. Our proof is by induction on $j - i$. When $i = j$, we noted above that J_i is an (i, i) -snake.

We now consider the case that $i < j$. We first establish that $\ell_i, \ell_{i+1}, \dots, \ell_m, r_m, \dots, r_{j-1}$ and r_j are in J_m . Suppose, for the sake of contradiction, that $r_m, r_{m+1}, \dots, r_{j'-1}$ are in J_m but $r_{j'}$ is not. By induction, there is a $(m + 1, j')$ snake $J_{m'}$. Then ℓ_{m+1} and $r_{j'}$ are in $J_{m'}$ and not in J_m , while $r_{m'-1}$ is in J_m and not $J_{m'}$ by Lemma 10.4. This means that $J_{m'} \sqsubset J_m$, a contradiction. We conclude that r_m, \dots, r_{j-1} , and r_j are in J_m ; similarly, $\ell_i, \dots, \ell_{m-1}$ and ℓ_m are in J_m .

We now establish that $(r_j, \ell_i) \cap (E \setminus J_m)$ is non-empty. We know that $(r_m, \ell_m) \cap (E \setminus J_m)$ is non-empty; let $e \in (r_m, \ell_m) \cap (E \setminus J_m)$. Suppose, for the sake of contradiction, that $(r_j, \ell_i) \cap (E \setminus J_m)$ is empty, and hence $e \in (r_{j'-1}, r_{j'})$ for some j' . (The case where $e \in (\ell_{i'}, \ell_{i'+1})$ is similar.) By induction, there is an $(m + 1, j')$ -snake $J_{m'}$. Let e' be an element of $(r_{j'}, \ell_{m+1}) \cap (E \setminus J_{m'})$.

The sets Ex and $J_{m'}$ are weakly separated. Considering $(\ell_{m+1}, e, r_{j'}, e')$, we see that $e \in J_{m'}$. But then ℓ_{m+1} and e are in $J_{m'}$ and not J_m , while $r_{m'-1}$ and e' are in J_m and not $J_{m'}$. By our choice of j' such that $e \in (r_{j'-1}, r_{j'})$, we know that $r_{m'} \subset (\ell_{m+1}, e)$. So $J_{m'}$ and J_m are not weakly separated, a contradiction. \square

As a corollary, we remove the need to assume that $(x, y) \setminus E$ is non-empty.

LEMMA 10.7. *Assume that \mathcal{C} is not weakly separated from any M^a nor any N^b , and recall the standing hypothesis that E is not the interval (x, y) . Then both $(x, y) \setminus E$ and $(y, x) \cap E$ are non-empty.*

Proof. By our standing hypotheses, either $(x, y) \setminus E$ or $(y, x) \cap E$ is non-empty. Without loss of generality, assume it is the former so the (J_m, ℓ_m, r_m) are defined. Lemma 10.6 applies, and we know that there is an $(1, q)$ -snake J_m . For this J_m , we know that $(r_q, \ell_1) \cap (E \setminus J_m)$ is non-empty. Since $(r_q, \ell_1) \cap (E \setminus J_m) = (y, x) \cap (E \setminus J_m) \subseteq (y, x) \cap E$, we conclude that $(y, x) \cap E$ is also non-empty. \square

We are now entering into the final stages of the proof of Lemma 10.1. Assume, for the sake of contradiction, that \mathcal{C} is not weakly separated from any M^a nor any N^b . So we have a sequence of witnesses (J_m, ℓ_m, r_m) , and a dual sequence (J'_m, ℓ'_m, r'_m) . We point out that ℓ'_m and r'_m are in $[y, x]$; we leave it to the reader to dualize the other parts of the definition of a witness.

Let J_μ and J'_μ be the \sqsubset -minimal elements among the sequences J_i and J'_i .

LEMMA 10.8. *The set J_μ contains both x and y , and $(y, x) \cap (E \setminus J_\mu)$ is non-empty. J'_μ contains neither y nor x and $(x, y) \cap (J'_\mu \setminus E)$ is non-empty.*

Proof. The first statement is Lemma 10.6 applied to $(i, j) = (1, q)$. The second statement is the dual of the first. \square

Let $H := (x, y) \cap (J'_\mu \setminus E)$. By Lemma 10.8, H is non-empty; let $a \in H$. By definition, $a \in (x, y) \setminus E$, so there is some (J_i, ℓ_i, r_i) which is a witness against M^a . In particular, $H \cap (\ell_i, r_i)$ contains a , and is thus non-empty.

Let (J_ν, ℓ_ν, r_ν) be \sqsubset -minimal such that $H \cap (\ell_\nu, r_\nu)$ is non-empty.

LEMMA 10.9. *With the above definition, $J_\nu \sqsubset J'_\mu$.*

Proof. For the sake of contradiction, assume that $J_\nu \not\sqsubset J'_{\mu'}$.

Let h be in $H \cap (\ell_\nu, r_\nu)$. Then h is in $J'_{\mu'}$ and not in J_ν . By the contrapositive of Lemma 10.4, either $J'_{\mu'}$ contains every element of $J_\nu \cap [x, h)$, or $J'_{\mu'}$ contains every element of $J_\nu \cap (h, y]$. Without loss of generality, assume that it is the latter.

Since $y \notin J'_{\mu'}$ (Lemma 10.8), we deduce that $y \notin J_\nu$. Now, $y = r_q$. Let j be such that $r_\nu, r_{\nu+1}, \dots, r_{j-1}$ are in J_ν and r_j is not.

Let J_m be the \square -minimal element of $J_\nu, J_{\nu+1}, \dots, J_j$. By Lemma 10.6, J_m is a (ν, j) -snake. Since $r_j \notin J_\nu$, we know that $J_m \neq J_\nu$ and thus we have the strict inequality $J_m \sqsubset J_\nu$.

Since $\nu \leq m - 1 < j$, we have $r_{m-1} \in J_\nu$. Since $J'_{\mu'} \supseteq J_\nu \cap (h, y]$, we have $r_{m-1} \in J'_{\mu'}$. Also, one of the conditions on the witness $(J_{m-1}, \ell_{m-1}, r_{m-1})$ is that $r_{m-1} \notin E$. So r_{m-1} is in H .

But r_{m-1} is in (ℓ_m, r_m) and we saw above that $J_m \sqsubset J_\nu$. This contradicts the minimality of J_ν . □

We have now shown that there is some J_ν with $J_\nu \sqsubset J'_{\mu'}$. By the minimality of J_μ , we have $J_\mu \sqsubseteq J_\nu$, so $J_\mu \sqsubset J'_{\mu'}$. But the dual argument shows that $J'_{\mu'} \sqsubset J_\mu$. We have reached a contradiction, and Lemma 10.1 is proved.

11. Proof of the purity conjecture

In this section, we prove Theorem 4.8. More precisely, we prove the following theorem.

THEOREM 11.1. *Let \mathcal{I} be a Grassmann necklace, \mathcal{M} the associated positroid. Let \mathcal{C} be a subset of $\binom{[n]}{k}$. Then \mathcal{C} is a maximal weakly separated collection in \mathcal{M} if and only if \mathcal{C} is of the form $\mathcal{F}(G)$ for a plabic graph G whose decorated permutation, π_G corresponds to the Grassmann necklace \mathcal{I} .*

See Section 4 to review the definitions of positroids, Grassmann necklaces and decorated permutations.

By Theorem 9.16, if \mathcal{C} is of the form $\mathcal{F}(G)$, then \mathcal{C} is a maximal weakly separated collection in \mathcal{M} . So our goal is to prove the converse. Suppose that \mathcal{C} is a maximal weakly separated collection in \mathcal{M} .

We first reduce to the case that \mathcal{I} is connected. Suppose that $I_i = I_j$. We reuse the notations I^1, I^2 and so forth from Section 5. So, by Proposition 5.13, \mathcal{C}^1 and \mathcal{C}^2 are maximal weakly separated collections in \mathcal{M}^1 and \mathcal{M}^2 . By induction, there are reduced plabic graphs G^1 and G^2 , with $f(G^r) = \mathcal{C}^r$, and with boundary regions labeled I^1 and I^2 . Gluing these graphs along these boundary regions, we have a reduced plabic graph G with $f(G) = \mathcal{C}$.

Thus, we may now assume that \mathcal{I} is connected. Form the CW -complex $\Sigma(\mathcal{C})$. By Proposition 9.4, the map π embeds $\Sigma(\mathcal{C})$ into \mathbb{R}^2 . By Proposition 9.10, for every $J \in \mathcal{C} \setminus \mathcal{I}$, the point $\pi(J)$ is inside the curve $\pi(\mathcal{I})$. Since these are the vertices of $\Sigma(\mathcal{C})$, we see that $\Sigma(\mathcal{C})$ lies within this curve.

Thus, $\Sigma(\mathcal{C})$ is a finite polyhedral complex in \mathbb{R}^2 , with outer boundary $\pi(\mathcal{I})$. We eventually take G to be the dual graph to $\Sigma(\mathcal{C})$. To do this, we need the following proposition.

PROPOSITION 11.2. *If \mathcal{C} is a maximal weakly separated collection for the connected positroid \mathcal{M} , then $\Sigma(\mathcal{C})$ fills the entire interior of the curve $\pi(\mathcal{I})$.*

EXAMPLE 11.3. Consider the weakly separated collection

$$\mathcal{C} = (123, 234, 345, 456, 156, 126, 125, 135, 235)$$




FIGURE 11. A plabic tiling for a non-maximal \mathcal{C} .

for the Grassmann necklace $(123, 234, 345, 456, 561, 612)$. This collection only has nine elements, not ten which a maximal collection would have. Figure 11 shows the corresponding $\Sigma(\mathcal{C})$, embedded in \mathbb{R}^2 . The grey region, labeled U , is not in $\Sigma(\mathcal{C})$, but note that all four edges of its boundary are in $\Sigma(\mathcal{C})$. In the vocabulary of the proof of Proposition 11.2, we take $E = \{1, 5\}$, $x = 3$ and $y = 6$. Our strategy is to use Lemma 10.1 to show that \mathcal{C} must be weakly separated from one of 145 , which is $E \cup 4$, or 356 , which is $Exy \setminus 1$. (In fact, it is weakly separated from both.)

Proof of Proposition 11.2. Let N denote the interior of $\pi(\mathcal{I})$. Suppose, to the contrary, that U is a connected component of $N \setminus \pi(\Sigma(\mathcal{C}))$. Let $(\pi(Ex), \pi(Ey))$ be an edge in ∂U , with U on the left-hand side as we look from $\pi(Ex)$ to $\pi(Ey)$. Note that (Ex, Ey) is not of the form (I_i, I_{i+1}) as, if it were, then U would be on the outside of $\pi(\mathcal{I})$. In particular, E is not the interval (y, x) .

By Lemma 10.1, there is a set J , either of the form Ea , with $a \in (x, y) \setminus E$, or the form $Exy \setminus b$, with $b \in (y, x) \cap E$, such that J is weakly separated from \mathcal{C} . The conditions on a and b imply that the triangle $(\pi(Ex), \pi(Ey), \pi(J))$ lies on the U side of $(\pi(Ex), \pi(Ey))$. (The reader is invited to check this in Example 11.3.) If $J \in \mathcal{M}$, then J must be in \mathcal{C} , by the maximality of \mathcal{C} . But then the triangle $(\pi(Ex), \pi(Ey), \pi(J))$ is part of $\Sigma(\mathcal{C})$, and lies on the U -side of $(\pi(Ex), \pi(Ey))$, contradicting the supposition that $(\pi(Ex), \pi(Ey))$ is a boundary edge of U . So $J \notin \mathcal{M}$. (At this point, we are already done if \mathcal{M} is the uniform matroid.)

We have shown that $J \notin \mathcal{M}$. By Proposition 9.10, $\pi(J)$ lies outside of $\pi(\mathcal{I})$. In particular, the triangle $(\pi(Ex), \pi(Ey), \pi(J))$ must cross some boundary edge $(\pi(I_i), \pi(I_{i+1}))$. Let us say that segments $(\pi(Ex), \pi(J))$ and $(\pi(I_i), \pi(I_{i+1}))$ cross (the case of Ey is practically identical). Note that (Ex, J) and (I_i, I_{i+1}) are both subsets of cliques in the weakly separated collection $\mathcal{C} \cup \{J\}$. It almost contradicts Proposition 9.4 for these two segments to cross. The only subtlety is that (Ex, J) and (I_i, I_{i+1}) might be contained in the same clique of $\mathcal{C} \cup \{J\}$, with (Ex, I_i, J, I_{i+1}) occurring in that circular order around the corresponding face of $\Sigma(\mathcal{C} \cup \{J\})$. But then the triangle (Ex, I_i, I_{i+1}) is contained in $\Sigma(\mathcal{C})$, and lies on the U -side of (Ex, Ey) , contradicting that (Ex, Ey) is supposed to be on the boundary of U . \square

So, $\Sigma(\mathcal{C})$ is a disc, whose two-dimensional faces are naturally colored black and white. Let G be the dual graph, so the faces of G are labeled by \mathcal{C} , and the boundary faces by \mathcal{I} . There are n vertices on the boundary of G , labeled by $[n]$, with vertex i between faces I_i and I_{i+1} .

We now show that G is reduced. Be warned that, at this point, we do not know that the labels of the faces of G inherited from $\Sigma(\mathcal{C})$ are the labels $\mathcal{F}(G)$. Let γ be a strand which separates Ea and Ex , for some x , and then goes on to separate Ea from some Ey . As we travel

along γ , the faces on the left are of the form $E_i a$ for some sequence E_i , with $|E_i \setminus E_{i+1}| = 1$, and the faces on the right are of the form $E_i \cup E_{i+1}$.

We have $E_i <_a E_{i+1}$ due to the definition of cliques in plabic tilings. Therefore, γ cannot loop or cross itself; it must travel from the boundary of N to the boundary of N . In particular, we see that there are only n strands in total, one ending between each pair of boundary faces of G .

The strand γ must start separating $I_{\pi^{-1}(a)}$ from $I_{\pi^{-1}(a)+1}$, and end separating I_a from I_{a+1} . That is because these are the only pairs of adjacent faces at the boundary of G which differ by an a . Thus, the strand we have called γ is the strand labeled by a in our standard way of labeling the strands of a plabic graph.

Now fix two strands α and β which intersect at least twice, labeled a and b . We want to show that the intersections occur in reverse order along α and β . At one crossing, let the strands separate faces Sa and Sb ; at the other crossing let the adjacent faces be Ta and Tb . From the fact that $Sa \parallel Tb$ and $Sb \parallel Ta$, we get that one of the following holds:

- (i) $S \setminus T \subset (a, b)$ and $T \setminus S \subset (b, a)$ or
- (ii) $T \setminus S \subset (a, b)$ and $S \setminus T \subset (b, a)$.

Assuming the first case holds, we obtain $S <_a T$ and $T <_b S$. So the intersections occur in opposite order along α and β . We also conclude that the intersections occur in opposite order in the second case. We have now checked that strands do not cross themselves, do not form closed loops, and that the intersections along any pair of strands occur in opposite order along the two strands. So G is reduced.

Finally, we must check that the labels coming from $\Sigma(\mathcal{C})$ are the same as the labels $\mathcal{F}(G)$. In this paragraph, when we refer to the label of a face of G , we mean the label I of the dual vertex of $\Sigma(\mathcal{C})$. Consider a strand γ , with label a . We must check that the faces to the left of γ all contain a , and those to the right do not. The faces immediately to the left of γ all contain a , and those immediately to the right do not. Let Δ be the union of all faces that contain a , so γ is part of the boundary of Δ , as is part of ∂N . We claim that, in fact, this is the entire boundary of Δ . If there were any other edge e in the boundary of Δ , since G is a disc with boundary ∂N , there would have to be a face on the other side of e from Δ , and the label of this face would not contain a . But then one of the strands passing through this edge would be γ . Similarly, the boundary of $G \setminus \Delta$ is also made up of γ and a piece of ∂N . It is now topologically clear that Δ is the part of N on one side of γ , as desired.

So starting from a maximal weakly separated collection \mathcal{C} for \mathcal{M} , we have obtained a reduced plabic graph G with $\mathcal{F}(G) = \mathcal{C}$. This completes the proof of Theorem 11.1.

We have now proved Theorem 11.1. By Theorem 6.7, we now know that every maximal weakly separated collection of $\mathcal{M}_{\mathcal{I}}$ has cardinality $k(n - k) - \ell(\mathcal{I}) + 1$ and any two maximal weakly separated collections of \mathcal{M} are mutation equivalent. This is Theorem 4.8 and our work is complete.

Before we end, we state a simple result that follows from Theorem 11.1.

COROLLARY 11.4. *Let \mathcal{I} be a Grassmann necklace, π the associated decorated permutation, and \mathcal{M} the associated positroid. We define $\mathcal{PH}(\mathcal{I})$ to be the union of $\mathcal{F}(G)$ for all reduced plabic graphs G of \mathcal{M} . Then $J \in \mathcal{M}$ is in $\mathcal{PH}(\mathcal{I})$ if and only if it obeys the following condition: For any alignment $\{i, j\}$ in π , if $i \in J$, then we have $j \in J$.*

Proof. One direction is obvious: For any $J \in \mathcal{F}(G)$ for some plabic graph G of \mathcal{M} , it should meet the condition due to the way the faces are labeled in plabic graphs.

So we only need to check the other direction. Pick any $J \in \mathcal{M}$ that satisfy the condition. Assume for the sake of contradiction that J is not weakly separated with some I_i . This means that there is some $k \in J \setminus I_i$ and $t \in I_i \setminus J$ such that i, k, t is circularly ordered. The

condition $k \notin I_i$ tells us that $\pi^{-1}(k) \leq_i k$, while the condition $t \in I_i$ tells us that $\pi^{-1}(t) \geq_i t$. So $i, \pi^{-1}(k), k, t, \pi^{-1}(t)$ are circularly ordered and $\{k, t\}$ forms an alignment of π . Then J should contain both k and t due to the condition, but this contradicts $t \notin J$. So J is weakly separated from all of the sets I_i , and $\{I_1, \dots, I_n, J\}$ is a weakly separated collection. Then Theorem 11.1 tells us that there is some $\mathcal{F}(G)$ that contains J . \square

12. Connection with w -chamber sets of Leclerc and Zelevinsky

In this section, we state Leclerc and Zelevinsky’s conjectures on w -chamber sets and explain how they follow from Theorem 4.8.

Given two sets of integers, I and J , we write $I \prec J$ if $i < j$ for all $i \in I$ and all $j \in J$. We write \prec_i for \prec with respect to the shifted order $<_i$ on $[n]$. The following definition is necessary only in this section; it is an extension of the definition of weak separation to sets of unequal cardinality.

DEFINITION 12.1. We say that $I, J \subset [m]$ are weakly separated in the sense of Leclerc and Zelevinsky if at least one of the following holds:

- (i) $|I| \geq |J|$ and $J - I$ can be partitioned into a disjoint union $J \setminus I = J' \sqcup J''$ so that $J' \prec I \setminus J \prec J''$;
- (ii) $|J| \geq |I|$ and $I - J$ can be partitioned into a disjoint union $I \setminus J = I' \sqcup I''$ so that $I' \prec J \setminus I \prec I''$.

We write $I \parallel_{LZ} J$ to denote that I, J are weakly separated in the sense of Leclerc and Zelevinsky.

For example, $\{1, 2\} \parallel_{LZ} \{3\}$ but $\{1, 4\} \not\parallel_{LZ} \{3\}$.

DEFINITION 12.2. Let w be a permutation in the symmetric group S_m . The w -chamber set, $\mathcal{H}(w)$, is defined to be the collection of those subsets I of $[m]$ that satisfy the following condition: for every pair $a < b$ with $w(a) < w(b)$, if $a \in I$, then $b \in I$.

REMARK 12.3. The attentive reader note that our definition of a w -chamber set differs from that in [11] by switching the directions of all inequalities. Thus, the w -chamber set in our notation is the image of Leclerc and Zelevinsky’s $(w_0 w w_0)$ -chamber set under the map $i \mapsto m + 1 - i$. This flip has to be included somewhere, due to an incompatibility in the sign conventions of [11, 13].

DEFINITION 12.4. A weakly separated collection \mathcal{C} of $\mathcal{H}(w)$ is a collection of subsets in $\mathcal{H}(w)$ such that the elements of \mathcal{C} are pairwise weakly separated (in the LZ sense). A maximal weakly separated collection of $\mathcal{H}(w)$ is a weakly separated collection of $\mathcal{H}(w)$ that is maximal under inclusion.

Here is the original conjecture of Leclerc and Zelevinsky.

THEOREM 12.5. Any maximal weakly separated collection of $\mathcal{H}(w)$ is of cardinality $n + \ell(w) + 1$, where $\ell(w)$ is the length of w .

Subsets of $[m]$ of any cardinality can be identified with certain m element subsets of $[2m]$ through a simple padding construction, as follows.

DEFINITION 12.6. Given $I \subset [m]$, we define $\text{pad}(I)$ to be $I \cup \{2m, \dots, m + |I| - 1\}$. Given a collection $\mathcal{C} \subset 2^{[m]}$, we define $\text{pad}(\mathcal{C})$ to be the collection of $\text{pad}(I)$ for each $I \in \mathcal{C}$.

For example, if $m = 4$, then $\text{pad}(\{1, 3\}) = \{1, 3, 7, 8\}$.

REMARK 12.7. This combinatorial definition is motivated by geometrical constructions. LeClerc and Zelevinsky expected w chamber sets to be coordinates on a unipotent cell or (essentially equivalently) on a Bruhat cell in the flag manifold $\mathcal{F}\ell(m)$. An open subvariety of this cell is a double Bruhat cell $GL_m^{e, w_0 w}$. As explained in [10, Section 6], GL_m double Bruhat cells are isomorphic to open positroid varieties in $G(m, 2m)$. The function $\text{pad}()$ describes the correspondence between the natural indexing sets for coordinates coming from $\mathcal{F}\ell_m$ and from $G(m, 2m)$.

Let $n = 2m$ and $k = m$. The padding allows us to reduce weak separation in the sense of Leclerc and Zelevinsky to the notion of weak separation of elements in $\binom{[n]}{k}$ which we use in the rest of the paper. Specifically, we have the following lemma.

LEMMA 12.8. Given I and $J \subset [m]$, we have $I \parallel_{\text{LZ}} J$ if and only if $\text{pad}(I) \parallel \text{pad}(J)$.

Proof. If $|I| = |J|$, then $I \parallel_{\text{LZ}} J$ if and only if $I \parallel J$ if and only if $\text{pad}(I) \parallel \text{pad}(J)$. Otherwise, we may assume that $|I| > |J|$. Then $\text{pad}(J) \setminus \text{pad}(I) = (J \setminus I) \cup \{m + |J| - 1, m + |J|, \dots, m + |I| - 1\}$. Note that $J \setminus I$ is contained in $[m]$, as is $\text{pad}(I) \setminus \text{pad}(J)$. So $\text{pad}(I)$ is weakly separated from $\text{pad}(J)$ if and only if there is no element of $J \setminus I$ which lies between two elements of $I \setminus J$ in the standard linear order on $[m]$. This precisely means that we can partition $J \setminus I$ as $J' \sqcup J''$ so that $J' \prec I \setminus J \prec J''$. □

Next, we explain the relation between positroids and w -chamber sets. For a permutation $w \in S_m$, let $\mathcal{I}(\hat{w}) = (I_1, \dots, I_{2m})$ be the Grassmann necklace corresponding to the decorated permutation

$$\hat{w} := [2m, 2m - 1, \dots, m + 1, w^{-1}(m), \dots, w^{-1}(1)].$$

The permutation \hat{w} has no fixed points, so we do not need to describe a coloring.

LEMMA 12.9. For any $w \in S_m$, we have $\ell(w) + m + 1 = m^2 - \ell(\mathcal{I}(\hat{w})) + 1$. Furthermore, for $J \subseteq [m]$, we have $J \in \mathcal{H}(w)$ if and only if $\text{pad}(J) \in \mathcal{PH}(\mathcal{I}(\hat{w}))$.

Proof. Let us compute the number of alignments of \hat{w} . All $\binom{m}{2}$ pairs $\{i, j\}$ with $i, j \in [m]$ form an alignment. A pair $\{2m + 1 - i, 2m + 1 - j\}$ with $i, j \in [m]$ forms an alignment if and only if (i, j) is not an inversion of w^{-1} , that is, $i < j$ and $w^{-1}(i) < w^{-1}(j)$. There are $\binom{m}{2} - \ell(w^{-1}) = \binom{m}{2} - \ell(w)$ alignments of this type. Finally, a pair $\{i, m + j\}$ with $i, j \in [m]$ never forms an alignment. In total, we obtain $\ell(\mathcal{I}(\hat{w})) = 2\binom{m}{2} - \ell(w)$. A bit of algebra shows this is the same as $\ell(w) + m + 1 = m^2 - \ell(\mathcal{I}(\hat{w})) + 1$.

Now let us check the second claim. Any pair $a < b$ with $w(a) < w(b)$ translates into an alignment $\{a, b\}$ of \hat{w} . So $\text{pad}(J) \in \mathcal{PH}(\mathcal{I}(\hat{w}))$ implies $J \in \mathcal{H}(w)$. The alignments $\{a, b\}$ of \hat{w}

such that $a, b \in [m + 1, 2m]$ translates into the restriction that any $H \in \mathcal{PH}(\mathcal{I}(\hat{w}))$ is obtained by padding some subset of $[m]$. And as we have checked above, a pair $\{i, m + j\}$ with $i, j \in [m]$ never forms an alignment. So $J \in \mathcal{H}(w)$ implies $\text{pad}(J) \in \mathcal{PH}(\mathcal{I}(\hat{w}))$. \square

Combining the above two lemmas, we have obtained the following claim.

PROPOSITION 12.10. *A collection \mathcal{C} is a maximal weakly separated collection of a w -chamber set $\mathcal{H}(w)$ if and only if $\text{pad}(\mathcal{C})$ is a maximal weakly separated collection inside the positroid $\mathcal{M}_{\mathcal{I}(w)}$.*

So Theorem 12.5 follows immediately from Theorem 4.8.

Acknowledgements. David E. Speyer is engaged in an as yet unpublished project with Andre Henriques and Dylan Thurston, which has some overlap with this work. The main definitions of Section 9 and the formulations of the main results come from conversations with Dylan and Andre. In particular, Proposition 9.4 was originally proved in that collaboration. David E. Speyer is grateful to his collaborators for letting these ideas appear here.

References

1. F. ARDILA, F. RINCÒN and L. WILLIAMS, ‘Positroids and non-crossing partitions’, *Trans. AMS*, Preprint, 2013, arXiv:1308.2698.
2. A. BERENSTEIN and A. ZELEVINSKY, ‘Quantum Cluster Algebras’, *Adv. Math.* 195 (2005) 405–455.
3. R. BOCKLANDT, ‘Calabi–Yau algebras and weighted quiver polyhedra’, *Math. Z.* 273 (2013) 311–329.
4. N. BROOMHEAD, ‘Dimer models and Calabi–Yau algebras’, *Mem. Amer. Math. Soc.* 215 (2012).
5. T. BRYLAWSKI, ‘Constructions’, *Chapter 7 of Theory of matroids*, Encyclopedia of Mathematics and its Applications 26 (ed. Neil White; 1986).
6. V. I. DANILOV, A. V. KARZANOV and G. A. KOSHEVOY, ‘On maximal weakly separated set-systems’, *J. Algebraic Combin.* 32 (2010) 497–531.
7. N. FORD, ‘The expected codimension of a matroid variety’, Preprint, 2013, arXiv:1309.0460.
8. C. GEISS B. LECLERC and J. SCHRÖER, ‘Cluster structures on quantum coordinate rings’, *Selecta Math.* 19 (2013) 337–397.
9. J. E. GRABOWSKI and S. LAUNOIS, ‘Graded quantum cluster algebras and an application to quantum Grassmannians’, Preprint, 2013, arXiv:1301.2133.
10. A. KNUTSON, T. LAM and D. E. SPEYER, ‘Positroid varieties: juggling and geometry’, *Compos. Math.* 149 (2013) 1710–1752.
11. B. LECLERC and A. ZELEVINSKY, ‘Quasicommuting families of quantum Plücker coordinates’, *Advances in Mathematical Sciences (Kirillov’s seminar)*, AMS Translations 181 (1998) 85–108.
12. B. LECLERC, ‘Cluster structures on strata of flag varieties’, *Adv. in Math.*, Preprint 2014, arXiv:1402.4435.
13. A. POSTNIKOV, ‘Total positivity, Grassmannians, and networks’, Preprint, 2006, <http://math.mit.edu/~apost/papers/tpgrass.pdf>.
14. J. SCOTT, ‘Quasicommuting families of quantum minors’, *J. Algebra* 290 (2005) 204–220.
15. J. SCOTT, ‘Grassmannians and cluster algebras’, *Proc. London Math. Soc.* 92 (2006) 345–380.
16. R. SIMION, ‘Non-crossing partitions’, *Discrete Math.* 217 (200) 367–409.

Suho Oh
Department of Mathematics
Texas State University
601 University Drive
San Marcos, TX 78666
USA
suhooh@txstate.edu

David E. Speyer
Department of Mathematics
University of Michigan
2074 East Hall
530 Church Street
Ann Arbor, MI 48109-1043
USA
speyer@umich.edu

Alex Postnikov
Department of Mathematics
Massachusetts Institute of Technology
Building E18, Room 369
77 Massachusetts Avenue
Cambridge, MA 02139-4307
USA
apost@math.mit.edu