

MIT Open Access Articles

Weak separation and plabic graphs

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Oh, Suho, Alexander Postnikov, and David E. Speyer. "Weak Separation and Plabic Graphs." *Proceedings of the London Mathematical Society* 110.3 (2015): 721–754.

As Published: <http://dx.doi.org/10.1112/plms/pdu052>

Publisher: Oxford University Press - London Mathematical Society

Persistent URL: <http://hdl.handle.net/1721.1/106463>

Version: Original manuscript: author's manuscript prior to formal peer review

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike



WEAK SEPARATION AND PLABIC GRAPHS

SUHO OH, ALEXANDER POSTNIKOV, 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 maximal by inclusion weakly separated collections of minors have the same cardinality, and that they can be related to each other by a sequence of mutations.

On the other hand, Postnikov studied *total positivity* on the Grassmannian. He described a stratification of the totally nonnegative Grassmannian into *positroid strata*, and constructed their parametrization 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 [8], 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 [8], 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. [8] *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 [8]. Similarly, 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

Key words and phrases. Weakly separated collection, Leclerc-Zelevinsky's purity conjecture, totally positive Grassmannian, plabic graph, positroid, plabic tiling.

The second author was supported in part by NSF CAREER Award DMS-0504629.

The third author was supported in part by a Clay Research Fellowship.

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 $(I \setminus J) \prec (J \setminus I)$.

Scott [10] 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. [10] *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 will present a stronger statement, which implies both these conjectures.

In [9], a stratification of the totally nonnegative Grassmannian $Gr^{tnn}(k, n)$, 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 will 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 $|\mathcal{C}| = 1 + \ell(\pi)$.*

We will 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 [8].

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, whose special cases were conjectured in [8] and [10].

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 [9]. Plabic graphs are variants of wiring diagrams of decompositions in the symmetric group, which deserve to be better understood. 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 $\binom{[n]}{k}$.

Actually, we show that 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:*

- (1) *The set of Plücker coordinates $\{p_I\}_{I \in \mathcal{C}}$ is a cluster in the cluster algebra structure on the coordinate ring of $G(k, n)$.*
- (2) *\mathcal{C} is a maximal weakly separated collection.*
- (3) *\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 [9].

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. Like 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 [5]. 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 [10], which is motivated by geometric considerations in the case of Grassmannians. A similar notion was introduced by Leclerc and Zelevinsky, [8], motivated by considerations in the study of flag manifolds. We will refer to Leclerc and Zelevinsky’s definition as **LZ weak separation**. As we explain in section 12, LZ weak separation is actually a special case of weak separation. 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.

1.3. Acknowledgments. The third author is engaged in an as yet unpublished project with Andre Henriques and Dylan Thurston, which has some overlap with this work. In particular, Proposition 9.4 was originally proved in that collaboration; the third author is grateful to his collaborators for letting it appear here.

2. NOTATION

Throughout the paper we use the following notation. Let $[n]$ denote $\{1, \dots, n\}$, and $\binom{[n]}{k}$ denote the set of k element subsets of $[n]$. We will 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 will 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 $[n]$, 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 will denote these by lower case letters, capital letters, and calligraphic letters, respectively.

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

3. WEAKLY SEPARATED COLLECTIONS

In this section, we define weak separation for collections of k element subsets and discuss the k subset analogue of Leclerc and Zelevinsky's conjectures. The relation of this approach to the original definitions and conjectures from [8] will be discussed in section 12.

Let us fix two nonnegative 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 [8], Scott observed the following claim.

Proposition 3.2. [10], cf. [8] *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 will 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, [9] introduced the positroid stratification of the Grassmannian (see also [7]). We will prove a version of Theorem 3.3 for every cell in this stratification. Theorem 3.3 itself will correspond to the case of the largest cell.

There are several combinatorial objects which can be used to index the cells of this stratification. We will use three of these – *Grassmann necklaces*, *decorated permutations*, and *positroids*. See [9] for more details.

Definition 4.1. [9, 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: $I_1 = \{1, 2, 4\}$, $I_2 = \{2, 4, 5\}$, $I_3 = \{3, 4, 5\}$, $I_4 = \{4, 5, 2\}$, $I_5 = \{5, 1, 2\}$.

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 [9].

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 $\mathcal{M}_{\mathcal{I}}$.

We have not shown yet that $\mathcal{I} \subseteq \mathcal{M}_{\mathcal{I}}$, or that \mathcal{I} is weakly separated, so there is a risk that there are no weakly separated collections in $\mathcal{M}_{\mathcal{I}}$. We remedy this by the following lemma.

Lemma 4.4. *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 $I_i \rightarrow I_{i+1} \rightarrow \cdots \rightarrow I_r \rightarrow I_{r+1} \rightarrow \cdots \rightarrow I_j$. 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 get 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 will say, 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 4.5. [9, Definition 13.3] A decorated permutation $\pi^\cdot = (\pi, col)$ 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 $\pi^\cdot = (\pi, col)$, 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 $col(i) = 1$. Finally, if $i \in I_i = I_{i+1}$ then $\pi(i) = i$ is a fixed point of color $col(i) = -1$.

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

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

For example, the decorated permutation $\pi^\cdot = (\pi, col)$ with $\pi = 81425736$ and $col(5) = 1$ 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 [9, Theorem 17.1].

Definition 4.6. [9, Section 17] For $i, j \in [n]$, we say that $\{i, j\}$ forms an *alignment* in π if $i, \pi(i), \pi(j), j$ are cyclically ordered (and all distinct). The *length* $\ell(\pi^\cdot)$ is defined to be $k(n-k) - A(\pi)$ where $A(\pi)$ is 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} .

We now state our result for an arbitrary Grassmann cell.

Theorem 4.7. *Fix any Grassmann necklace \mathcal{I} . Every maximal weakly separated collection inside $\mathcal{M}_{\mathcal{I}}$ has cardinality $\ell(\mathcal{I}) + 1$. 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° , \mathcal{I} , and \mathcal{M} are certain decorated permutations, Grassmann necklaces, and positroids, whose ground sets may no longer be $[n]$ but rather subsets of $[n]$, which inherit their circular order from $[n]$.

Example 5.1. The arrows in the left hand side of Figure 10 form a decorated permutation with $(k, n) = (5, 10)$. The connected components are 019, 2378 and 456. The 5 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, will have her curiosity satisfied 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 **noncrossing** if, for any circularly ordered (a, b, c, d) , we have $\{a, c\} \subseteq S_i$ and $\{b, d\} \subseteq S_j$ then $i = j$.

See, for example, [12] for background on noncrossing partitions. The following is obvious from the definition:

Proposition 5.3. *The common refinement of two noncrossing partitions is noncrossing.*

Therefore, the following definition makes sense.

Definition 5.4. Let π° be a decorated permutation. Let $[n] = \bigsqcup S_i$ be the finest noncrossing partition of $[n]$ such that, if $i \in S_j$ then $\pi(i) \in S_j$.

Let $\pi_{(j)}^\circ$ be the restriction of π° 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 $\pi_{(j)}^\circ$ the *connected components* of π° , and $\mathcal{I}_{(j)}$ the *connected components* of \mathcal{I} .

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

Lemma 5.5. *The decorated permutation π° 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 noncrossing partition preserved by π . So there is a nontrivial noncrossing partition preserved by π and π is not connected. Conversely, any nontrivial noncrossing 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° (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° 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. In order 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 will explain how to reduce computations about positroids to the connected case.

So, 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 will also write $n^1 = |[i, j)|$ 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; see [4, Theorem 7.6.4] for the fact that this rank condition implies that \mathcal{M} is the direct sum of $\mathcal{M}|_{[i, j)}$ and $\mathcal{M}|_{[j, i)}$. See [4, Proposition 7.6.1] for the interpretation of 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. *\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 $J \cup I^2 \geq_k I_k$ for every $k \in [n]$. This immediately implies that $J \geq_k I_k^1$ 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 $J \cup I^2 \geq_k I_k$ for all $k \in [i, j)$. For $k \in [j, i)$, we know that $I^2 \geq_k I_k^2$ and, since \geq_k and \geq_i coincide on $[i, j)$, we know that $J \geq_k I^1$. So $J \cup I^2 \geq_k I_k$ for k in $[j, i)$ as well. So $J \in \mathcal{M}$. \square

It is easy to check the following:

Proposition 5.11. *With the above definitions, $\ell(\mathcal{I}) = \ell(\mathcal{I}^1) + \ell(\mathcal{I}^2)$.*

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. We have partially proved this in this section. In order to fully establish this result we would also need to show that, if π° is connected, then the matroid \mathcal{M} is connected. We don't need this fact, but it is true.

6. PLABIC GRAPHS

Plabic graphs were introduced in [9]. We will use these graphs to represent weakly separated collections and the properties of these graphs will be a major key in our proof of the conjecture. For more details on plabic graphs, see [9].

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.





FIGURE 1. Labeling faces of a plabic graph

Let G be a plabic graph in the disc D . We will 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 that figure will be explained below. 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 towards each endpoint. As we will 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 will 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

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

do not restrict in this paper to bipartite graphs for two reasons: (1) We want to be compatible with the definitions in [9], 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* [9, 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 .
- (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 [2, Section 8] and are called “marginally geometrically consistent” in [3, Section 3.4].

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

Definition 6.5. [9, 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 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: $col(i) = 1$ if the corresponding loop is counter-clockwise; and $col(i) = -1$ if the loop is clockwise. In this way, we assign the *decorated strand permutation* $\pi_G^c = (\pi_G, col)$ to each reduced plabic graph G .

We will label the faces of a reduced plabic graph with subsets of $[n]$; this construction was first published in [10]. 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 will eventually play 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 and the corresponding Grassmann necklace $\mathcal{I} = \mathcal{I}(\pi)$, 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 .*

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 $w = [k + 1, k + 2, \dots, n, 1, 2, \dots, k]$.

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 3 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.




FIGURE 3. (M1) Square move

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




FIGURE 4. (M2) Unicolored edge contraction

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

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 the next Theorem:




FIGURE 5. (M3) Vertex removal

Theorem 6.7. [9, 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:*

- G can be obtained from G' by moves (M1)-(M3).
- 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. So, once we prove Theorem 6.6, we will establish that all maximal weakly separated collections within $\mathcal{M}(\mathcal{I})$ are connected to each other by mutations.

Notice 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 [9]:

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

Proof. By [9, Proposition 17.10], $\ell(\pi)$ is the dimension of $S_{\mathcal{M}}^{t\pi n}$, a manifold whose definition we do not need to know. By [9, Theorem 12.7], $S_{\mathcal{M}}^{t\pi n}$ is isomorphic to $\mathbb{R}^{|F(G)|-1}$. \square

Thus, proving Theorem 6.6 will establish all parts of 4.7, 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 in [11]. 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:*

- (1) *The set of Plücker coordinates $\{p_I\}_{I \in \mathcal{C}}$ is a cluster in the cluster algebra structure on the coordinate ring of $G(k, n)$.*
- (2) *\mathcal{C} is a maximal weakly separated collection.*
- (3) *\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.

The implication (2) \implies (3) is the main result of this paper; the implication (3) \implies (1) is [11, Theorem 2]. We now show (1) \implies (3).

Let I and $J \in \mathcal{C}$; we will show that $I \parallel J$. Let $\mathbb{C}_q(G(k, n))$ be the quantization of the coordinate ring of $G(k, n)$ introduced in [1]; since p_I and p_J are in a common cluster, they quasi-commute in this ring. By [6], the ring $\mathbb{C}_q(G(k, n))$ is isomorphic to the ring of “quantum minors” of a $k \times (n - k)$ matrix (see [10] or [8]), so the

quantum minors $\Delta(I)$ and $\Delta(J)$ quasi-commute. As computed in [10], if $\Delta(I)$ and $\Delta(J)$ quasi-commute, then 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 [10]. \square

We expect that an analogous statement holds for all positroids. We have not proved it here for two reasons. The first is that, although positroid varieties are widely expected to have a cluster structure, this has not yet been verified in print. The second is that, since the coordinate ring of a positroid variety is a quotient of the coordinate ring of $G(k, n)$, it is possible that two minors which do not quasi-commute on $G(k, n)$ do quasi-commute when restricted to this smaller subvariety. Deriving the analogues of [10] and [8] for positroid varieties strikes us as an excellent project.

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 will 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.


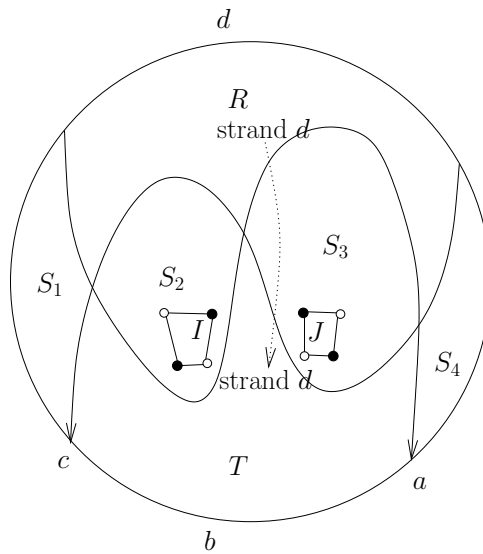
Proposition 8.1. *Let G be a reduced plabic graph. Then $\mathcal{F}(G)$ is a weakly separated collection.*

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 don't 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 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 decreasing order. Let face I be in S_i and J be in S_j ; note that $i \neq j$. We will 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

FIGURE 6. The case where strands a and c don't crossFIGURE 7. The case where strands a and c do cross

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.

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




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

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

We now use Proposition 8.1, together with Theorem 6.7, to demonstrate some facts about the face labels of plabic graphs.

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

Proposition 8.2. *Let G be a reduced plabic graph with decorated permutation π° such that $\ell(\pi^\circ) > 0$. There is some pair of strands (i, j) in π° 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 [9, Lemma 18.9], a *removable* edge has these properties. (See [9] for definitions.) According to [9, Corollary 18.10], the removable edges of G are in bijection with the decorated permutations covered by π° . Since $\ell(\pi^\circ) > 1$, it covers at least one decorated permutation. \square

Proposition 8.3. *If G is a reduced plabic graph with decorated permutation π° , then the following properties hold:*

- (1) *The boundary cells of G are labeled by $\mathcal{I}(\pi^\circ)$.*
- (2) *Every face of G receives a separate label in $\mathcal{F}(G)$.*
- (3) *$\mathcal{F}(G)$ is contained in $\mathcal{M}_{\mathcal{I}(\pi^\circ)}$.*

Proof. We are going to use induction on $\ell(\pi^\circ)$. When $\ell(\pi^\circ) = 0$, 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 Proposition 8.2. By induction hypothesis, all the properties hold for G' . We will use μ° 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° to π° 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 $\mathcal{I}(\pi^\circ)$.

Now we check the second property. Since i and j only cross once, if two faces are assigned the same label in G , that 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 get $j_t \notin I_i$, and hence $\pi^{-1}(j_t) \leq_i j_t$. And for 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 will be 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 will define a **plabic tiling**. For a weakly separated collection \mathcal{C} , we will construct a 2-dimensional CW-complex embedded in \mathbb{R}^2 , and denote this complex by $\Sigma(\mathcal{C})$. Maximal plabic tilings will turn out to be dual to reduced bipartite plabic graphs. We will 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 will 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 **nontrivial** if it has at least three elements. Observe that, if \mathcal{X} is a nontrivial 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 $\{L \setminus b_1, L \setminus b_2, \dots, L \setminus b_s\}$, with the b_i 's cyclically ordered. If $\mathcal{W}(K)$ is nontrivial, we define the boundary of $\mathcal{W}(K)$ to be the cyclic graph

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




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

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

$$(L \setminus b_1) \rightarrow (L \setminus b_2) \rightarrow \cdots \rightarrow (L \setminus b_s) \rightarrow (L \setminus b_1).$$

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 nontrivial, 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 i and j occur among the a 's, and among the 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 nonempty, 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 nontrivial, so at least one of S_1 and S_2 is nonempty. Similarly, at least one of S_3 and S_4 is nonempty.

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 the S_i 's are empty. Checking the 6 possibilities, the ones 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 $\Sigma(\mathcal{C})$ will be the elements of \mathcal{C} . There will be an edge (I, J) if

- (1) $\mathcal{W}(I \cap J)$ is nontrivial and (I, J) appears in the boundary of $\mathcal{W}(I \cap J)$ or
- (2) $\mathcal{B}(I \cup J)$ is nontrivial and (I, J) appears in the boundary of $\mathcal{B}(I \cup J)$ or
- (3) $\mathcal{W}(I \cap J) = \mathcal{B}(I \cup J) = \{I, J\}$.¹

There will be a (two-dimensional) face of $\Sigma(\mathcal{C})$ for each nontrivial clique \mathcal{X} of \mathcal{C} . The boundary of this face will be 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 will 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 will 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 nonnegative for a in $[\ell, r)$ and is nonpositive for a in $[r, \ell)$. We have

$$(1) \quad \pi(e) - \pi(f) = \sum_{a \in [\ell, r)} (e_a - f_a)v_a - \sum_{b \in [r, \ell)} (f_b - e_b)v_b.$$

Since $\sum_{a \in [n]} e_a = \sum_{a \in [n]} f_a$, 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 v_a 's 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 nonzero, and we deduce that $\pi(e) \neq \pi(f)$. \square

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 $\pi(e_I)$ 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 .*

¹This 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.

Proof. Suppose, to the contrary, that there are two faces of $\Sigma(\mathcal{C})$ whose interiors have overlapping image. We will 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 will 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.² So we may assume that $K \setminus L$ is nonempty. Also, $|L| = k + 1 > |K| = k - 1$, so $L \setminus K$ is nonempty. 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) , 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 $a \geq_x b$. 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 will study the construction $\Sigma(\mathcal{C})$ when \mathcal{C} is a weakly separated collection for a particular positroid \mathcal{M} . Propositions 9.7 through 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:

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.

²The 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.

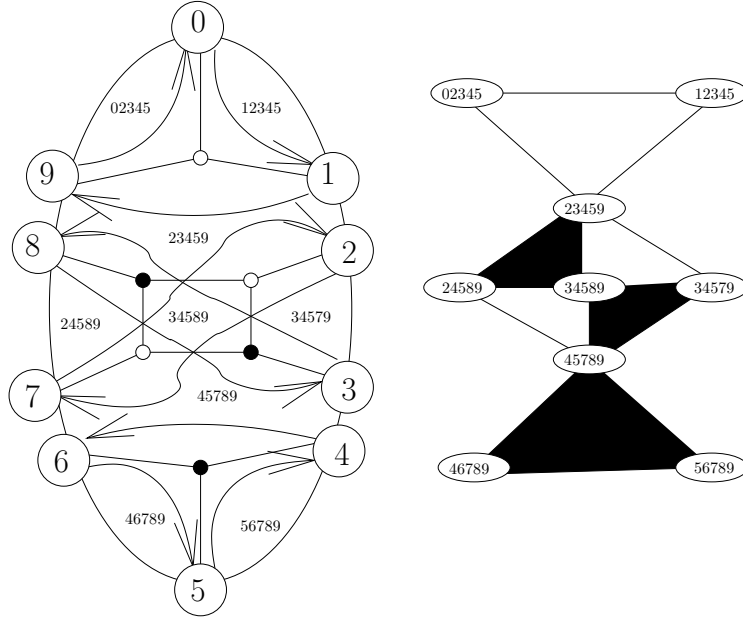


FIGURE 10. A disconnected plabic graph and the corresponding plabic tiling

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, $I_i \neq I_j$ so, by Lemma 9.3, the vertices $\pi(I_i)$ and $\pi(I_j)$ are disjoint. 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 4.4. 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 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} .*

Let $\phi(t)$ be the parameterization of the simple closed curve $\pi(\mathcal{I})$, with $\phi(r/n) = \pi(I_r)$, and $\phi(r/n + u) = (1 - u)\pi(I_r) + u\pi(I_{r+1})$ for $u \in (0, 1/n)$. 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 will 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 nonempty. 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's 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. \square

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 $y(t)$, travelling 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 $J \not\subseteq M$. 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:

- (1) the minimal element of $J \setminus I_s$ is less than the minimal element of $I_s \setminus J$ or
- (2) 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. \square

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.

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, homomorphic 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 $\psi(D(G))$. 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 $J^1 \cup I^2$ 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 LEMMA 10.1

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} :*

- (1) $M^a := E \cup \{a\}$ with $a \in (x, y) \setminus E$ or
- (2) $N^b := E \cup \{x, y\} \setminus \{b\}$ 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:*

- (1) ℓ and r are in $J \setminus E$
- (2) $x \leq_x \ell <_x r \leq_x y$
- (3) $(r, \ell) \cap (E \setminus J) \neq \emptyset$.
- (4) $(\ell, r) \cap (J \setminus E) = \emptyset$.
- (5) $a \in (\ell, r)$

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 nonnegative to negative or nonpositive 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 nonempty, 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,*

- (1) $J_m \in \mathcal{C}$
- (2) ℓ_m and r_m are in $J_m \setminus E$
- (3) $x \leq_x \ell_m <_x r_m \leq_x y$
- (4) $(r_m, \ell_m) \cap (E \setminus J_m) \neq \emptyset$.
- (5) $(\ell_m, r_m) \cap (J_m \setminus E) = \emptyset$.

and such that $x = \ell_1 <_x \ell_2 <_x \cdots <_x \ell_q$ and $r_1 <_x r_2 <_x \cdots <_x r_q = y$ 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) will obey the numbered conditions above.

Choose a minimal collection $(J_1, \ell_1, r_1), \dots, (J_q, \ell_q, r_q)$, obeying the numbered conditions, 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 nonempty.) 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 $\ell_1 <_x \ell_2 <_x \cdots <_x \ell_q$ and $r_1 <_x r_2 <_x \cdots <_x r_q = y$.

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 nonempty and that no M^a is weakly separated from \mathcal{C} . For clarity, we will 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 inconclusive³, 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.

³Which decision we make in this case is completely immaterial to the proof; we just made a decision so that the order would be total.

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 nonempty.

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 nonempty 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 $i = j$, the result is clear.

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 nonempty. We know that $(r_m, \ell_m) \cap (E \setminus J_m)$ is nonempty; 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 $(r_{j'}, \ell_{m+1}) \cap (E \setminus J_{m'})$.

E_x 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 nonempty.

Proposition 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 nonempty.*

Proof. By our standing hypotheses, either $(x, y) \setminus E$ or $(y, x) \cap E$ is nonempty. 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 nonempty. 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 nonempty. \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 $J'_{\mu'}$ be the \sqsubset -minimal elements among the sequences J_i and J'_i .

Lemma 10.8. *J_μ contains both x and y , and $(y, x) \cap (E \setminus J_\mu)$ is nonempty. $J'_{\mu'}$ contains neither y nor x and $(x, y) \cap (J'_{\mu'} \setminus E)$ is nonempty.*

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 nonempty; 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 nonempty.

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

Proposition 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 \sqsubset -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 . \square

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. \square

11. PROOF OF THE PURITY CONJECTURE


In this section, we prove Theorem 4.7. More precisely, we prove:

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

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

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 will eventually take G to be the dual graph to $\Sigma(\mathcal{C})$. In order 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

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

for the Grassmann necklace $(123, 234, 345, 456, 156, 126)$. This collection only has 9 elements, not 10 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 lefthand 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 $\pi(\mathcal{T})$. In particular, the triangle $(\pi(Ex), \pi(Ey), \pi(J))$ must cross some boundary edge $(\pi(I_i), \pi(I_{i+1}))$. Let's 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 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 can not 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:

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

Assuming the first case holds, we get $S <_a T$ 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. \square

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 $\ell(\mathcal{I}) + 1$ and any two maximal weakly separated collections of \mathcal{M} are mutation equivalent. This is Theorem 4.7 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 I_i 's, 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 4.7.

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 will write \prec_i for \prec with respect to the shifted order $<_i$ on $[n]$. The following definition will only be necessary 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:

- (1) $|I| \geq |J|$ and $J - I$ can be partitioned into a disjoint union $J \setminus I = J' \sqcup J''$ so that $J' \prec I \setminus J \prec J''$.
- (2) $|J| \geq |I|$ 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 will write $I \parallel_{\text{LZ}} J$ to denote that I, J are weakly separated (in the sense of Leclerc and Zelevinsky).

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

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 will notice that our definition of a w -chamber set differs from that in [8] 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 $i \mapsto n + 1 - i$. This flip has to be included somewhere, due to an incompatibility in the sign conventions of [8] and [9].

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 $\text{pad}(\{1, 3\}) = \{1, 3, 7, 8\}$.

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:

Lemma 12.7. *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 $J' \prec I \setminus J \prec J''$. \square

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 12.8. *For any $w \in S_m$, we have $\ell(w) + m + 1 = \ell(\mathcal{I}(\hat{w})) + 1$. Furthermore, for $J \subseteq [m]$, we have $J \in \mathcal{H}(w)$ if and only if $\mathbf{pad}(J) \in \mathcal{PH}(\mathcal{I}(\hat{w}))$.*

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 get $\ell(\mathcal{I}(\hat{w})) = m^2 - A(\hat{w}) = m^2 - 2\binom{m}{2} + \ell(w) = m + \ell(w)$.

Now let's check the second argument. Any pair $a < b$ with $w(a) < w(b)$ translates into an alignment $\{a, b\}$ of \hat{w} . So $\mathbf{pad}(J) \in \mathcal{PH}(\mathcal{I}(\hat{w}))$ 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 $\mathbf{pad}(J) \in \mathcal{PH}(\mathcal{I}(\hat{w}))$. \square

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

Proposition 12.9. *\mathcal{C} is a maximal weakly separated collection of a w -chamber set $\mathcal{H}(w)$ if and only if $\mathbf{pad}(\mathcal{C})$ is a maximal weakly separated collection inside the positroid $\mathcal{M}_{\mathcal{I}(w)}$.*

So Theorem 12.5 follows immediately from Theorem 4.7.

REFERENCES

- [1] Arkady Berenstein and Andrei Zelevinsky. Quantum Cluster Algebras. *Advances in Mathematics* **195** no. 2, 405 – 455, (2005).
- [2] Raf Bocklandt. Calabi Yau algebras and weighted quiver polyhedra. [arXiv:0905.0232](https://arxiv.org/abs/0905.0232)
- [3] Nathan Broomhead. Dimer models and Calabi-Yau algebras. [arXiv:0901.4662](https://arxiv.org/abs/0901.4662)
- [4] Thomas Brylawski. Constructions. Chapter 7 of *Theory of Matroids*, Encyclopedia of Mathematics and its Applications Volume 26, edited by Neil White, 1986.
- [5] Vladimir I. Danilov, Alexander V. Karzanov, and Gleb A. Koshevoy. On maximal weakly separated set-systems, 2009. *Journal of Algebraic Combinatorics*, **32**, 497–531, 2010.
- [6] Christof Geiss, Bernard Leclerc and Jan Schröer. Cluster structures on quantum coordinate rings, 2011. [arXiv:1104.0531](https://arxiv.org/abs/1104.0531)
- [7] Allen Knutson, Thomas Lam, and David E. Speyer. Positroid varieties: juggling and geometry, 2009. [arXiv:0903.3694](https://arxiv.org/abs/0903.3694)
- [8] Bernard Leclerc and Andrei Zelevinsky. Quasicommuting families of quantum Plücker coordinates. In *Advances in Math. Sciences (Kirillov's seminar)*, *AMS Translations*, **181** 85–108, 1998.
- [9] Alexander Postnikov. Total positivity, grassmannians, and networks, 2006. <http://math.mit.edu/~apost/papers/tpgrass.pdf>
- [10] Joshua Scott. Quasi-commuting families of quantum minors. *Journal of Algebra*, **290(1)**, 204 – 220, 2005.
- [11] Joshua Scott. Grassmannians and cluster algebras. *Proceedings of the London Mathematical Society*, **92(3)**, 345–380, (2006).
- [12] Rodica Simion. Noncrossing partitions *Discrete Mathematics* **217** no. 1–3, 367–409, 2000.

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF MICHIGAN, ANN ARBOR, MI 48109

DEPARTMENT OF MATHEMATICS, MIT, CAMBRIDGE, MA 02139

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF MICHIGAN, ANN ARBOR, MI 48109