Cluster algebras and classical invariant rings

by

Kevin Carde

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Mathematics)
in the University of Michigan
2014

Doctoral Committee:

Professor Sergey Fomin, Chair
Professor Gordon Belot,
Professor Thomas Lam,
Professor Karen E. Smith,
Professor David E. Speyer

ACKNOWLEDGEMENTS

I would first like to thank my advisor, Sergey Fomin, for his support these past several years. Besides introducing me to this beautiful corner of combinatorics and providing mathematical guidance throughout, he has been unfailingly supportive even of my pursuits outside of research. I am tremendously grateful on both counts.

I would also like to thank David Speyer, whose insight and interest in my work has been one of the best parts of my time at the University of Michigan. I am also grateful to him for suggesting to me a much shorter proof of the irreducibility of mixed invariants. His suggestion, rather than my longer, much less elegant original effort, appears as the end of the proof of Proposition 4.1.

Besides these two major mathematical inspirations, I would like to thank the University of Michigan math department at large. From the wonderful people of the graduate office whose patience and care kept me on track on a number of administrative matters, to the precalculus and calculus coordinators I've worked with who helped spark in me a love of teaching, to the professors whose friendliness and enthusiasm come through both in and out of the classroom, to the graduate students who are always willing to work together through difficult problems (mathematical or otherwise!), everyone I've met here has contributed to the warm, collaborative environment of the department. I couldn't have asked for a better place to study.

Finally, I would like to thank my family and friends, who have made the best of times so much better and the tougher times so much more bearable.

TABLE OF CONTENTS

LIST OF FIGURES

Figure

ABSTRACT

Let V be a k-dimensional complex vector space. The Plücker ring of polynomial $SL(V)$ invariants of a collection of n vectors in V can be alternatively described as the homogeneous coordinate ring of the Grassmannian $\mathrm{Gr}(k,n)$. In 2003, using combinatorial tools developed by A. Postnikov, J. Scott showed that the Plücker ring carries a cluster algebra structure. Over the ensuing decade, this has become one of the central examples of cluster algebra theory.

In the 1930s, H. Weyl described the structure of the "mixed" Plücker ring, the ring of polynomial $SL(V)$ invariants of a collection of n vectors in V and m covectors in V^*. In this thesis, we generalize Scott's construction and Postnikov's combinatorics to this more general setting. In particular, we show that each mixed Plücker ring carries a natural cluster algebra structure, which was previously established by S. Fomin and P. Pylyavskyy only in the case $k = 3$. We also introduce mixed weak separation as a combinatorial condition for compatibility of cluster variables in this cluster structure and prove that maximal collections of weakly separated mixed subsets satisfy a purity result, a property proved in the Grassmannian case by Oh, Postnikov, and Speyer.

CHAPTER 1

Introduction

Cluster algebras, first introduced by Fomin and Zelevinsky [3] over ten years ago, arose as a potential framework for studying canonical bases in quantum groups. While this application is still a work in progress, cluster algebras as a field of mathematics has exploded over the last decade, with connections to integrable systems, total positivity, Teichmüller theory, and even aspects of physics. Moreover, identifying a classically-studied ring as a cluster algebra gives a wealth of new structure to it, often with rich associated combinatorics that suggests further connections and directions.

A cluster algebra is generated by recursively-defined elements called *cluster variables*. We start with an initial *seed*, consisting of a *cluster* of cluster variables along with additional data (usually in the form of a quiver) specifying how to *mutate* the seed to form new seeds. Clusters in these new seeds are created from the old cluster by replacing one of the old cluster variables by a rational function in the old cluster variables. This process generates a (usually infinite) pattern of seeds connected by mutations, and the cluster algebra itself is the algebra generated by all cluster variables in all seeds connected to the initial seed.

Algebraic geometry has proved to be a rich source of rings which turn out to be cluster algebras. In one of the first major such successes, Scott [12] showed that the homogeneous coordinate ring of any Grassmann manifold carries a beautiful cluster algebra structure, intricately related to the combinatorics of plabic graphs developed by Postnikov in his study of totally nonnegative Grassmannians [11]. While these plabic graphs describe only a small portion of the cluster structure, the geometric properties of the Grassmannian guarantee that its homogeneous coordinate ring coincides with the resulting cluster algebra.

Once a cluster structure is established and some of its cluster variables are described, it is natural to ask if there is a purely combinatorial description of compatibility; i.e., of when two cluster variables can appear together in the same cluster. In the case of a Grassmannian, compatible Plücker coordinates exhibit a combinatorial property on subsets of $[n] := \{1, 2, \ldots, n\}$ called *weak separation*. The existence of a cluster structure on the Grassmannian brought about the *purity conjecture*, which states that all maximal collections of pairwise weakly separated subsets of $[n]$ have the same cardinality. This conjecture goes back to work by Leclerc and Zelevinsky [8] who suggested that purity holds for weakly separated collections of w-chamber subsets; for Grassmannians, it was explicitly formluated by Scott in the study of quasi-commuting quantum minors [13]. Danilov, Karzanov, and Koshevoy [1] proved Leclerc–Zelevinsky's conjecture; independently, Oh, Postnikov, and Speyer [9] proved Scott's conjecture by showing that all maximal collections do indeed correspond to plabic graphs.

Since Scott's result, cluster structures have been identified in coordinate rings of more general partial flag varieties [4], as well as quantizations of these rings [6]. In a different direction, Fomin and Pylyavskyy [2] suggested that Scott's work can be generalized to classical rings of invariants of other kinds. Viewing the homogeneous coordinate ring of the Grassmannian as the classical ring of SL_k vector invariants, it is natural to ask whether rings of *mixed invariants* of SL_k, that is, polynomial invariants of SL_k actions on collections of vectors and covectors, also carry a cluster structure. These invariant rings were studied comprehensively almost a century ago by Hermann Weyl, who succeeded in describing them completely in terms of generators and relations [14]. Fomin and Pylyavskyy constructed cluster structures on all SL_3 mixed invariant rings, often constructing more than one cluster structure on the same ring. Instead of plabic graphs, their constructions depend heavily on the diagrammatic calculus of webs developed by Kuperberg [7]. Unfortunately, there is no adequate counterpart to this calculus beyond SL_3; consequently, the approach of Fomin and Pylyavskyy has not been extended to higher dimensions.

In this thesis, we demonstrate that, for any k, each ring of SL_k invariants of at least k vectors and k covectors carries a natural structure of a cluster algebra. Rather than using webs, we return to plabic graphs as the underlying combinatorial foundation, introducing a generalization called *mixed plabic graphs* which correspond to special clusters. Just as Plücker coordinates provide a special set of cluster variables in the Grassmannian which form clusters corresponding to plabic graphs, we introduce a set of special cluster variables in mixed invariant rings; these special generators form the clusters corresponding to mixed plabic graphs. Similarly to the case of the Grassmannian, even though we have virtually no understanding of the rest of the (usually infinitely many) clusters, geometric properties of mixed invariant rings allow us to prove in this thesis that the resulting cluster algebra coincides with the invariant ring at hand.

Once the cluster structure is established, we introduce *mixed weak separation*, a purely combinatorial criterion for compatibility of the special cluster variables arising in mixed plabic graphs. This suggests an analogue of the purity conjecture stating that all maximal collections of weakly separated mixed subsets have the same cardinality. Analogous to Oh, Postnikov, and Speyer, we prove this result by constructing a mixed plabic graph corresponding to any such maximal collection.

Several proofs in this thesis generalize (and are sometimes similar to) the proofs in the existing literature for the ordinary (not mixed) setting. Nevertheless, we provide complete proofs here as well as references. Many of the original proofs use slightly different language, and none of them are written from the perspective of the mixed setting. Including proofs that utilize this perspective hopefully serves to reinforce the language and ideas developed in this thesis while making it as self-contained as possible. The main exception to the rule outside Chapter 2 is Corollary 5.16, borrowed from [9] without reproducing the detailed construction of a plabic graph as a consequence of the key lemma. Chapter 3 also uses a few results from Postnikov's work on plabic graphs [11].

While this thesis generalizes the cluster structures on Grassmannians and associated combinatorics to mixed setting, many open questions remain. No characterization of all cluster variables is known even in the Grassmannian case, let alone mixed invariant rings. (Fomin and Pylyavskyy [2] do conjecture such a characterization for SL_3.) The combinatorics of mixed plabic graphs suggests further generalizations and further directions of study; some of them are outlined in Chapter 6.

This thesis is organized as follows:

Chapter 2 is a brief overview of the background necessary for this thesis, with references to more complete treatments.

Chapter 3 introduces *mixed plabic graphs*, our primary combinatorial tool for establishing cluster structures on mixed invariant rings. The definition of a mixed plabic graph in Definition 3.2 is difficult to check directly, so Theorem 3.9 establishes easily verifiable criteria in terms of the *threading* of the graph, discussed in Section 3.3. We then prove (Theorem 3.11) that two mixed plabic graphs are move-equivalent if and only if they have the same trip permutation, analogous to [11, Theorem 13.4] in the ordinary, unmixed setting. Starting in Section 3.6, we restrict our attention to mixed plabic graphs with a certain trip permutation called the *mixed Grassmann permutation*. We introduce a special L-shaped graph and prove that it is a mixed plabic graph whose trip permutation is the mixed Grassmann permutation (Theorem 3.20). This chapter ends with a characterization of the possible face labels of mixed plabic graphs for the mixed Grassmann permutation (Theorem 3.22).

Chapter 4 is devoted to proving that each ring of mixed SL_k invariants carries a cluster structure, provided we have at least k vectors and k covectors. Analogous to Scott's strategy for the Grassmannian [12], we establish a correspondence between face labels in mixed plabic graphs for the mixed Grassmann permutation and certain special invariants which generate the invariant ring. As a result, each mixed plabic graph for the mixed Grassmann permutation gives rise to a seed whose cluster variables are special invariants. We then show that *combinatorial exchanges* in mixed plabic graphs arising from square moves coincide with *algebraic exchanges* arising from cluster algebra exchange relations (Theorem 4.6). The L-shaped mixed plabic graph from Theorem 3.20 then provides us with a special seed to which the geometric criterion of Proposition 2.2 applies, allowing us to prove that the cluster algebra generated by this seed is precisely the invariant ring at hand. After the proof, we give an explicit example of an initial seed giving rise to a cluster structure on a ring of mixed invariants (Example 4.8).

Chapter 5 introduces and studies the purely combinatorial condition of *mixed weak* *separation*. The goal of this chapter is to prove that every maximal (by inclusion) collection of pairwise weakly separated mixed subsets arises as the collection of face labels of a mixed plabic graph for the Grassmann permutation, and hence every such maximal collection has the same cardinality (Theorem 5.5). To do so, we begin by adapting Oh, Postnikov, and Speyer's key lemma [9, Lemma 10.1] to the mixed setting (Lemma 5.13). Section 5.4 discusses using this lemma to extend a maximal collection of pairwise weakly separated mixed subsets to a larger collection consisting of the face labels of three disjoint plabic graphs. Armed with the results of Section 5.4, we finally prove Theorem 5.5 in Section 5.5 by cutting these three plabic graphs apart and gluing together the pieces to form a mixed plabic graph whose collection of face labels is the desired maximal collection.

Finally, Chapter 6 briefly addresses some further directions of study, outlining a notion of total positivity in these invariant rings and sketching potential further generalizations of mixed plabic graphs and mixed weak separation.

CHAPTER 2

Background

2.1 Plabic Graphs

All definitions and results in this section come from Posnikov's paper [11].

Definition 2.1. A *plabic graph* (short for planar bicolored graph) is a planar graph drawn in a disk whose vertices are colored either white or black. Edges incident to the boundary are called *boundary edges*, which we label clockwise as $b_1, b_2, \ldots, b_n$ (often abbreviated just $1, 2, \ldots, n$).

Starting on any boundary edge b_i, we define the ith *trip* as follows. We start by proceeding into the graph along b_i and then we continue according to the "rules of the road": whenever we encounter a white vertex, we turn left, continuing along the edge immediately clockwise from our current edge, and similarly, whenever we encounter a black vertex, we turn right. Eventually we exit the graph along a boundary edge b'_i. This establishes a permutation $b_i \mapsto b'_i$ called the *trip permutation* associated to the plabic graph.

The following three types of local *moves* on plabic graphs do not change the trip permutation:

1. The square move (M1), which flips the colors of the vertices of a square in our plabic graph (see Figure 2.1).

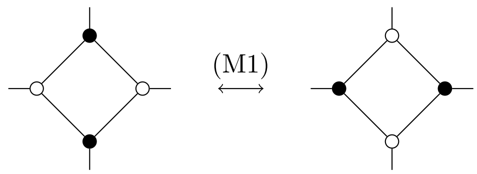

Figure 2.1: The square move (M1).

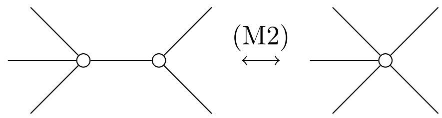

Figure 2.2: Monochromatic edge expansion/contraction (M2).

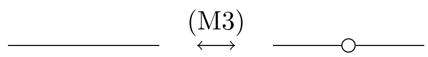

Figure 2.3: Vertex addition/removal (M3).

2. Monochromatic edge expansion/contraction (M2), which expands or contracts an edge along vertices of the same color (see Figure 2.2).
3. Vertex addition/removal (M3), which adds or removes a degree 2 vertex in the interior of any edge (see Figure 2.3).

These moves partition plabic graphs into *move-equivalence classes*. It is easy to see that two move-equivalent plabic graphs have the same trip permutation; the converse, that two plabic graphs with the same trip permutation are move-equivalent, is also true.

Finally, we label faces of plabic graphs with subsets of $[n] := \{1, 2, \ldots, n\}$ as follows. Every trip $b_i \to b'_i$ splits the disk into two regions; any face lying on the left side of this path receives i as part of its label. For any given plabic graph P, it turns out that every face label of P has the same cardinality, called the *rank* of P. The rank is unchanged by the moves, and hence depends only on the trip permutation.

Let π_n^k be the *Grassmann permutation* mapping each i to $i+k$ (modulo n). Plabic graphs whose trip permutation is π_n^k have rank k, and every possible k-element subset of $[n]$ occurs as a face label in some plabic graph with this trip permutation.

2.2 Invariant Theory

Let V be a k-dimensional vector space over $\mathbb{C}$. The special linear group $SL(V) \cong SL_k(\mathbb{C})$ naturally acts on $\mathbb{C}[V]$. Similarly, $SL(V)$ also acts on $\mathbb{C}[V^n]$, the ring of polynomials in the coordinates of n vectors in V treated as indeterminates. By the $SL(V)$ *invariants of* n *vectors* we mean elements of $\mathbb{C}[V^n]^{SL(V)}$, the polynomials invariant under the action of $SL(V)$.

Let V^* be dual to V. The group $SL(V)$ also acts on V^*; given $M \in SL(V)$ and $x \in V^*$, we define Mx to be the linear form $(Mx)(v) = x(M^{-1}v)$. This gives an action of $SL(V)$ on $\mathbb{C}[(V^*)^m]$. We call elements of $\mathbb{C}[(V^*)^m]^{SL(V)}$ $SL(V)$ *invariants of* m *covectors*.

Finally, by combining these actions, $SL(V)$ acts on $\mathbb{C}[(V^*)^m \times V^n]$. We will be focused throughout this thesis on $\mathbb{C}[(V^*)^m \times V^n]^{SL(V)}$, the ring of $SL(V)$ *invariants of* n *vectors and* m *covectors*, or *mixed invariants* for short. Following [2], we denote this ring by $R^k_{n,m}$; by convention, we will usually use x_j (with components $x_{j,1}, \ldots, x_{j,k}$) to denote one of our covectors and v_i (with components $v_{1,i}, \ldots, v_{k,i}$) for one of our vectors. When working in $R^k_{n,m}$, we will usually write SL_k for $SL(V)$ to emphasize the role played by k.

Hermann Weyl [14] stated and proved the First and Second Fundamental Theorems for $R^k_{n,m}$. The First Fundamental Theorem states that $R^k_{n,m}$ is generated by three kinds of invariants:

1. Plücker coordinates $\Delta_S = \det(v_{j,i})_{1 \le j \le k, i \in S}$, for S a k-element subset of $[n]$;
2. "dual Plücker coordinates" $D^S = \det(x_{j,i})_{j \in S, 1 \le i \le k}$, for S a k-element subset of $[m]$;

3. pairings $x_j(v_i) = \sum_{\ell=1}^{k} x_{j,\ell} v_{\ell,i}$, for $j \in [m]$, $i \in [n]$.

We will call these invariants the *Weyl generators* of $R_{n,m}^k$. The Second Fundamental Theorem describes the relations among these generators (see, e.g., [10]), but we will not need it in what follows.

It is clear that $R_{n,m}^k$ is a domain. In fact, it is a unique factorization domain; see, e.g., [10, Theorem 3.17].

2.3 Cluster Algebras

In this brief section, we do not intend to give a full introduction to cluster algebras; for this, see, e.g., [5]. Here we deal only with *skew-symmetric* cluster algebras over $\mathbb{C}$ of geometric type.

Let R be a commutative domain containing $\mathbb{C}$. A *seed* $(\mathbf{x}, Q)$ is a collection of *cluster variables* $\mathbf{x} \subset \mathbf{R}$ along with a quiver (i.e., a directed graph) whose vertices are labeled by $\mathbf{x}$. We designate some of the vertices of Q as *mutable*, and the rest as *frozen*. For any mutable vertex v_i of Q labeled by $x_i \in \mathbf{x}$, we can define a new seed $(\mathbf{x}', Q')$ as follows:

1. Perform an *algebraic exchange*, replacing x_i by the new cluster variable x_i' defined by

$$x_i' x_i = \prod_{v_j \to v_i} x_j + \prod_{v_i \to v_j} x_j$$

   where the first product is over all j such that Q has a directed edge from v_j to v_i, and likewise the second product is over all j such that Q has a directed edge from v_i to v_j. Note that x_i' may live in R's field of fractions, rather than R itself, since we must divide by x_i. The other cluster variables in $\mathbf{x}'$ remain the same as in $\mathbf{x}$.

2. *Mutate* Q at v_i, forming a new quiver Q' as follows:

(a) For any directed path $v_j \to v_i \to v_k$ in Q consisting of exactly two edges, add an edge $v_j \to v_k$.

(b) Reverse the direction of every edge incident to v_i.

(c) Remove directed 2-cycles, one by one.

We call two seeds *adjacent* if they are related by such a mutation, and *mutation-equivalent* if they are related by a sequence of mutations. The *cluster algebra* $\mathcal{A}(\mathbf{x}, Q)$ determined by a seed $(\mathbf{x}, Q)$ is the (unital) $\mathbb{C}$-algebra generated by all cluster variables in all seeds mutation-equivalent to $(\mathbf{x}, Q)$. The following proposition gives a useful criterion for the containment of a cluster algebra in the original domain R.

Proposition 2.2 ([2], Corollary 3.7)**.** *Let R be a unique factorization domain finitely generated over $\mathbb{C}$. Let $(\mathbf{x}, Q)$ be a seed such that all cluster variables in it and in all adjacent seeds are irreducible elements of R. Then $\mathcal{A}(\mathbf{x}, Q) \subset R$.*

Cluster algebras have many remarkable properties, and many "classical" algebras turn out to be cluster algebras. In particular, using combinatorial objects equivalent to plabic graphs, Scott [12] showed that the homogeneous coordinate ring of any Grassmannian (with respect to its Plücker embedding) has a cluster structure. Since this coordinate ring can be identified with a ring of vector invariants, it is natural to ask whether rings of mixed SL_k invariants are also cluster algebras. For $k = 3$, Fomin and Pylyavskyy [2] describe multiple cluster structures on each ring $R^3_{n,m}$. Their strategy does not use plabic graphs, and they leave open the question of whether $R^k_{n,m}$ has a cluster structure for $k \geq 4$.

In this thesis, we introduce *mixed plabic graphs* and use them to construct a cluster structure on $R^k_{n,m}$ for $n, m \geq k$.

CHAPTER 3

Mixed Plabic Graphs

In this chapter, we lay out the combinatorial groundwork for everything that follows. Just as Scott's cluster structure on the Grassmannian depends on the combinatorics established by Postnikov [11], so too will the cluster structures on mixed invariant rings depend on the combinatorics of mixed plabic graphs developed here.

3.1 Preliminary Definitions

Definition 3.1. A *biplabic graph* (short for biplanar bicolored graph) is a pair (G, s) where G is a (not necessarily planar) bicolored graph drawn in the plane with boundary labels $B = \{b_1, \ldots, b_{n+m}\}$, and s is a path in G called the *seam*, such that the following properties hold:

1. s is a path from one boundary edge b_s to another b_e;
2. Removing the edges in s divides G into three components: the *mixed part*, the *positive part*, and the *negative part*. The positive part is connected only to white vertices along the seam; the negative part is connected only to black vertices along the seam.
3. The union of the positive part, the seam, and the mixed part forms a plabic graph with clockwise boundary labels $b_1, b_2, \ldots, b_n, b_e, b_{e+1}, \ldots, b_{n+m}$, as does

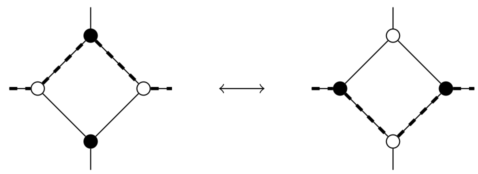

Figure 3.1: The square move (M1) along the seam of a biplabic graph.

the the union of the negative part, the seam, and the mixed part with clockwise boundary labels $b_1, b_2, \ldots, b_s, b_{n+1}, b_{n+2}, \ldots, b_{n+m}$.

By abuse of notation, we often will denote a biplabic graph by G, with the seam understood. In figures, the seam will be indicated by a thick dashed line overlaying its edges.

There is one other plabic graph associated to a biplabic graph which we will occasionally consider. The *mixed graph* of G consists of the mixed part, the seam, and all edges connecting one of the other two parts to the seam (which are now boundary edges in the mixed graph).

The permutation π_G corresponding to a biplabic graph G is determined by following trips as usual.

Moves (M1)-(M3) on biplabic graphs are identical to the corresponding moves on usual plabic graphs, with slight changes. Most square moves behave exactly as on usual plabic graphs, with the exception of squares incident to the seam. Such squares share either one or two edges with the seam, as we will see momentarily. If a square shares one edge with the seam, the square move proceeds as usual after making the vertices of the square trivalent. If a square has two edges along the seam, then performing the square move must also flip the seam to use the other two edges of the square (see Figure 3.1). This is necessary to make sure that the signed parts retain only connections to the correct color vertices of the seam. Additionally, (M2)

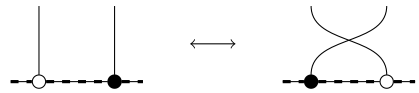

Figure 3.2: The seam swap move (M4).

and (M3) have to be adjusted in obvious ways if they take place along the seam; for example, if we expand a new edge at a vertex on the seam using (M2), then we treat the new edge as part of the seam if and only if we must do so in order to keep the seam connected.

In addition to the usual moves on plabic graph, we have one additional move (M4) on biplabic graphs. Because we think of the positive and negative parts as independent of each other, and because their edges are allowed to cross, we can freely swap edges connecting the positive part to the seam with edges connecting the negative part to the seam (see Figure 3.2).

We will require all our biplabic graphs to be *reduced*, a notion analogous to that of reduced plabic graphs. For convenience in the following definition, call an edge *bicolored* if its endpoints are different colors. Then we call a biplabic graph G reduced if the four conditions of Postnikov's Theorem 13.2 in [11] hold:

1. G has no round trips; i.e., following the rules of the road starting at any vertex (not necessarily a boundary vertex) eventually leads to the boundary.

2. G has no trips with essential self-intersections; i.e., if a trip uses a bicolored edge e in one direction, it does not come back later to use e in the other direction.

3. G has no bad crossings; i.e., there do not exist a pair of trips and a pair of bicolored edges e_1 and e_2 such that both trips use e_1 and then later e_2. (If two trips use the same two edges, one must use e_1 before e_2 and the other e_2 before e_1.)

4. G If $\pi_G(i) = i$, then G has a single leaf attached to the edge labeled i.

It is easy to see that if G is reduced, then so is any G' move-equivalent to G.

3.2 Mixed Plabic Graphs

Definition 3.2. A *mixed plabic graph* is a biplabic graph G together with a partition of the boundary labels $B = P \cup N$ into two distinct parts (the "positive" and "negative" labels) such that two conditions hold:

1. The boundary edges of G lying in the positive (resp. negative) part all have positive (resp. negative) labels.
2. For any biplabic graph in the move-equivalence class of G, trips corresponding to positive labels (resp. negative labels) only travel towards b_e (resp. b_s) when they use edges on the seam.

We will often use positive integers to represent the elements of P and negative integers for the elements of N.

As an immediate consequence of the definition, note that $b_s \in P$ and $b_e \in N$, as the trip starting at b_s must travel along the first edge of the seam, and hence towards b_e (and similarly for b_e towards b_s).

The definition also implies that positive trips never enter the negative part of the graph at all. Since a positive trip always travels towards b_e along the seam, it can only leave the mixed part by turning left at a white vertex, which by definition of a biplabic graph must take it into the positive part. Similarly, negative trips never enter the positive part at all.

We define the *face labels* of a mixed plabic graph G similarly to the face labels in an ordinary plabic graph, with one important change. The trips corresponding to positive labels contribute their index to all faces in the mixed part and positive part lying to their left (and never to the negative part), whereas the trips corresponding

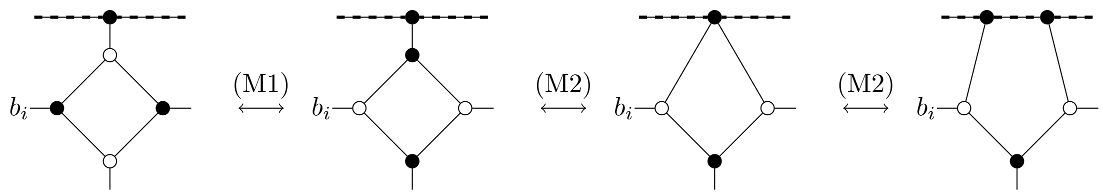

Figure 3.3: The index b_i's trip in the leftmost graph does not touch the seam, but its trip in the move-equivalent rightmost graph does.

to negative labels contribute their index to all faces in the mixed part and negative part lying to their right.

Warning 3.3. Unfortunately, it does *not* suffice to impose Condition 2 above only on G itself. There may be a trip b_i which does not touch the seam in G but does in some move-equivalent G'; from the perspective of G, b_i could be in either P or N, and we do not know which without considering G' (see Figure 3.3).

3.3 Threads in Mixed Plabic Graphs

Given this unfortunate warning, we might wonder how we can ever check the definition of a mixed plabic graph. In this section, we introduce *threads* and prove that certain (more easily checkable) conditions on threads guarantee that we have a mixed plabic graph.

Definition 3.4. Let G be a biplabic graph, and let the boundary vertices be partitioned $B = P \cup N$. The *threading* of G is the subgraph formed by all edges of G which are used in one direction by a positive trip and the other by a negative trip. We direct edges in the threading according to which direction is used by the positive trip. We call edges not in the threading *empty* if used only by negative trips and *crossed out* if used only by positive trips.

To indicate the threading of a graph, we will put an arrow to indicate which directions of edges are used by positive trips. An edge, then, is in the threading if

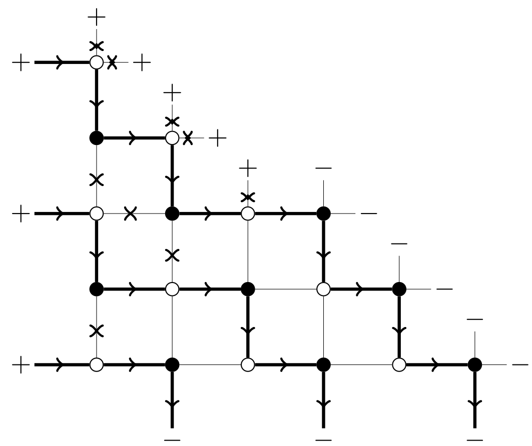

Figure 3.4: A graph with its threading. Positive boundary labels are indicated by $+$, and likewise negative labels with $-$. Arrows show the (directed) edges used by positive trips. The edges in the threading are bolded; empty edges have no decoration, and crossed out edges are decorated with a $\times$ formed by two arrowheads.

it has one (not zero or two) arrows on it, and that arrow gives the direction. See Figure 3.4 for an illustration of these conventions.

Proposition 3.5. *Let G be a trivalent biplabic graph with boundary edges partitioned into positive and negative parts. Then each vertex in the threading has indegree* 1 *and outdegree* 1*.*

Proof. We simply check what the threading looks like locally, at each vertex and its three incident edges. There are (up to obvious symmetry) two possibilities:

1. All incoming edges are of the same sign, from which it follows that all outgoing edges are also the same sign, and none of the three edges contribute to the threading.

2. Two incoming edges are of the same sign (say, positive), and the remaining is different (say, negative). Then the situation looks like Figure 3.5, and we see that this vertex does indeed have indegree 1 and outdegree 1 in the threading.

$\square$

Corollary 3.6. *Let G be trivalent. Then the threading consists solely of* threads

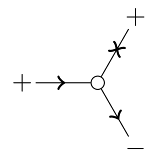

Figure 3.5: A vertex of degree two in the threading.

(directed paths from one boundary vertex to another) and (directed) cycles.

Proof. This is immediate from the proposition, as regular graphs of degree two are unions of lines and cycles. □

Since threads by definition can only occur in the mixed part which is required to be planar, every thread T divides the plane into two regions. We say that T is oriented *forwards with respect to a region* R if empty edges in R are incident to T only at white vertices of T and crossed out edges in R are incident only at black vertices; we similarly say that T is oriented *backwards with respect to* R if the reverse is true. Note that for a given T and R, the thread T is either forwards or backwards with respect to the region R.

Given any two threads, note that there is a well-defined notion of the region R "between" them. We call these two threads *parallel* if one of the threads is oriented forwards with respect to R and the other is oriented backwards with respect to R. In Figure 3.4, all threads are oriented forwards with respect to the region to their southwest and backwards with respect to the region to their northeast, and hence all pairs of threads are parallel.

The complement of the threading consists of several connected components of empty edges and several connected components of crossed out edges. We say that a thread T is *simple* if each one of these connected components which is incident to T is also incident to some other thread as well. All of the connected components of the complement of the threading in Figure 3.4 consist of only a single edge between two different threads, and hence all threads in Figure 3.4 are simple.

We now have enough properties of threadings to say something about how they behave under moves.

Proposition 3.7. *Let G be a biplabic graph with boundary edges partitioned into positive and negative parts. Suppose that the threading of G consists only of threads which are simple and parallel. Then the threadings of all G' move-equivalent to G consist only of simple and parallel threads.*

Remark 3.8. Note that since we are not requiring G to be trivalent, we could have vertices of higher degree in the threading. Hence saying that the threading consists only of threads is stronger than just forbidding cycles—we also forbid, e.g., degree 4 vertices in the threading.

Proof. Note that it suffices to prove the proposition just for G' one move away from G; the result for any graph G'' in the move-equivalence class of G follows by repeatedly applying this to every adjacent pair of graphs in some chain of moves connecting G to G''.

Given a thread T and a region R, the moves cannot change the direction T is oriented with respect to R. This is very clear for all moves except the square move (M1). For (M1), we check the following four cases (which up to symmetry are exhaustive):

1. All trips involved in the square move have the same sign. Then no part of the square (either before or after the move) appears in the threading, so nothing changes.
2. Three trips involved in the square move have the same sign. Then the square has one thread passing through it both before and after, and its orientations

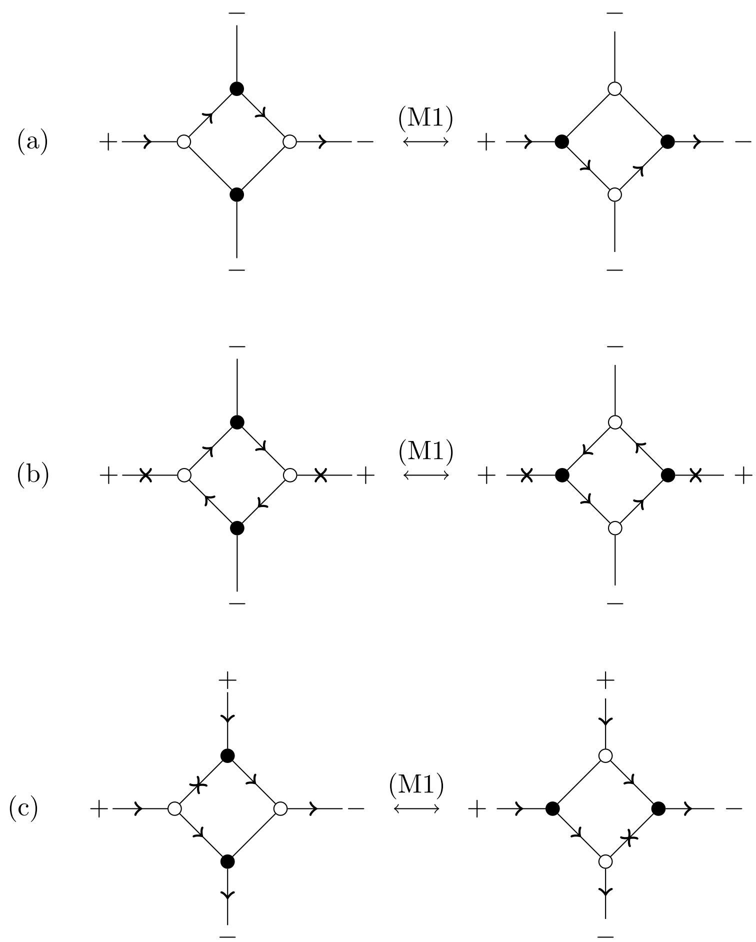

Figure 3.6: The square move's effect on the threading.

with respect to the regions on either side do not change (see Figure 3.6a).

3. Two trips at opposite vertices of the square have the same sign. Then the square forms a cycle, which is impossible in G (see Figure 3.6b). Thus this case does not occur at all.

4. Two trips at adjacent vertices of the square have the same sign. Then applying (M1) has absolutely no effect on the threading (see Figure 3.6c).

Hence all threads in G' are parallel.

Notice that the only way to create a new nonsimple thread is by applying (M2) to a vertex of degree greater than 2 in the threading, and the only way to create a new

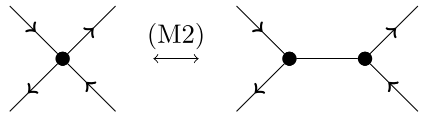

Figure 3.7: “Bad” features of the threading can only arise from preexisting bad features.

vertex of degree greater than 2 in the threading is by contracting a monochromatic edge between threads (see Figure 3.7). But neither side of Figure 3.7 can appear in G: vertices of degree greater than 2 do not show up in the threading of G, and the right hand side of Figure 3.7 depicts either two nonparallel threads or one nonsimple thread, neither of which show up in G. Hence G' cannot have nonsimple threads or vertices of degree greater than 2 in the threading, as desired.

□

We now are prepared to show that such a biplabic graph is in fact a mixed plabic graph, as long as it itself satisfies the conditions of Definition 3.2.

Theorem 3.9. *Let G be a biplabic graph with boundary edges partitioned into positive and negative parts such that the boundary edges of G lying in the positive (resp. negative) part all have positive (resp. negative) labels, and trips corresponding to positive labels (resp. negative labels) only travel towards b_e (resp. b_s) when they use edges on the seam. Finally, suppose that the threading of G consists only of threads which are simple and parallel. Then G is a mixed plabic graph.*

Proof. According to Definition 3.2, we additionally need the condition on the seam for every G' move-equivalent to G. Notice that this condition is equivalent to requiring the seam to be a thread oriented from b_s to b_e.

Since we know this condition holds for G, we begin with the seam coinciding with a thread. Further, notice that the rules for how threads change under the moves coincide precisely for how the seam changes under the moves (see Figure 3.6 and

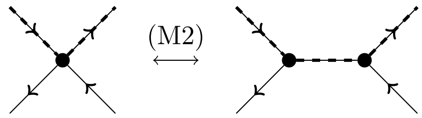

Figure 3.8: The seam may fail to coincide with a thread after applying (M2) to a high degree vertex in the threading.

Figure 3.1 for the square move on threads and seams respectively), with the exception of how they behave on vertices of degree greater than 2 of the threading (see Figure 3.8). However, by the previous proposition, no G' has vertices of degree greater than 2 in the threading. Therefore, we conclude that the seam always coincides with a thread for any G' move-equivalent to G, as desired.

$\square$

3.4 Extensions and Restrictions of Mixed Plabic Graphs

Given a mixed plabic graph P, we will sometimes want to look at smaller portions of it or at certain larger graphs built from it. Given a region R of the plane bounded by a simple closed curve which does not pass through any of P's vertices, the *restriction of P to R* is simply the intersection of P with R. Any edge of P that intersects the boundary of R becomes a new boundary edge in the restriction; this boundary edge inherits its sign from the trip in P which uses it to travel from outside R to inside R.

An *extension of P* is any new biplabic graph formed by adding edges to the positive and/or negative parts of P (but not to the mixed part). Any new boundary edges added must have the appropriate sign (e.g. positive if added to the positive part).

Proposition 3.10. *Suppose P is a mixed plabic graph satisfying the hypotheses of Theorem 3.9. Then extensions and restrictions of P do as well, and hence they are* *mixed plabic graphs.*

Proof. This is an immediate consequence of the definition and Theorem 3.9. Restrictions inherit their threads from the original graph, and hence its threads will be simple and parallel since P's are. Extensions do not change the threads at all, since threads never venture into the positive or negative parts. Therefore, since P satisfies the hypotheses of Theorem 3.9, so too do any extensions and restrictions. □

3.5 Connectivity

Just like plabic graphs, we want reduced mixed plabic graphs with the same trip permutation to be connected by moves. This is still true, as long as we impose a slight additional restriction in order to accommodate the seam.

Theorem 3.11. *Two reduced mixed plabic graphs are move-equivalent if and only if they have the same trip permutation and the same boundary edge b_s as the first edge of the seam.*

Remark 3.12. If we have two mixed plabic graphs with the same trip permutation and the same b_s, then they also share the same last edge of the seam b_e: b_e is forced to be the first boundary edge in the threading clockwise from b_s.

Proof. Since the moves do not change the trip permutation, one direction is trivial.

We will make use of a lemma of Postnikov's:

Lemma 3.13 ([11] Lemma 13.5 and its proof)**.** *Let G be a plabic graph whose trip permutation is π_G. Let $i < i'$ be two indices such that $\pi_G(i) = i'$ or $\pi_G(i') = i$ and there is no pair $j, j' \in [i+1, i'-1]$ such that $\pi_G(j) = j'$. Then one can transform G by moves (M1)-(M3) into a graph such that the trip between i and i' travels along the boundary of G (see Figure 3.9).*

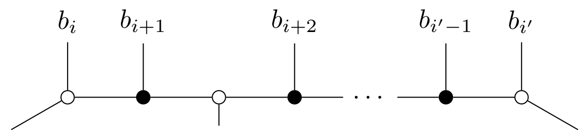

Figure 3.9: Boundary trip for $\pi_G(i) = j$.

Define a *seam trip* to be a trip in the mixed graph which starts and ends at a boundary edge incident to the seam. Note that we do not consider b_s and b_e, which actually lie along the seam, to be incident to the seam. Our first application of Postnikov's lemma is the following proposition, which we will need shortly:

Proposition 3.14. *Let G be a reduced mixed plabic graph. Then G is move-equivalent to a reduced mixed plabic graph without seam trips.*

Proof. Consider the set of mixed plabic graphs move-equivalent to G with the fewest possible number of internal faces in the mixed graph. Among these, pick G' to be one whose mixed graph has the fewest number of boundary edges incident to the seam. We claim that G' has no seam trips.

Suppose not. Let $i \to i'$ be a seam trip in the mixed graph of G' which satisfies the hypotheses of Lemma 3.13; such a seam trip must exist if any seam trips exist, as if a given seam trip $i \to i'$ fails to satisfy these hypotheses, some trip starting and ending between i and i' (and hence necessarily a seam trip) must satisfy them instead. By the lemma, we can transform the mixed graph to have a boundary trip between i and i'.

A priori, this boundary trip may no longer coincide with part of the seam, as the seam may have changed when performing these moves. However, this cannot occur in our situation, as any move inside the mixed graph which changes the seam necessarily reduces the number of internal faces by one.

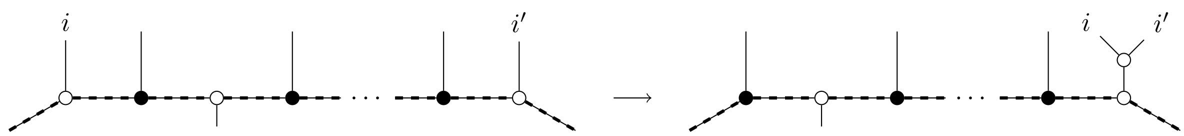

Figure 3.10: Moves (M2) and (M4) turn a boundary trip into one with one fewer boundary edge in the mixed graph.

Hence all internal edges in the boundary trip between i and i' are on the seam. Using (M4), we can move the boundary edge i past all the intervening negative boundary edges, and using (M2), we can move i past all the intervening white vertices. We continue until the boundary edges i and i' share the same vertex along the seam, at which point one further application of (M2) leaves us with a mixed graph with one fewer boundary edge (see Figure 3.10). This contradicts the minimality of G' and completes the proof. □

We are now ready to prove Theorem 3.11. We will apply induction on the number of faces in our mixed plabic graphs. Suppose G and G' are mixed plabic graphs with the same trip permutation. Remove any fixed points from G and G', as they obviously make no difference. By the Proposition 3.14, we may assume G and G' have no seam trips (as otherwise, we can replace them with move-equivalent graphs with no seam trips).

Suppose we can find a trip between a and a' such that both a and a' are boundary edges in the positive part or mixed part. Pick i, i' between a, a' satisfying the conditions of Lemma 3.13. Since G and G' have no seam trips, the trips between i and i' must stay entirely within either the positive part or the mixed graph, depending on where p and p' lived. Applying Lemma 3.13 to the relevant part and removing the resulting boundary trip from G and G' results in graphs G_0 and G_0' with the same trip permutation and fewer faces. By induction, G_0 and G_0' are move-equivalent, and

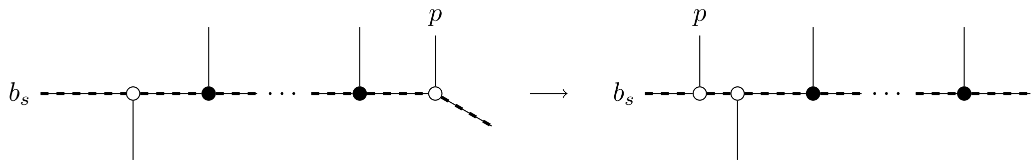

Figure 3.11: Moving p next to b_s.

hence by gluing the boundary trip back on, so too are G and G'.

Suppose instead no such trip exists. Consider the trips in G and G' starting at b_s. They must end up in the positive part: they cannot end in the negative part (since b_s is a positive label), and they cannot end in the mixed part (since we assumed no such trips existed). Hence the trip starting at b_s must leave the mixed graph at some point. Denote by p (resp. p') the boundary edge along which it leaves the mixed graph of G (resp. G'). Since G and G' have no seam trips, these trips in the mixed graph satisfy the hypotheses of Lemma 3.13, so we can turn this into a boundary trip. This boundary trip necessarily travels along the seam, and so similar to Figure 3.10, we can use (M2) and (M4) to put p and p' right next to b_s (see Figure 3.11). From now on, given that p and p' are now both immediately clockwise from b_s, we refer to them both as p.

Now consider the trips in the positive parts between b_s and $b_{\pi(s)}$. Because we arranged for p to be so close to b_s, the only way that $b_s \to b_{\pi(s)}$ could fail to satisfy the hypotheses of Lemma 3.13 is if we had a trip which both starts and ends in the positive part. Since we are assuming no such trips exist, we can apply Lemma 3.13 to $b_s \to b_{\pi(s)}$. As a result, we now have boundary trips $b_s \to b_{\pi(s)}$ in both G and G', so by the same argument, we can remove them, use the inductive hypothesis, and glue them back on to conclude G and G' are move-connected.

$\square$

3.6 Mixed Grassmann Permutations

For the rest of this chapter, we will consider only mixed plabic graphs whose partition of boundary edges into positive and negative labels is $B = P \cup N$, with $P = \{b_1, b_2, \ldots, b_n\}$ and $N = \{b_{n+1}, b_{n+2}, \ldots, b_{n+m}\}$. For convenience, we will identify b_i with i for $1 \leq i \leq n$ and $b_{n+m+1-j}$ with $-j$ for $1 \leq j \leq m$. Hence we think of P as $[n]$ and N as $[-m]$.

Our study will focus on mixed plabic graphs whose trip permutation is the *mixed Grassmann permutation* $\pi^k_{n,m}$, the permutation on $[n] \cup [-m]$ defined by the following rules:

- If $1 \leq i \leq n-k$, then $\pi^k_{n,m}(i) = i + k$.
- If $n - k < i \leq n$, then $\pi^k_{n,m}(i) = i - n - 1$.
- If $-1 \geq i \geq -(m-k)$, then $\pi^k_{n,m}(i) = i - k$.
- If $-(m-k) > i \geq -m$, then $\pi^k_{n,m}(i) = i + m + 1$.

We will require that $k \leq \min(m, n)$. This is our analog of the Grassmann permutation π^k_n defined by $i \mapsto n + i$ (modulo n) in S_n.

A Mixed Plabic Graph for $\pi^k_{n,m}$

Our first task is to prove that any mixed plabic graph with this trip permutation exists at all. It turns out that the graph we construct now will in fact be planar, and, more importantly for the cluster structure discussed in Chapter 4, all its interior faces will be quadrilateral.

We will define an L-shaped plabic graph $L^k_{n,m}$ as follows. Consider the subgraph of the lattice $\mathbb{Z}^2$ consisting of all (x, y) such that either $0 \leq x < k$ and $0 \leq y < n$ or $0 \leq x < m$ and $0 \leq y < k$. Color the vertex (x, y) white if $x + y$ is even and black otherwise. Add boundary edges to all white vertices on vertical boundary segments

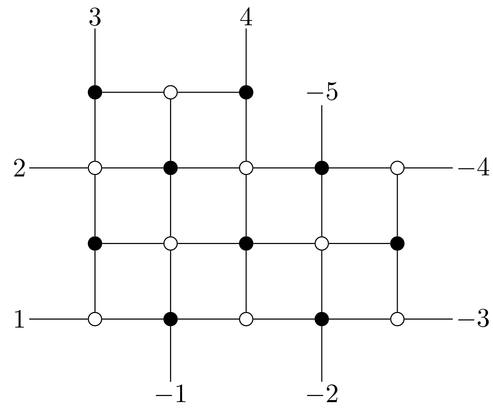

Figure 3.12: The L-shaped plabic graph $L^3_{4,5}$.

and to all black vertices on horizontal boundary segments. Label the boundary edge incident to the origin by 1, and continue clockwise, labeling boundary edges $2, \ldots, n$ in order. Then continue labeling boundary edges clockwise by $-m, -(m-1), \ldots, -1$ (see Figure 3.12).

We claim that, with an appropriate choice of the seam, $L^k_{n,m}$ is a mixed plabic graph for $\pi^k_{n,m}$. To prove this claim, we must check two things: that $L^k_{n,m}$ has the correct trip permutation, and that $L^k_{n,m}$ satisfies the definition of a mixed plabic graph.

The Trip Permutation of the Graph $L^k_{n,m}$

To check the trip permutation of $L^k_{n,m}$, we will use the following lemma about trip permutations in rectangular plabic graphs:

Lemma 3.15. *Let P be a rectangular grid with vertices at (x, y) for all $0 \leq x < k$ and $0 \leq y < n-k$. Let (x, y) be colored white if $x + y$ is even and black otherwise, and attach boundary edges to all white vertices on the vertical sides of the rectangle and to all black vertices on the horizontal sides. Label the boundary edge attached to $(0, 0)$ with 1, and label the remaining boundary edges $2, \ldots, n$ in order clockwise.*

Then P has trip permutation equal to the (usual, not mixed) Grassmann permutation $\pi_n^k \in S_n$.

Remark 3.16. This is a translation of the quadrilateral arrangement as described by Scott [12] into the language of plabic graphs. We verify here that it has the proper trip permutation in order to have a self-contained proof in terms of plabic graphs.

Proof. For convenience, we will zero-index our boundary labels so that the boundary edge incident to the origin is 0, increasing clockwise up to a maximum label of $n-1$. Except for the relabeling, this obviously does not affect the permutation, but it has the convenience of making the boundary edge incident to $(0, 2i)$ have label i.

Consider the trip corresponding to a boundary edge i. Assume that i is on the left edge of the rectangle; by symmetry, the other sides will behave analogously. The trip corresponding to i will consist of a zigzag path $(0, 2i)$, $(0, 2i+1)$, $(1, 2i+1)$, etc. until reaching either the right or top side of the rectangle. Then the trip will "bounce" and continue on an orthogonal zigzag path up and left (if it bounced off the right edge) or down and right (if it bounced off the top). The next time the trip hits a side of the rectangle, it will leave along a boundary edge. We can easily keep track of this bounce by drawing a second copy of our rectangle reflected along the edge we bounce off, with the vertex colors switched. Instead of bouncing, then, we simply continue the trip in the flipped rectangle (see Figure 3.13). Notice that the white vertices along the path will always be of the form $(j, 2i+j)$ and the black vertices of the form $(j, 2i+j+1)$. The trip will finish at one of four locations:

- At the right side of a reflected rectangle to the right, or $(2k-1, 2i+2k)$, which has the same label as $(0, 2i+2k)$, or $i+k$.
- At the top side of a reflected rectangle to the right, or $(n-k-2i, n-k)$, which has

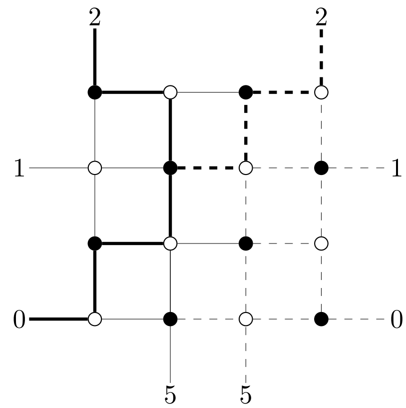

Figure 3.13: Treating a bounce off the right edge as continuing onto a reflected rectangle (drawn with dashed lines).

the same label as $(3k-1-n+2i, n-k)$ in the original plabic graph. If $(x, n-k)$ is a black vertex on the top of P, the boundary edge incident to it has label $(n-k+x+1)/2$. Hence this point has label $(3k-1-n+2i+n-k+1)/2 = i+k$.

- At the right side of a reflected rectangle above, or $(k-1, 2i+k)$. This has the same label as $(k-1, 2n-3k-2i-1)$ in the original plabic graph. If $(k-1, y)$ is a white vertex on the right side of P, the boundary edge incident to it has label $(2n-k-1-y)/2$. Hence this point has label $(2n-k-1-2n+3k+2i+1)/2 = i+k$.
- At the top side of a reflected rectangle above, or $(2n-2k-2i-1, 2n-2k-1)$, which has the same label as $(2n-2k-2i-1, 0)$ in the original plabic graph. If $(x, 0)$ is a black vertex on the bottom of P, the boundary edge incident to it has label $n-(x+1)/2$. Hence this point has label $n-(2n-2k-2i-1+1)/2 = i+k$.

Our final location depends on i, n, and k; luckily, however, since all these possible final boundary edges would be labeled with $i+k$, we need not concern ourselves with which case we are in, as i's trip ends up at $i+k$ in all cases, as desired. $\square$

Now we can handle the L-shaped graph $L^k_{n,m}$.

Proposition 3.17. *The L-shaped graph $L^k_{n,m}$ has trip permutation $\pi^k_{n,m}$.*

Proof. Let R_v be the vertical rectangle $0 \leq x < k$, $0 \leq y < n$ with boundary edges labeled clockwise from the origin as $1, \dots, n+k$, and let R_h be the horizontal rectangle $0 \leq x < m$, $0 \leq y < k$ with boundary edges labeled counterclockwise from $(1,0)$ by $-1, \dots, -(m+k)$. We think of $L^k_{n,m}$ as the union of R_v and R_h. Notice that the last $\lfloor k/2 \rfloor$ of the extra labels $n+1, \dots, n+k$ coincide with the negative labels $-1, \dots, -\lfloor k/2 \rfloor$, and the rest of these extra labels occur in the interior of $L^k_{n,m}$ (and similarly for the extra negative labels $-(m+1), \dots, -(m+k)$).

We will argue for a positive label i mapping to the appropriate place; the negative case is obviously symmetric.

If $i \leq n-k$, then it is easy to see that i's trip never leaves R_v, and hence i maps to $i+k$ as the previous lemma describes trips in R_v.

If $i > n-k$, then we would exit R_v along the edge labeled $i+k$. If $i+k$ is one of the last $\lfloor k/2 \rfloor$ of the extra labels, then it is indeed a boundary edge of $L^k_{n,m}$, coinciding with the negative edge $i+k-(n+k+1) = i-n-1$. Hence $i \mapsto i-n-1$ in this case.

Finally, if $i+k$ is one of the extra labels but not one of these last ones, consider the coordinates of the point where the edge labeled $i+k$ exits R_v. If ℓ is the label of a boundary edge on the right side of R_v, then its coordinates are $(k-1, 2n+k-2\ell+1)$; in our case, then, the coordinates are $(k-1, 2n-2i-k+1)$. Continuing this path in R_h, we have that this path eventually finishes on the (black) vertex $(2n-2i+1, 0)$, which has label $-(2n-2i+2)/2 = i-n-1$ as desired. (Note that if $2n-2i+1 \geq m$, then this point does not actually lie in R_h; however, it is easy to see that the vertex through which we leave R_h on our way towards $(2n-2i+1, 0)$ has the same label as the point $(2n-2i+1, 0)$ would have if m were greater than $2n-2i+1$.) $\square$

The Graph $L^k_{n,m}$ as a Mixed Plabic Graph

Our goal will be to give $L^k_{n,m}$ a seam such that it satisfies the conditions of Theorem 3.9. As in the previous section, it will be convenient to consider a different family of graphs which is easier to study.

Definition 3.18. The *rank k mixed staircase* δ_k is a biplabic graph whose vertices are the points $(x, y) \in \mathbb{Z}^2$ with $x, y \geq 0$ such that $x + y \leq 2k - 2$ for $x \leq k - 1$ and $x+y \leq 2k-1$ for $x \geq k$. (x, y) is white if $x+y$ is even and black otherwise. There is a boundary edge attached to all white vertices on the left and all black vertices on the bottom. Additionally, we attach extra positive boundary edges to white vertices and negative boundary edges to black vertices so that every vertex along the "staircase" on top (i.e., those vertices where the relevant inequality for $x + y$ is an equality) has degree 4.

Figure 3.4 shows the rank 3 mixed staircase and its threading. Note that all the threads are simple and parallel, and hence by Theorem 3.9 it is a mixed plabic graph. This is true in general.

Proposition 3.19. *The rank k mixed staircase is a mixed plabic graph.*

Proof. Since it is clear that we have labeled the boundary edges with the correct signs, the only nontrivial condition in Theorem 3.9 is that the threads are simple and parallel.

We prove a slightly stronger statement by induction. In particular, let the augmented staircase $\tilde{\delta}_k$ be formed by adding extra positive boundary edges to all vertices on the left that do not have one, and extra negative boundary edges to all vertices on the bottom that do not have one. We prove the following three statements simultaneously:

1. δ_k has simple parallel threads.

2. $\tilde{\delta}_k$ has the same threading as δ_k.

3. The trips leaving the left side of $\tilde{\delta}_k$ alternate between negative and positive, starting with negative at $(0, 0)$ and proceeding up; the trips leaving the bottom side of $\tilde{\delta}_k$ also alternate between negative and positive, starting at $(0, 0)$ and proceeding right.

All these statements are almost vacuous for the base case $k = 1$. Proceeding by induction, let us consider the rank $k + 1$ staircase assuming the statements for the rank k case. We form δ_{k+1} by gluing an L-shaped outer shell to $\tilde{\delta}_k$. This outer shell consists of the left and bottom sides of δ_{k+1}, along with extra "boundary" edges which connect it to $\tilde{\delta}_k$ (see Figure 3.14). From the above condition on trips leaving $\tilde{\delta}_k$, we know that these "boundary" edges alternate in sign between negative and positive. It is easy to see, then, that the threading of this outer layer is as in Figure 3.14; the threading of all of δ_{k+1}, then, is simply a slight extension of the threads of δ_k plus a new thread near $(0, 0)$ which is easily seen to be simple and parallel to the existing threads. Since we have created no loops and added a simple parallel thread, we conclude that δ_{k+1} has simple parallel threads.

Finally, since positive trips use all the edges directed downwards along the left side of the outer shell, the extra boundary edges we would add to black vertices on the left to form $\tilde{\delta}_{k+1}$ do not change the threading, and the trips leaving along these new boundary edges are all positive, alternating with the negative trips which leave along the boundary edges on the left of δ_k. Analogously, the extra boundary edges added along the bottom to form $\tilde{\delta}_{k+1}$ behave similarly.

$\square$

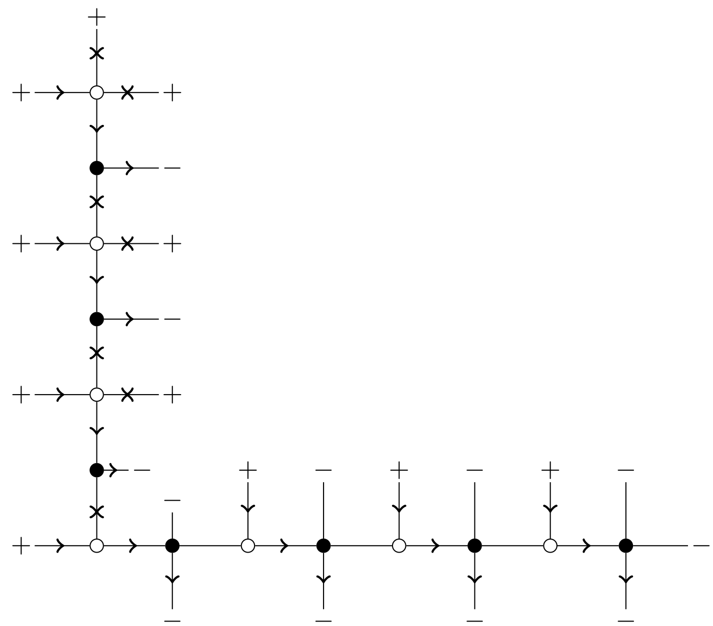

Figure 3.14: The outer shell of δ_{k+1}; fitting $\tilde{\delta}_k$ in would form δ_{k+1}.

Theorem 3.20. *The L-shaped graph $L^k_{n,m}$ is a mixed plabic graph whose trip permutation is $\pi^k_{n,m}$, where the seam of $L^k_{n,m}$ is the thread between k and $-k$.*

Proof. Let V be the set of vertices in both $L^k_{n,m}$ and δ_k (where the vertices of each graph are considered to be points in $\mathbb{Z}^2$ as described in their definitions). Let P be the restriction of δ_k to a simply connected open region around these vertices. We see that $L^k_{n,m}$ is an extension of P, which in turn is a restriction of δ_k. By Proposition 3.10, it follows that $L^k_{n,m}$ is a mixed plabic graph. The seam in δ_k is the thread between k and $-k$, and hence the same is true of $L^k_{n,m}$. Finally, we have already seen in Proposition 3.17 that $L^k_{n,m}$ has trip permutation $\pi^k_{n,m}$ $\square$

Corollary 3.21. *Every mixed plabic graph with trip permutation $\pi^k_{n,m}$ has exactly $k(n+m-k)+2$ face labels.*

Proof. By Theorem 3.20 and Theorem 3.11, it suffices to check this count for $L^k_{n,m}$. But it is easy to see that L has $n+m$ boundary faces and $(k-1)(n+m-k-1)$

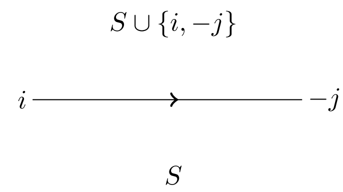

Figure 3.15: The face labels around an edge in a thread.

interior faces, for a total of

$$n + m + nk + mk - k^2 - k + 1 - n - m + k + 1 = k(n + m - k) + 2$$

faces and hence face labels. $\square$

3.7 Face Labels in Mixed Plabic Graphs for $\pi^k_{n,m}$

Let $\mathcal{P}$ be the set of mixed plabic graphs with trip permutation $\pi^k_{n,m}$. In this section, we precisely describe all the subsets of $[n] \cup [-m]$ which occur as face labels of elements of $\mathcal{P}$.

In the mixed part of any mixed plabic graph, consider the two faces as we cross an edge in a thread. Their face labels differ by one positive index and one negative index (see Figure 3.15). Crossing an edge not in the thread replaces one positive index with another (if the edge is crossed out) or one negative index with another (if the edge is empty). For notational convenience, given any face label S, we will let S_+ be the set of positive indices in S and S_- the set of negative indices. We have just observed, then, that crossing over a thread increases $|S_+|$ and $|S_-|$ by one, whereas crossing over any other edge leaves these cardinalities the same.

Now consider a mixed plabic graph P with trip permutation $\pi^k_{n,m}$. The face label of the boundary face between 1 and -1 is empty, as all positive trips occur to its left and all negative trips occur to its right. By the above observations on crossing edges, we see then that every face label in the mixed part of P consists of an equal number of positive and negative indices. Moreover, notice that there are k threads in $L^k_{n,m}$, and hence k threads in every $P \in \mathcal{P}$ (since simple parallel threads stay simple and parallel under moves). To get to a face F in the mixed part which is incident to the seam, we must cross over the other $k-1$ threads. Hence F's face label has $k-1$ positive indices and $k-1$ negative indices. Passing across the seam, we see that the face labels in the positive part consist of k positive indices, and likewise face labels in the negative part consist of k negative indices.

Hence face labels in mixed plabic graphs with trip permutation $\pi^k_{n,m}$ come in three types: positive face labels consisting of k positive indices; negative face labels consisting of k negative indices; and mixed face labels $S = S_+ \cup S_-$ where $|S_+| = |S_-| \leq k-1$. In fact, all such sets occur as face labels in some $P \in \mathcal{P}$:

Theorem 3.22. *$S = S_- \cup S_+$, $S_+ \subset [n], S_- \subset [-m]$, is a face label of some reduced mixed plabic graph with trip permutation $\pi^k_{n,m}$ if and only if one of three cases holds:*

1. *$|S_+| = k$ and $S_- = \emptyset$;*
2. *$S_+ = \emptyset$ and $|S_-| = k$;*
3. *$|S_+| = |S_-| \leq k-1$.*

Proof. We argue by induction on m. Since we require $n, m \geq k$, our base case is $m = k$. In this case, the negative part of any mixed plabic graph P is trivial (consisting of the single unbounded face labeled by $[-k]$), and hence P is in fact plabic. Let us identify the negative label $-j \in [m]$ with $n+m+1-j$ so that we can treat P as a plabic graph with boundary labels $[n+m]$; under this identification, the mixed Grassmann permutation $\pi^k_{n,k}$ is just the (non-mixed) Grassmann permutation $\pi^k_{n,k}$. To avoid ambiguity, we will use P_{plab} and P_{mix} to mean P considered as either a plabic and mixed plabic graph respectively.

Now consider how face labels of P_{plab} compare to face labels of P_{mix} as a mixed plabic graph. $S \subset [n]$ is a face label in (the positive part of) P_{mix} if and only if S is the corresponding face label in P_{plab}. $S = S_+ \cup S_-$ is a face label in P_{mix} if and only if $S_+ \cup S'$ where $S' = \{j \in [n+1, n+m] : j - n - m - 1 \in S_-\}$ is the corresponding face label in P_{plab}. Since all k-element subsets of $[n+m]$ occur as face labels in some P_{plab} with trip permutation π_n^k, all the corresponding face labels occur as face labels of some P_{mix} with trip permutation $\pi_{n,m}^k$.

This covers all face labels of type (1) or (2) above. The only face label of type (3) is $[-k]$, which occurs in every P_{mix} as the only face of the negative part.

Now suppose the theorem holds for $m-1 \geq k$; we will conclude it for m. Given a mixed plabic graph P_0 with trip permutation $\pi_{n,m-1}^k$ and an index $-j \in [-m]$, adding a border trip $-j \mapsto \pi_{n,m}^k(-j)$ to P_0 and bumping every label $-\ell \leq -j$ to $-(\ell+1)$ (to make room for the new $-j$ added) yields a mixed plabic graph P with trip permutation $\pi_{n,m}^k$. P is reduced as long as long as there is no trip $i \to i'$ with $i, i' \in (-j, \pi_{n,m}^k(-j))$, true in our case (as no trips in mixed plabic graphs for $\pi_{n,m}^k$ are nested at all). Further, the face labels of P_0 become face labels of P after changing every $-\ell$ in a face label of P_0 to $-(\ell+1)$.

We will use this observation to construct a mixed plabic graph with an arbitrary face label S as in the statement of the theorem. Pick some $-j \in [-m] \backslash S$; since S contains at most k elements of $[-m]$ and $m > k$, such a $-j$ exists. Let S' be formed from S by replacing all $-\ell \leq -j$ in S by $-(\ell-1)$. Since $-j \notin S$, S' still has the same number of negative elements as S. Further, $-m \notin S'$ since if $-m \in S$, it would have become $-(m-1) \in S'$. Therefore, S' is a subset of $[n] \cup [-(m-1)]$ of one of the types in the theorem, and by the inductive hypotheses, there exists some mixed plabic graph P_0 with trip permutation $\pi_{n,m-1}^k$ which has S' as a face label. By the observation above, adding a boundary trip $-j \mapsto \pi^k_{n,m}(-j)$ to P_0 will yield a mixed plabic graph with the desired face label S.

$\square$

CHAPTER 4

A Cluster Structure on Mixed Invariant Rings

In this chapter, we establish a cluster structure on mixed invariant rings. Analogous to the Grassmannian case, we do so by identifying face labels in mixed plabic graphs with special invariants in our ring, then proving that square moves in mixed plabic graphs correspond to valid exchange relations in the mixed invariant ring. We use a strategy we call *complementation*, whereby we essentially use the isomorphism $\bigwedge^r \mathbb{C} \cong \bigwedge^{k-r} \mathbb{C}^*$ so that relations in mixed invariant rings can be proved by appealing to relations in the Grassmannian case.

4.1 Special Invariants in $R_{n,m}^k$

In our study of mixed invariants, we will primarily consider three different kinds of elements of $R_{n,m}^k$, which we collectively refer to as *special invariants*:

- Plücker coordinates Δ_I, corresponding to k-element subsets $I \subset [n]$;
- "Dual Plücker coordinates" D^J, corresponding to k-element subsets $J \subset [-m]$;
- Mixed invariants W_S, corresponding to unions $S = S_+ \cup S_-$ where $S_+ \subset [n]$, $S_- \subset [-k]$, and $|S_+| = |S_-| \leq k-1$. We call $r = |S_+| = |S_-|$ the *size* of W_S. If X_{S_-} is the $r \times k$ matrix whose rows are x_{-j} for $-j \in S_-$ and V_{S_+} is the $k \times r$

matrix whose columns are v_i for $i \in S_+$, then we define

$$W_S = \det(X_{S_-} V_{S_+}).$$

It is easy to see that in terms of the Weyl generators, $W_S = \det(x_{-j}(v_i))_{-j \in S_-, i \in S_+}$. In particular, size 1 mixed invariants are themselves Weyl generators.

Proposition 4.1. *Every special invariant is an irreducible polynomial.*

Proof. Let W be any special invariant (not necessarily mixed) with a factorization $W = AB$. To conclude W is irreducible, we must show either A or B is a unit.

The invariant W is multilinear in each individual vector or covector, and it is alternating among the vectors and among the covectors. Without loss of generality, let $v_1, \ldots, v_r$ be the (possibly empty) set of vectors appearing in W. Since $W = AB$, by multilinearity, either A or B (but not both) involve each $v_1, \ldots, v_r$. If v_a appears in A and v_b appears in B, then neither A nor B vanishes upon setting $v_a = v_b$ but W does, contradicting that $R^k_{n,m}$ is a domain. Hence if some v_a appears in A, all of $v_1, \ldots, v_r$ appear in A, and none in B. A similar argument shows that all covector contributions are contained entirely in A or in B.

If one contains both the vector and the covector contributions (as is trivially the case if W is a Plücker or dual Plücker coordinate), then the other is forced to be a unit, as desired.

It remains to consider $W = AB$ where W is mixed with $|S_-| = |S_+| = r$, for $1 \le r \le k-1$. We show by contradiction that it is impossible for this factorization to be nontrivial. Without loss of generality, suppose A contains only the vectors $v_1, \ldots, v_r$, and B contains only the covectors $x_1, \ldots, x_r$. Let $e_1, e_2, \ldots, e_k$ be a basis of V and let $f_1, f_2, \ldots, f_k$ be the corresponding dual basis of V^*. Then

$$W(e_1, e_2, \ldots, e_r, f_1, f_2, \ldots, f_r) = W(e_2, e_3, \ldots, e_{r+1}, f_2, f_3, \ldots, f_{r+1}) = 1,$$

as both are determinants of the $r \times r$ identity matrix. Then since

$$1 = W(e_1, e_2, \ldots, e_r, f_1, f_2, \ldots, f_r) = A(e_1, e_2, \ldots, e_r)B(f_1, f_2, \ldots, f_r),$$

we must have $A(e_1, e_2, \ldots, e_r) \neq 0$; similarly, we have $B(f_2, f_3, \ldots, f_{r+1}) \neq 0$. However, $W(e_1, e_2, \ldots, e_r, f_2, f_3 \ldots, f_{r+1})$ must be zero, as $f_{r+1}(e_i) = 0$ for $i = 1, 2, \ldots, r$. But then

$$0 = W(e_1, e_2, \ldots, e_r, f_2, f_3 \ldots, f_{r+1}) = A(e_1, e_2, \ldots, e_r)B(f_2, f_3, \ldots, f_{r+1}) \neq 0,$$

a contradiction. Hence there is no factorization $W = AB$ of mixed invariants, as desired.

$\square$

Remark 4.2. Note that size k mixed invariants in $R^k_{n,m}$ are *not* irreducible: in fact, in this case, $W_S = \Delta_{S_+} D^{S_-}$. The above proof breaks down because since $r = k$, we do not have basis elements e_{r+1} and f_{r+1}.

4.2 Mixed Plabic Graphs and Special Invariants

We now have two collections which are indexed by the same family of sets:

1. Face labels in mixed plabic graphs with trip permutation $\pi^k_{n,m}$ (Theorem 3.22);
2. Special invariants in $R^k_{n,m}$.

Given this correspondence, we can think of face labels in mixed plabic graphs as special invariants. The square move in mixed plabic graphs gives us a notion of a *combinatorial exchange*: given any square in a mixed plabic graph, we can replace the invariant labeling it by the new invariant determined by the result of the square move.

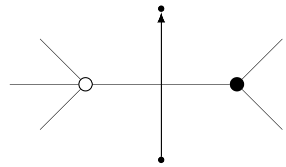

Figure 4.1: Edges in the quiver associated to a mixed plabic graph.

We associate a seed $(\mathbf{W}, Q)$ to a mixed plabic graph P as follows. The cluster variables $\mathbf{W}$ are the special invariants corresponding to face labels of P. Each face of P is a vertex of Q, labeled by the special invariant corresponding to its face label. The mutable vertices are the interior faces of P. Given two faces of P separated by a bicolored edge, we add an edge between the corresponding vertices oriented so that the bicolored edge's white vertex is on its left (see Figure 4.1). Finally, we remove any directed 2-cycles, one by one; such 2-cycles might arise near the seam of P, as well as elsewhere if P has degree 2 vertices.

Since we have a quiver labeled with elements of $R^k_{n,m}$, recall that we have a notion of *algebraic exchange* as well. If W_i is the invariant attached to the ith vertex, we can define W_i' by

$$W_i' W_i = \prod_{j \to i} W_j + \prod_{i \to j} W_j.$$

A priori, W_i' lives in the field of fractions of $R^k_{n,m}$, though we will (eventually) see that it always lives in $R^k_{n,m}$ itself.

Our goal in this section is to justify the correspondence between face labels and special invariants by verifying that every combinatorial exchange coincides with the corresponding algebraic exchange. Our strategy is called *complementation.*

Complementation when $m = k$

Suppose to start that we are in $R^k_{n,k}$, the ring of SL_k invariants of n vectors and k covectors. As usual, our special invariants come in three forms (Plücker coordinates, dual Plücker coordinates, and mixed invariants). Instead of thinking of the Plücker and dual Plücker coordinates as separate, we will instead think of the product of each Plücker coordinate with the unique dual Plücker coordinate. The "index set" of the result, then, is $S = S_+ \cup S_-$, where $S_+ \in \binom{[n]}{k}$ and $S_- = [-k]$. Hence every set $S \in \mathcal{P}([n] \cup [-k])$ of the form $S_+ \cup S_-$ where $|S_+| = |S_-|$ is the index set of some invariant, and every corresponding invariant is some W_S. Let $\mathcal{S} \subset \mathcal{P}([n] \cup [-k])$ be the collection of all such S.

We claim that mutations in $R^k_{n,k}$ are closely related to mutations in R^k_{n+k}, the ring of SL_k invariants in $n+k$ vectors (and no covectors). Instead of $[n+k]$, it will be convenient to use $[n] \cup [-k]$ as the index set for R^k_{n+k}. We form a correspondence between $\mathcal{S}$ and k-element subsets of $[n] \cup [-k]$. Given any $S \in \mathcal{S}$, we form $S' = S_+ \cup S_-^c$, where $S_-^c = [-k] \backslash S_-$. We have $|S'| = |S_+| + |S_-^c| = |S_+| + (k - |S_-|) = k$, so S' is indeed a k-element subset of $[n]$.

It is clear that this is reversible and hence a bijection; call it the *complementation map* on sets. Let $c : \binom{[n]\cup[-k]}{k} \to \mathcal{S}$ be one direction of this bijection.

Remark 4.3. This bijection is a combinatorial interpretation of the identity

$$\sum_{i=0}^{k} \binom{n}{i}\binom{k}{i} = \binom{n+k}{k},$$

which itself is a special case of

$$\sum_{i=0}^{k} \binom{a}{i}\binom{b}{k-i} = \binom{a+b}{k}.$$

By abuse of notation, let us also denote by c the map which sends the Plücker coordinate indexed by S' to the invariant W_S.

Proposition 4.4. *The map c extends to an algebra homomorphism $c : R^k_{n+k} \to R^k_{n,k}$ whose image is the subalgebra generated by the W_S for $S \in \mathcal{S}$.*

Proof. To verify that c extends to a homomorphism, we must check that c sends all elements of the Plücker ideal to 0.

Denote our vectors in $R^k_{n,k}$ by $v_1, \ldots, v_n$ and our covectors as $x_{-1}, \ldots, x_{-k}$. Let the extra k vectors in R^k_{n+k} be denoted $v_{-1}, \ldots, v_{-k}$. Let the components of v_i be indeterminates $v_{i,1}, \ldots, v_{i,k}$, and likewise for the x_{-j}. We will think of the v_i as column vectors and the x_{-j} as row vectors. Given any $S_+ \subset [n]$, let V_{S_+} be the matrix whose columns are the v_i for $i \in S_+$, and similarly, given $S_+ \subset [-k]$, let X_{S_-} be the matrix whose rows are the x_{-j} for $-j \in S_-$.

Consider any W_S for $S \in \mathcal{S}$, and write $S = S_+ \cup S_-$ as usual. Let $r = |S_+| = |S_-|$. We have

$$W_S = \det(x_j(v_i)) = \det(X_{S_-}V_{S_+})$$

which, by Cauchy–Binet, equals

$$\sum_{R \in \binom{[k]}{r}} D^R(X_{S_-})\Delta_R(V_{S_+}).$$

This in turn is the the induced pairing on the rth exterior powers, so we have

$$W_S = \left(\bigwedge_{-j \in S_-} x_{-j}\right)\left(\bigwedge_{i \in S_+} v_i\right)$$

where here (and in what follows) wedges over negative indices go in decreasing order.

Now consider evaluating W_S at any point where $D = \det(X_{[-k]}) \neq 0$. We claim that we can choose values of $v_{-1}, \ldots, v_{-k}$ such that the value of W_S equals the value of $\operatorname{sgn}(S_-)\Delta_{S'}$, where $S' = c^{-1}(S)$ and $\operatorname{sgn}(S_-)$ is the sign of the permutation taking $-1, -2, \ldots, -k$ to the elements of $[-k]\setminus S_-$ in decreasing order followed the elements in S_- in decreasing order. Indeed, take $v_{-1}, \ldots, v_{-k}$ to be $D^{1/k}$ times the dual basis to $x_{-1}, \ldots, x_{-k}$ (which is in fact a basis, since we are assuming their determinant D is not 0); we claim that this choice of the v_{-j} works.

To see this, consider the linear functional f_{S_-} on the $\bigwedge^r \mathbb{C}^k$ defined by

$$f_{S_-}(u_1 \wedge \cdots \wedge u_r)\,\mathrm{Vol} = \mathrm{sgn}(S_-) D^{r/k} \left(\bigwedge_{-j \notin S_-} v_{-j} \right) \wedge u_1 \wedge \cdots \wedge u_r$$

where Vol is the standard volume form. Notice that $f_{S_-}\left(\bigwedge_{-j\in S_-} v_{-j}\right) = D^{r/k}$, as the determinant of all the v_{-j} in (decreasing) order is $(D^{1/k})^k / \det(x_{-1}, \ldots, x_{-k}) = 1$, and $\mathrm{sgn}(S_-)$ is, by the definition above, the sign of the permutation which puts this wedge back in order.

We claim that $f_{S_-} = \bigwedge_{-j\in S_-} x_{-j}$. To see this, consider the values of these two functionals on the basis of $\bigwedge^r \mathbb{C}^k$ induced by the basis $v_{-1}, \ldots, v_{-k}$ of $\mathbb{C}^k$. If we define $v_R = \bigwedge_{-j\in R} v_{-j}$, then by definition, $f_{S_-}(v_R) = 0$ unless $R = S_-$, in which case it equals $D^{r/k}$. On the other hand, it is easy to see that

$$\bigwedge_{-j\in S_-} x_{-j}(v_R) = \det(X_{S_-} V_R)$$

behaves the same way; $X_{S_-}V_R$ will have a row of all zeros if R contains some $-j \notin S_-$, and it will be diagonal with $D^{1/k}$ along the diagonal if $R = S_-$.

Hence we have that

$$\begin{aligned}
W_S\,\mathrm{Vol} &= \left(\bigwedge_{-j\in S_-} x_{-j} \right) \left(\bigwedge_{i\in S_+} v_i \right) \mathrm{Vol} \\
&= f_{S_-} \left(\bigwedge_{i\in S_+} v_i \right) \mathrm{Vol} \\
&= \mathrm{sgn}(S_-) D^{r/k} \bigwedge_{-j\notin S_-} v_{-j} \wedge \bigwedge_{i\in S_+} v_i \\
&= \mathrm{sgn}(S_-)(-1)^{\binom{r}{2}} D^{r/k} \Delta_{S'}\,\mathrm{Vol}
\end{aligned}$$

where the factor of $(-1)^{\binom{r}{2}}$ comes from putting the wedge over S_- in increasing order instead (in order to make Δ'_S). We conclude, then, that $W_S = \mathrm{sgn}(S_-)(-1)^{\binom{r}{2}} D^{r/k} \Delta_{S'}$.

Before we proceed, let us investigate this sign further. If the elements of S_- are, in decreasing order, $a_1 > a_2 > \cdots > a_r$ and the remaining elements of $[-k]$ are $b_1 > b_2 > \cdots > b_{k-r}$, then the permutation whose sign we are investigating is $b_1 \ldots b_{k-r} a_1 \ldots a_r$. It takes $|b_{k-r}| - (k-r)$ adjacent transpositions to put b_{k-r} in the correct place among the a_i; then $|b_{k-r-1}| - (k-r-1)$ for b_{k-r-1}, etc. In total, then, it takes

$$\begin{aligned}\sum |b_i| - (k-r)(k-r+1)/2 &= k(k+1)/2 - \sum |a_i| - (k-r)(k-r+1)/2 \\ &= kr - \sum |a_i| - \binom{r}{2}\end{aligned}$$

adjacent transpositions to put this permutation in order, and hence

$$\operatorname{sgn}(S_-) = (-1)^{kr - \sum |a_i| - \binom{r}{2}}.$$

Combining this with the $(-1)^{\binom{r}{2}}$ in our formula for W_S, we see

$$\begin{aligned}W_S &= (-1)^{kr - \sum |a_i|} D^{r/k} \Delta_{S'} \\ &= (-1)^{\sum_{i=1}^{r} (k+a_i)} D^{r/k} \Delta_{S'}.\end{aligned}$$

Now we are almost done. Our map c clearly extends to an algebra homomorphism on the ring of polynomials in Plücker coordinates. Since R^k_{n+k} is a quotient of this polynomial ring by the Plücker ideal, it remains to check that every element of the Plücker ideal maps to 0. Suppose we have an element p' of the Plücker ideal; i.e., a polynomial in Plücker coordinates which is identically zero. Under c, each of these Plücker coordinates $\Delta_{S'}$ maps to W_S, and hence p' itself maps to some polynomial p in the W_S.

On a dense set (i.e,. where $\det(x_{-1}, \ldots, x_{-k}) \neq 0$), we have just shown that the value of every W_S in fact equals the value of $(-1)^{\sum_{i=1}^{r}(k+a_i)} D^{r/k} \Delta_{S'}$ at some point. Since p is multihomogeneous (since p' is), we see that the combined power of D on each term of p is the same. Similarly, the combined power on (-1) is also the same. Hence on a dense set, the value of p equals a constant multiple of the value of p'. Since p' is identically zero, we have that p equals zero on a dense set, and hence p itself is identically zero. □

Armed with this homomorphism, we can now prove the main result of this section.

Proposition 4.5. *Combinatorial exchanges and algebraic exchanges on squares in mixed plabic graph with trip permutation* $\pi_{n,k}^{k}$ *coincide.*

Proof. Exchanges come in three types:

- Exchanges which take place entirely in the positive part of the graph;
- Exchanges which occur along the seam, which involve some Plücker coordinates, some mixed invariants, and possibly the one dual Plücker coordinate;
- Exchanges which occur entirely in the mixed part, which involve only mixed invariants.

The first type of exchange needs no checking, as these are exactly the same as exchanges which occur in (non-mixed) plabic graphs.

The second type of exchange can be reduced to the third case. Let D be the one dual Plücker coordinate. Since exchange relations are multihomogeneous, the difference between the vector degree and the covector degree of every term must be constant. Since D can appear in at most one term and is the only source of a negative degree difference, we see that the exchange relation must have homogeneous positive degree difference. Since Plücker coordinates are the only source of positive degree difference, this difference must be a multiple of k. Therefore, we can multiply the entire exchange relation by enough copies of D so that the degree difference becomes zero. The new equation will be valid if and only if the original exchange relation is (as all we have done is multiply by a polynomial which is not identically zero), so it suffices to check this new equation. But now, all Plücker coordinates Δ_{S_+} have been paired with a copy of D, and since $\Delta_{S_+}D = W_S$ for $S = S_+ \cup [-k]$, we see that our new equation involves only mixed invariants, and we have succeeded in reducing these exchanges to the third case.

Finally, the third case follows immediately from our complementation homomorphism $c : R^k_{n+k} \to R^k_{n,k}$. After applying set complementation, we have an exchange relation in a plabic graph, known to be true. Pushing it forward by c yields the original (mixed) exchange relation, which holds since c is a homomorphism. $\square$

Complementation in $R^k_{n,m}$

We will now extend the complementation technique beyond the case $m = k$ to any $m \geq k$. First, recall that the image of the complementation homomorphism $c : R^k_{n+k} \to R^k_{n,k}$ is generated by the mixed invariants W_S (including those mixed invariants which are products of one Plücker and one dual Plücker coordinate). Let $W^k_{n,k}$ be the algebra generated by these mixed invariants. If $k' \geq k$, we have the specialization $W^{k'}_{n,m} \to W^k_{n,m}$ obtained by projecting $\mathbb{C}^{k'}$ onto $\mathbb{C}^k$ (in coordinates, we can set the last $k' - k$ coordinate variables of each vector and covector to zero). Roughly, then, our strategy is to prove identities in $W^k_{n,m}$ by treating them as images of identities in $W^m_{n,m}$, which in turn are images under complementation of identities in R^m_{n+m}.

Theorem 4.6. *Combinatorial exchanges and algebraic exchanges on squares in mixed plabic graph with trip permutation $\pi^k_{n,m}$ coincide.*

Proof. Once again, we consider three cases:

- Exchanges which take place entirely in one of the signed parts;

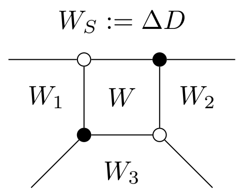

Figure 4.2: Exchanges along the seam involving three mixed invariants.

- Exchanges which occur along the seam, which involve some Plücker and/or dual Plücker coordinates and some mixed invariants;
- Exchanges which occur entirely in the mixed part, which involve only mixed invariants.

The first case continues to be easy. The third is also straightforward by the projection idea described above: every $W_S \in R^k_{n,m}$ is a projection of the "same" (as a polynomial in pairings $x_{-j}(v_i)$) $W_S \in R^m_{n,m}$, where the relation is guaranteed by Proposition 4.5.

The second case now requires more work. Projecting from $R^m_{n,m}$ does not immediately solve this problem, as the Plücker and dual Plücker coordinates in each are not as easily related as the W_S.

Instead, consider what combinations of Plücker coordinates, dual Plücker coordinates, and mixed invariants are possible among the invariants entering into the exchange binomial. Locally, exchanges in a square along the seam must be one of three types:

- One Plücker coordinate, one dual Plücker, and three mixed invariants, in which case the invariant we are exchanging must be mixed (see Figure 4.2). In this case, we can pair the Plücker and dual Plücker coordinates to form a single mixed invariant W_S. Now the exchange relation is simply an exchange relation between several mixed invariants, and it holds by the third case.
- Two Plücker coordinates and two mixed invariants (see Figure 4.3). The invari-

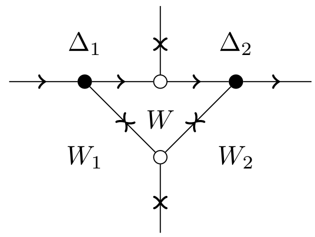

Figure 4.3: Exchanges along the seam involving two Plücker coordinates and mixed invariants. The threading is indicated to make it easier to see that W, W_1, and W_2 have the same negative indices.

ant we are exchanging could either be mixed or it could be a Plücker coordinate itself, depending on which side of the seam we are on. Regardless, notice that all the mixed invariants involved have exactly the same set S_- of covectors. Since these mixed invariants have $\leq k$ covectors (in fact, these should have exactly $k-1$), we can ignore all but k of the covectors and treat this as an exchange in $R^k_{n,k}$, which again holds by Proposition 4.5.

- Two dual Plücker coordinates and two mixed invariants. This is entirely symmetric to the previous case.

Since we have verified all the possibilities for square exchange relations, the theorem is proved. $\square$

4.3 The Mixed Invariant Ring $R^k_{n,m}$ as a Cluster Algebra

We associate to any mixed plabic graph P with trip permutation $\pi^k_{n,m}$ the seed $(\mathbf{x}_P, \tilde{Q}_P)$ whose cluster variables are the special invariants appearing as face labels in P and whose quiver is the reduced quiver of P. Let $\mathcal{A}_P = \mathcal{A}(\mathbf{x}_P, \tilde{Q}_P)$ be the cluster algebra determined by this seed. We now show that all these seeds give rise to the same cluster algebra, and that this cluster algebra is $R^k_{n,m}$.

Theorem 4.7. *Let P be a mixed plabic graph whose trip permutation is $\pi^k_{n,m}$. Then $\mathcal{A}_P = R^k_{n,m}$.*

Proof. Since all mixed plabic graphs with the same trip permutation are connected by square moves (Theorem 3.11) and these combinatorial exchanges correspond with algebra exchanges (Theorem 4.6), any P' with trip permutation $\pi^k_{n,m}$ corresponds to a seed $(\mathbf{x}_{P'}, \tilde{Q}_{P'})$ in $\mathcal{A}_P$. Since all possible face labels occur in some mixed plabic graph (Theorem 3.22), all special invariants appear as cluster variables in $\mathcal{A}_P$. In particular, the Weyl generators of $R^k_{n,m}$ are cluster variables, and hence $R^k_{n,m} \subseteq \mathcal{A}_P$.

Since $L = L^k_{n,m}$ is a mixed plabic graph with trip permutation $\pi^k_{n,m}$ (Theorem 3.20), $(\mathbf{x}_L, \tilde{Q}_K)$ is a seed in $\mathcal{A}_P$. Moreover, since all internal faces of L are quadrilateral, L and all adjacent clusters consist entirely of special invariants, which are irreducible (Proposition 4.1). Since $R^k_{n,m}$ is a unique factorization domain, Proposition 2.2 applies, and we conclude $\mathcal{A}_P = R^k_{n,m}$ as desired. $\square$

Example 4.8. The proof of Theorem 4.7 uses the L-shaped mixed plabic graph $L^k_{n,m}$ to construct an initial seed for the cluster structure on $R^k_{n,m}$. $L^3_{5,5}$ is shown in Figure 4.4; the cluster variables in this seed are the special invariants corresponding to the face labels.

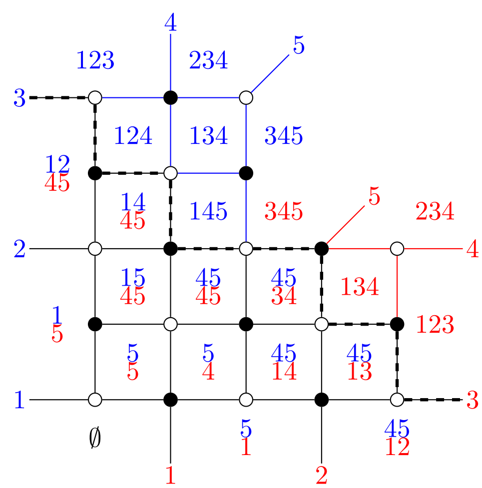

Figure 4.4: The L-shaped plabic graph for $R^3_{5,5}$. Blue numbers indicate positive indices, and red numbers indicate negative indices.

By following the construction in Section 4.2, $L^3_{5,5}$ gives rise to an L-shaped quiver (Figure 4.5). In general, the mutable part of the quiver is an "L" $k-1$ vertices in thickness, $n-1$ vertices tall, and $m-1$ vertices wide.

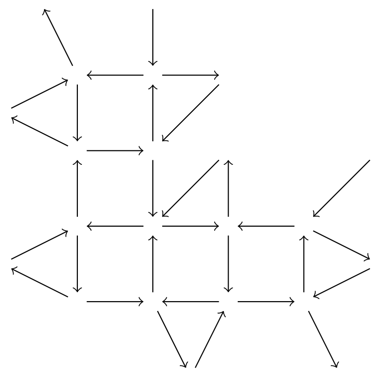

Figure 4.5: The quiver corresponding to the L-shaped plabic graph for $R^3_{5,5}$.

CHAPTER 5

Mixed Weak Separation

Compatibility of Plücker coordinates on cluster structures of Grassmannians is equivalent to the corresponding subsets being weakly separated. In this chapter, we generalize the notion of weak separation to apply to face labels of mixed plabic graphs, which are the special subsets of $[n] \cup [-m]$ described in Theorem 3.22. The goal in this chapter is to prove that each maximal (by inclusion) collection of pairwise weakly separated special subsets corresponds to a cluster in the cluster structure on $R^k_{m,n}$, and hence all such maximal collections have the same cardinality.

5.1 Weak Separation for Subsets

Ordinary Weak Separation

Definition 5.1. Let $S, T \in \binom{[n]}{k}$. Then S and T are *weakly separated*, denoted $S||_u T$, if there do not exist $a, c \in S \backslash T$ and $b, d \in T \backslash S$ such that $a < b < c < d$ or $b < c < d < a$.

Remark 5.2. Occasionally, we will wish to consider S and T to be different sizes. For our purposes, we will consider such S and T to be weakly separated according to the same rule just given.

Maximal collections of pairwise weakly separated k-element subsets of $[n]$ enjoy the following "purity" result, implied by a conjecture of Leclerc and Zelevinsky [8] and conjectured directly by Scott [13]:

Theorem 5.3. *All maximal collections of pairwise weakly separated k-element subsets of $[n]$ consist of precisely $k(n-k)+1$ subsets. Further, every such collection occurs as the collection of face labels of some plabic graph for the Grassmann permutation $\pi_{k,n}$.*

This theorem was proved by Oh, Postnikov, and Speyer [9] using their notion of *plabic tilings*. Their result also implies Leclerc and Zelevinsky's conjecture, which Danilov, Karzanov, and Koshevoy [1] also proved independently using a different sort of tiling.

In this chapter, we will take a similar strategy to Oh, Postnikov, and Speyer to extend this result to the mixed setting.

Mixed Weak Separation

Recall that in the mixed setting, our subsets may consist both of positive and of negative elements. For positive integers k, m, and n, let $\mathcal{C}^k_{n,m}$ denote the collection of all sets S of one of three types:

1. S is a k-element subset of $[n]$, which we will call a *Plücker subset*;
2. S is a k-element subset of $[-m]$, which we will call a *dual Plücker subset*;
3. $S = S_+ \cup S_-$, where $S_+ \subset [n]$, $S_- \subset [-m]$, and $|S_+| = |S_-| < k$, which we will call a *mixed subset*.

We will throughout require that $n, m \geq k$ so that $\mathcal{C}^k_{n,m}$ contains at least one set of each of the three types.

Definition 5.4. Let $S, T \in \mathcal{C}^k_{n,m}$. Then S and T are *weakly separated*, also denoted $S||T$, if the following holds. If both S and T are mixed, then $S||T$ as mixed subsets if and only if $S_+ \cup T_- ||_u T_+ \cup S_-$ as *unmixed* subsets. Otherwise, $S||T$ if and only if $S||_u T$.

Just as maximal (unmixed) collections of weakly separated subsets appear as the face labels of plabic graphs, we will see that maximal collections of weakly separated subsets in $\mathcal{C}^k_{n,m}$ appear as face labels of the mixed plabic graphs from Chapter 3. Our goal in this chapter is to prove a purity result analogous to Theorem 5.3:

Theorem 5.5. *All maximal collections of pairwise weakly separated subsets in $\mathcal{C}^k_{n,m}$ have cardinality $k(n+m-k)+2$. Every such collection appears as the collection of face labels of a mixed plabic graph for the Grassmann permutation $\pi^k_{n,m}$.*

Section 5.3 will be devoted to proving this theorem, but first we will recast the problem in terms of slightly different combinatorial objects called sign strings.

5.2 Sign Strings

Definitions

A length n *sign string* is a sequence of n symbols, each one a `+`, `-`, or `0`. A sign string is *segregated* if it avoids the patterns `+-+-` and `-+-+`; i.e., if neither of these patterns appears as a (not necessarily consecutive) substring. Note that this notation is cyclically invariant: cyclically shifting an entire sign string does not change whether it is segregated. The *sum* of a sign string is simply the number of pluses minus the number of minuses. Given a subinterval I, the sum of a sign string s along this interval, denoted $\sigma_I(s)$, is simply the number of pluses in I minus the number of minuses in I. As a special case of this notation, set $\sigma_i(s) = \sigma_{[i,i]}(s)$ so that σ_i gives the sign of the ith position of s.

We say that sign strings s and t are *consistent* if there is no position where one has a plus and the other has a minus. If additionally t has no zeros, we say that s is *supported* by t. Note that if we have a collection of pairwise consistent sign strings, then there exists a sign string supporting them all.

Note also that the all zeros string $\texttt{0}^n$ is consistent with all sign strings of the same length. Occasionally, we will refer to the all zeros string as $\mathbf{0}$ to avoid reference to its length.

If s and t are consistent, we can form their *difference* $s \ominus t$ component-wise in the obvious way: $\texttt{+} \ominus \texttt{+} = \texttt{0}$, $\texttt{+} \ominus \texttt{0} = \texttt{+}$, $\texttt{0} \ominus \texttt{+} = \texttt{-}$, $\texttt{0} \ominus \texttt{0} = \texttt{0}$, $\texttt{0} \ominus \texttt{-} = \texttt{+}$, $\texttt{-} \ominus \texttt{0} = \texttt{0}$, and $\texttt{-} \ominus \texttt{-} = \texttt{0}$. Note that since s and t are assumed consistent, we never encounter $\texttt{+} \ominus \texttt{-}$ or $\texttt{-} \ominus \texttt{+}$, and hence $\ominus$ behaves exactly as normal subtraction between $+1$, 0, and -1. We say s and t are *weakly separated*, denoted $s||t$, if they are consistent and $s \ominus t$ is segregated.

Example 5.6. Let C_n^k be the collection of all length n sign strings with k pluses and all remaining $n-k$ positions zero. In other words, C_n^k is the collection of all sign strings of sum k supported by $\texttt{+}^n$. The standard bijection between C_n^k and $\binom{[n]}{k}$ preserves our notion of weak separation: if $S, T \in \binom{[n]}{k}$ are weakly separated, then the corresponding sign strings s and t are weakly separated.

As an explicit example, let $S = \{1, 3\}$, $T = \{1, 4\}$, and $U = \{2, 4\}$. Then the corresponding sign strings are $s = \texttt{+0+0}$, $t = \texttt{+00+}$, and $u = \texttt{0+0+}$. We see that S and T are weakly separated, and indeed, $s \ominus t = \texttt{00+-}$ is segregated. However, S and U are not weakly separated, as $s \ominus u = \texttt{+-+-}$ is not segregated.

Weak Separation for Sign Strings

In the mixed setting, it is still clear that we can form a bijection between $\mathcal{C}_{n,m}^k$ and some collection of sign strings that we will call $C_{n,m}^k$ supported by $\texttt{+}^n\texttt{-}^m$, though we must be careful if we wish for our two notions of weak separation to line up. The desired bijection between $S \in \mathcal{C}^k_{n,m}$ and the corresponding sign string $s \in C^k_{n,m}$ is:

1. If S is a Plücker subset (i.e., a k-element subset of $[n]$), then s is the length $(n+m)$ sign string whose first n positions indicate with pluses whether the corresponding positive number is in S, and whose remaining m positions are all minuses. For example, $S = \{1, 2, 4\} \in \mathcal{C}^3_{4,4}$ corresponds to $s = \texttt{++0+----} \in C^3_{4,4}$.

2. If S is a dual Plücker subset (i.e., a k-element subset of $[-m]$), then s is the length $(n+m)$ sign string whose first n positions are all pluses, and the last m positions, *in reverse order*, indicate the negative elements of S. For example, $S = \{-1, -2, -4\} \in \mathcal{C}^3_{4,4}$ corresponds to $s = \texttt{++++-0--} \in C^3_{4,4}$, whose first, second, and fourth positions *from the end of the string* correspond to the -1, -2, and -4 in S.

3. If S is mixed, then the first n positions of s indicate the positive elements of S and the last m positions, *in reverse order*, indicate with $\texttt{-}$ whether the corresponding negative number is in S. For example, $S = \{1, 4, -4, -5\}$ and $T = \{4, 5, -1, -3\}$ in $\mathcal{C}^3_{5,5}$ have the corresponding sign strings

$$s = \texttt{+00+0--000}$$

$$t = \texttt{000++00-0-}$$

$$s \ominus t = \texttt{+000---+0+}.$$

Notice that $S||T$, and also $s||t$ as $s \ominus t$ is segregated. This is true in general:

Proposition 5.7. *Let $S, T \in \mathcal{C}^k_{n,m}$ correspond to sign strings $s, t \in C^k_{n,m}$. Then $S||T$ if and only if $s||t$.*

Proof. We divide into cases based on which of S and T are mixed, assuming that neither S nor T is the empty set (in which case the statement is trivial).

- Both S and T are unmixed of the same sign. Then this situation is exactly the unmixed setting, where the preservation of weak separation across this correspondence is clear.
- Both S and T are unmixed of different signs. Then S and T are automatically weakly separated, and $s \ominus t$ uses only one kind of sign and is thus automatically segregated.
- Exactly one is mixed. Without loss of generality, let T be mixed and S be positive. Then S and T will be weakly separated if and only if there does not exist $a, c \in S \backslash T$ and $b \in T \backslash S$ such that $a < b < c$ (as any negative element of T would act as the d used in the definition of weak separation). As sign strings, we see that $s \ominus t$ will have a minus somewhere towards the end (as we assume that $n, m \geq k > |T_-|$, and hence T cannot cancel out all the negative signs in s), so it will be segregated if and only if the first n characters avoid the pattern `+-+`. This is easily seen to be equivalent to the above condition on S and T.
- Both are mixed. Then it is straightforward to see that the difference $s \ominus t$ has a `+` precisely at the positions of $S \backslash T$'s positive elements and $T \backslash S$'s negative elements (i.e., at all the elements of $S_+ \cup T_- \backslash T_+ \cup S_-$), and likewise for the negative ones. Hence $s \ominus t$ has the pattern `+-+-` or `-+-+` if and only if $S_+ \cup T_-$ and $T_+ \cup S_-$ have interlacing elements as in the definition of weak separation, as desired. $\square$

Properties of Differences

We will now begin to study collections of pairwise weakly separated sign strings. The following straightforward lemma summarizes important properties of differences:

Lemma 5.8. *Let s, t, u be pairwise consistent sign strings.*

1. $s \ominus u$ and $t \ominus u$ are consistent.

2. $(s \ominus u) \ominus (t \ominus u) = s \ominus t$.

3. $s||t$ if and only if $(s\ominus u)||(t\ominus u)$. In particular, if s, t, and u are pairwise weakly separated, then so are $s \ominus u$ and $t \ominus u$.

Proof. 1. Consider any fixed position in the sign string. We will show that $s \ominus u$ and $t \ominus u$ do not disagree in sign at this position. If all three of s, t, u are 0 at this position, then the differences will both be zero there, fine for consistency. Suppose, then, that one of them is nonzero, say +, in this position. By consistency, it is impossible for any of s, t, u to be - in this position. In order to end up with a + sign in $s\ominus u$ or $t\ominus u$, then, we must have that u is 0 in this position. However, it is now impossible for $s \ominus u$ or $t \ominus u$ to be - in this position, as in both cases, we have a nonpositive minus 0, never -.

2. Since $\ominus$ behaves as ordinary subtraction as long as it is applied to consistent sign strings (which we know by the previous part of this proposition), this is clear.

3. This follows immediately from the previous part of the proposition, as we have that $(s \ominus u) \ominus (t \ominus u)$ is segregated if and only if $s \ominus t$ is (since they are in fact equal).

□

Example 5.9. Let p be a sign string without zeros, and let s be any length n sign string supported by p. Let p_- be p with each + replaced by 0. Suppose $s\ominus p_-$ consists of one (cyclically) contiguous subsequence of pluses, with zeros everywhere else. As a contiguous subsequence of pluses, $s \ominus p_-$ is evidently weakly separated from every sign string supported by $p \ominus p_- = +^n$. We conclude by the third part of Lemma 5.8

that s is weakly separated from every sign string supported by p. We call such s *contiguous with respect to* p; hence we have demonstrated that contiguous substrings are weakly separated from any sign string with the same support.

Given a collection C of pairwise consistent sign strings and a sign string $s \in C$, we will want to consider the collection $C \ominus s = \{t \ominus s : t \in C\}$. Lemma 5.8 immediately implies the following important corollary:

Corollary 5.10. *Let C be a collection of pairwise weakly separated sign strings, and let $s \in C$. Then $C \ominus s$ is also a collection of pairwise weakly separated sign strings.*

Remark 5.11. In fact, the same is true if s is not in C to begin with, but the individual elements of $C \ominus s$ will not in general be segregated. $C \ominus s$ can still be pairwise weakly separated without individual elements being segregated, as in this case the all zeros string $\mathbf{0}$ is *not* in $C \ominus s$.

Hence whenever we wish to focus on a single sign string s in such a collection C, we may assume s is $\mathbf{0}$ by considering instead the collection $C \ominus s$.

Adjacency

We will now introduce various notions of nearness.

For the following definitions, fix s to be the all $\mathbf{0}$; by the observation at the end of the previous section, any general s can be translated to this situation.

We say that t is *adjacent* to s if t consists of a single + and a single - (and the rest 0); we then imagine s and t as connected by an edge.

Suppose that t is adjacent to s, and t's plus occurs at position i and its minus at j. We say that u lies *left* of the edge $s \to t$ if u if u is also adjacent to s and either

1. u's plus also occurs at i, but its minus occurs in the (cyclic) interval (i, j); or
2. u's minus also occurs at j, but its plus occurs in the (cyclic) interval (j, i).

If the first case holds, we will sometimes say that u is *positively left* of $s \to t$, and likewise, if the second case holds, we will say that u is *negatively left* of $s \to t$.

We can think of u as the result of moving either (but not both) t's plus or t's minus "leftwards" relative to the opposite sign.

For example, consider the following strings:

$$\begin{aligned} s = \mathbf{0} &= \texttt{00000000} \\ t &= \texttt{00+00-00} \\ u_1 &= \texttt{+0000-00} \\ u_2 &= \texttt{00000-0+} \\ u_3 &= \texttt{00+0-000} \\ u_4 &= \texttt{000+0-00}. \end{aligned}$$

u_1 and u_2 are both positively left of $s \to t$ and u_3 is negatively left of $s \to t$, but u_4 is not (as the + has moved right relative to the -).

We define *right* similarly so that in the above example, u_4 is positively right of $s \to t$. Note u is left of $s \to t$ if and only if u is right of $t \to s$.

If s and t are both contiguous (with respect to some support p) and adjacent, then either $s \ominus t$ or $t \ominus s$ looks (cyclically) like $\texttt{-0}^a\texttt{+0}^b$, supported by $\texttt{--}^a\texttt{++}^b$. Without loss of generality, say $s \ominus t$. Then there is no sign string whatsoever left of $s \to t$, as such a sign string would either need a minus somewhere in the $\texttt{+}^b$ portion of the support or a plus somewhere in the $\texttt{-}^a$ portion of the support. We say, then, that the edge $s \to t$ is on the *left boundary*; by our similar definitions for right, we simultaneously have $t \to s$ on the *right boundary*.

Note that in a collection of pairwise weakly separated sign strings, it is impossible for there to exist simultaneously a u left of $s \to t$ obtained by moving the +, *and* a u'

left of $s \to t$ obtained by moving the -, as u and u' would not be weakly separated. However, in a maximal collection of pairwise weakly separated sign strings, it is almost always the case that some u exists to the left and some v exists to the right of any given $s \to t$. In the unmixed case, the existence of such a u and v is Lemma 10.1 of [9], which we state here in the language of sign strings:

Lemma 5.12 ([9] Lemma 10.1)**.** *Let C be a collection of pairwise weakly separated sign strings with the same sum. Suppose that C is maximal; i.e., that there is no sign string $s \notin C$ with the same sum which is weakly separated from every element of C. Let $s \to t$ be any edge between adjacent elements of C. Then there exists some $u \in C$ left of $s \to t$ unless $s \to t$ is on the left boundary.*

In the next section, we will state and prove a generalization of this result for the mixed setting.

5.3 A Lemma for the Mixed Setting

In the next three sections, we prove Theorem 5.5. We will do this from the perspective of sign strings, so we wish to show that every maximal (by inclusion) collection C of pairwise weakly separated sign strings in $C^k_{n,m}$ has cardinality equal to $k(n+m-k)+2$ (and hence is also maximal by cardinality). As in the unmixed case, it suffices to show that each such collection arises as the collection of face labels of a mixed plabic graph; by Corollary 3.21, any such collection has precisely $k(n+m-k)+2$ faces.

Our strategy for the proof is to rely on our knowledge of (ordinary) plabic graphs and unmixed purity as much as possible. To do this, we will divide the sign strings in our collection into subcollections of three types: the Plücker ones, with sum $k-m$; the dual Plücker ones, with sum $n-k$, and the mixed ones, with sum 0. We will then extend each type of collection to a full maximal collection with that sum such that the whole result, all three extensions together, is pairwise weakly separated. Finally, we will cut and glue together the plabic graph arising from each extended subcollection to make the desired mixed plabic graph.

To carry out the extensions of collections just described, we need to strengthen Lemma 5.12 for use in the mixed setting.

Lemma 5.13. *Let C be a maximal collection of pairwise weakly separated sign strings. Let $s \to t$ be any edge between adjacent elements of C. Then unless $s \to t$ is on the left boundary, there exists some $u \in c$ left of $s \to t$.*

Remark 5.14. This lemma implies Lemma 5.12, which required the collection C to be maximal among sign strings with the same sum. Given such a collection C, we can extend it to a maximal collection C' of pairwise weakly separated sign strings with no condition on the sum. Then given $s \to t$ in C, as long as $s \to t$ is not on the left boundary, this new lemma guarantees we can find $u \in C'$ left of $s \to t$. But since u has the same sum as s and t, u is also in C, as required by Lemma 5.12.

Proof. Even though this lemma is stronger, we follow closely the Oh, Postnikov, and Speyer's proof of their lemma [9, Lemma 10.1] with only slight modifications.

Our goal is to show that there exists some $u \in C$ left of $s \to t$. Assume for the sake of contradiction that no such u exists.

Subtract off an appropriate sign string from all elements of C so that we may assume s and t each consist of a single plus, and re-index so that s's plus occurs at position 1 and t's at position x. Let u_y denote the sign string with a single plus at position $y \in (1, x)$. Then u is positively left of $s \to t$ if and only if it is of the form u_y. Let us consider first what we must have in C to prevent these u_y from being part of our collection.

If v is not weakly separated from u_y, then v must have at least two pluses and one minus. Let ℓ and r be positions of pluses in v such that v has no plus in (ℓ, r) but $y \in (\ell, r)$; then v must also have a minus at position $d \in (r, \ell)$. If further $v \in c$, then v must be weakly separated from s and t, and thus we must have $\ell, r \in [1, x]$. Following [9], we call (v, ℓ, r) a *witness* against u_y.

Since no u_y for $y \in (1, x)$ exists in C, we must have in our collection C such a witness against every u_y. Take a minimal subcollection of C that serve as witnesses against every u_y, and label them $(v_1, \ell_1, r_1), \ldots, (v_q, \ell_q, r_q)$ such that $\ell_1 \leq \cdots \leq \ell_q$. By minimality of the subcollection, we in fact have much more: we have the inequalities $\ell_1 < \cdots < \ell_q$ and simultaneously $r_1 < \cdots < r_q$, as otherwise one of our witnesses would be redundant, contradicting minimality. Since there is a witness against every u_y, we must also have $\ell_1 = 1$, $r_q = x$, and $\ell_{i+1} < r_i$.

We now introduce a total order on any consistent collection of sign strings. For consistent sign strings w_1 and w_2, our total order is defined based on the properties of $w = w_1 - w_2$. To determine which of w_1 and w_2 is larger in this order, we apply the following tests, moving on to the next only when a test is inconclusive:

1. If $\sigma_{(1,x)}(w) > \sigma(w)$, set $w_1 > w_2$; if $\sigma_{(1,x)}(w) < \sigma(w)$, set $w_1 < w_2$.
2. If $\sigma_{(x,1)}(w) > 0$, set $w_1 < w_2$; if $\sigma_{(x,1)}(w) < 0$, set $w_1 > w_2$.
3. If the first nonzero sign of w in $(1, x)$ is a plus, set $w_1 < w_2$.

   If it is instead a minus, set $w_1 > w_2$.
4. If the first nonzero sign of w in $(x, 1)$ is a plus, set $w_1 > w_2$.

   If it is instead a minus, set $w_1 < w_2$.
5. If all of the above are inconclusive, then w is 0 except possibly at 1 and x. Since the sum of w must be 0 to reach this case, there are only three possibilities:

- w is the all zeros string $\mathbf{0}$, in which case $w_1 = w_2$.
- w has a plus at 1, in which case set $w_1 < w_2$.
- w has a minus at 1, in which case set $w_2 < w_1$.

(Note that while it does depend on x, this order remains unchanged if we subtract off the same sign string from every sign string under consideration.)

We must check that this is an order at all; i.e., that transitivity holds. This is straightforward, as σ distributes over subtraction. If $\sigma_{(1,x)}(w_1 \ominus w_2) \leq \sigma(w_1 \ominus w_2)$ and $\sigma_{(1,x)}(w_2 \ominus w_3) \leq \sigma(w_2 \ominus w_3)$, then

$$\begin{aligned}
\sigma_{(1,x)}(w_1 \ominus w_3) &= \sigma_{(1,x)}(w_1 \ominus w_2 \ominus (w_3 \ominus w_2)) \\
&= \sigma_{(1,x)}(w_1 \ominus w_2) - \sigma_{(1,x)}(w_3 \ominus w_2) \\
&\leq \sigma(w_1 \ominus w_2) - \sigma(w_3 \ominus w_2) \\
&= \sigma(w_1 \ominus w_2 \ominus (w_3 \ominus w_2)) = \sigma(w_1 - w_3)
\end{aligned}$$

etc.

Now suppose $w_1 || w_2$ and $w = w_1 \ominus w_2$ contains `+-+` as a substring in $[1, x]$; we will show $w_1 < w_2$. Since $w_1 || w_2$, w cannot contain any minuses in $[x, 1]$, and hence $\sigma_{[x,1]}(w) \geq 0$. It follows that $\sigma_{(1,x)}(w) \leq \sigma(w)$, which by the first test yields $w_1 < w_2$ or is inconclusive. If inconclusive (i.e., if $\sigma_{(1,x)}(w) = \sigma(w)$), then since there are no minus signs in $[x, 1]$, we must have that $[x, 1]$ is identically zero. Hence $\sigma_{(x,1)}(w) = 0$, so the second test is also inconclusive. In this case, since the substring `+-+` must be contained entirely in $(1, x)$, the first sign of w in $(1, x)$ must be plus to avoid forming `-+-+`, so the third test finally tells us that $w_1 < w_2$.

Again following [9], we call v_m an (i, j)-*snake* if v_m has a plus at positions $\ell_i, \ell_{i+1}, \ldots, \ell_m$, $r_m, r_{m+1}, \ldots, r_j$ and at least one minus in the interval (r_j, ℓ_i). We now describe how the total order relates to (i, j)-snakes.

Proposition 5.15 ([9] Lemma 10.6). *The minimal witness among $v_i, \ldots, v_j$ under the total order just described is an (i,j)-snake.*

Proof. We show this by induction on $j - i$. Note that (v_i, ℓ_i, r_i) is indeed an (i,i)-snake, establishing the base case.

Now let v_m be the minimal witness among $v_i, \ldots, v_j$. We wish to show (v_m, ℓ_m, r_m) is an (i,j)-snake. Note that m is in at least one of $[i,j)$ and $(i,j]$; without loss of generality, suppose $m \in [i,j)$. Then by the induction hypothesis, (v_m, ℓ_m, r_m) is an $(i, j-1)$-snake.

We wish first to show that v_m does indeed have a plus at r_j. Let v_n be the minimal witness among $v_{m+1}, \ldots, v_j$ and hence an $(m+1, j)$-snake. v_n has a plus at ℓ_{m+1} and v_m does not, and likewise v_m has a plus at r_{n-1} and v_n does not. If v_m failed to have a plus at r_j, then $v_n - v_m$ would have a +-+ substring formed out of ℓ_{m+1}, r_{n-1}, and r_j. By our observation above, this implies $v_n < v_m$, contradicting that v_m is minimal among $v_i, \ldots, v_j$. Therefore, v_m does indeed have a plus at r_j.

Finally, we must show that v_m has a minus in (r_j, ℓ_i). Since v_m is an $(i, j-1)$-snake, it must have a minus in (r_{j-1}, ℓ_i). If this fails to lie in (r_j, ℓ_i), then it lies in (r_{j-1}, r_j), at some position z. Now consider $v_n - v_m$ again. Notice that to be weakly separated from s and t, v_n cannot have a minus at z. Therefore, $v_n - v_m$ has a +-+ substring formed by positions ℓ_{m+1}, r_{n-1}, and z, contradicting that $v_m < v_n$. This contradiction implies that v_m must in fact have a - in (r_j, ℓ_i), as desired. □

We have just seen how the lack of any positively left element to $s \to t$ leads to the existence of a $(1,x)$-snake v_m. We now switch to a dual perspective of sorts. By subtracting off a different element, we could have begun with the assumption that s and t each consist of a single minus rather than a single plus. We have dual notions of witnesses and snakes just by exchanging the role of plus and minus, and our total order in the dual setting remains the same. We therefore have a dual set of witnesses $(v'_1, \ell'_1, r'_1), \ldots, (v'_{q'}, \ell'_{q'}, r'_{q'})$, with a dual $(x, 1)$-snake $v'_{m'}$. For consistency, we will consider every sign string to be in the setting where s and t consist of a single plus, where our $(x, 1)$-snake $v'_{m'}$ has a zero at both position 1 and x.

By definition, the $(x, 1)$-snake $v'_{m'}$ contains at least one plus in $(1, x)$. Hence there exists a minimal witness (v_n, ℓ_n, r_n) such that $v'_{m'}$ contains a plus at some position h in (ℓ_n, r_n).

We claim now that $v_n < v'_{m'}$. Suppose not. Then consider $v = v_n - v'_{m'}$, which therefore cannot contain any +−+ substring. Since $v'_{m'}$ has a plus at h which v_n lacks, we must have that $v'_{m'}$ contains either a plus at all the pluses of v_n in $[1, h)$ or a plus at all the pluses of v_n in $(h, x]$. Without loss of generality, we assume the latter. Since $v'_{m'}$ has a zero at x, so too must v_n. Therefore, v_n is not a (n, q)-snake and hence not the minimum witness among $v_n, v_{n+1}, \ldots, v_q$.

Let p be the smallest index greater than n such that $v_p < v_n$. Since v_n is therefore minimal among $v_n, \ldots, v_{p-1}$, v_n is an $(n, p-1)$-snake, and hence v_n has a plus at r_{p-1}. Since $r_{p-1} \geq r_n > h$ is a plus in v_n, it must also be a plus in $v'_{m'}$. Since $\ell_p < r_{p-1} < r_p$, we have found a plus in $v'_{m'}$ in (ℓ_p, r_p) with $v_p < v_n$, contradicting that v_n was the minimal witness with this property.

As a result, we have the chain of inequalities $v_m \leq v_n < v'_{m'}$, and hence $v_m < v'_{m'}$. But by symmetry, we can also show $v'_{m'} < v_m$. We have achieved a contradiction, and we conclude that there does indeed exist some sign string left of $s \to t$ in C.

$\square$

As a consequence of this lemma, we find that such maximal collections are just unions of maximal collections with the same sum:

Corollary 5.16. *Let C be a maximal collection of pairwise weakly separated sign strings, and let C_σ be its subcollection consisting of all sign strings in C with sum σ. Then C_σ is a maximal collection of pairwise weakly separated sign strings with sum σ.*

Proof. By Lemma 5.13, as in the arguments in [9], we can "grow" a plabic graph starting with the contiguous elements of C_σ whose collection of face labels is precisely C_σ; see Theorem 11.1 of [9] for details. Since plabic graphs whose boundary faces are labeled by contiguous sign strings correspond to the Grassmann permutation, any such collection of face labels is a maximal collection of pairwise weakly separated sign strings. □

Applying this to the unmixed setting by considering C to be a maximal collection of pairwise weakly separated sign strings supported by $+^n$, we immediately obtain the following enumerative result:

Corollary 5.17. *All maximal collections of pairwise weakly separated subsets of $[n]$ have $\frac{1}{6}(n+1)(n^2-n+6)$ elements, exactly $k(n-k)+1$ of which have cardinality k.*

Proof. By Corollary 5.16, the elements of cardinality k form a maximal collection of weakly separated k-element subsets of $[n]$, and hence there are $k(n-k)+1$ of them. We conclude that the total number of elements is

$$\sum_{k=0}^{n}(k(n-k)+1) = \frac{1}{6}(n+1)(n^2-n+6).$$

□

5.4 Extensions of Maximal Collections in $C_{n,m}^k$

We now consider maximal collections of weakly separated sign strings in $C_{n,m}^k$; i.e., those with sum 0 but at most k pluses and at most k minuses; those with k pluses and m minuses; and those with n pluses and k minuses. Our goal is to show that every such collection arises as the collection of face labels of a mixed plabic graph.

First let us fix some conventions. Let C be a maximal collection of weakly separated sign strings in $C^k_{n,m}$ (hence supported by $\texttt{+}^n\texttt{-}^m$), with C_0, C_+, and C_- its mixed (sum 0), Plücker (sum $k-m$), and dual Plücker (sum $n-k$) subcollections respectively. We extend C to a maximal collection C' of weakly separated sign strings each with sum 0, $k-m$, or $n-k$. By Corollary 5.16, $C' = C'_0 \cup C'_+ \cup C'_-$ where each component is a maximal weakly separated extension of the corresponding component of C. Hence the three components correspond to the face labels of three plabic graphs, which we call p'_0, p'_+, and p'_-.

Note that it is important that we begin with C and extend it to C'. If we begin with C' and restrict to $C^k_{n,m}$ to get C, the result may *not* be maximal in $C^k_{n,m}$.

To prove Theorem 5.5, we will need to construct a mixed plabic graph by gluing together parts of three separate plabic graphs. Before we begin, we demonstrate this process by an example.

Example 5.18. Let C consist of the following 17 sign strings in $C^3_{4,4}$, written with their corresponding set in $\mathcal{C}^3_{4,4}$:

$\texttt{00000000} \leftrightarrow \emptyset$

$\texttt{+000000-} \leftrightarrow \{1,-1\}$

$\texttt{0+00000-} \leftrightarrow \{2,-1\}$

$\texttt{000+-000} \leftrightarrow \{1,-4\}$

$\texttt{++00--00} \leftrightarrow \{1,2,-3,-4\}$

$\texttt{++00-00-} \leftrightarrow \{1,2,-1,-4\}$

$\texttt{++000--0} \leftrightarrow \{1,2,-2,-3\}$

`++000-0-` $\leftrightarrow \{1, 2, -1, -3\}$

`+0+00-0-` $\leftrightarrow \{1, 3, -1, -3\}$

`+00+0-0-` $\leftrightarrow \{1, 4, -1, -3\}$

`00++00--` $\leftrightarrow \{3, 4, -1, -2\}$

`00++0-0-` $\leftrightarrow \{3, 4, -1, -3\}$

`+++0----` $\leftrightarrow \{1, 2, 3\}$

`+0++----` $\leftrightarrow \{1, 3, 4\}$

`0+++----` $\leftrightarrow \{2, 3, 4\}$

`++++---0` $\leftrightarrow \{-2, -3, -4\}$

`++++0---` $\leftrightarrow \{-1, -2, -3\}$.

The Plücker sign strings have sum $3 - 4 = -1$, the dual Plücker sign strings have sum $4 - 3 = 1$, and the mixed sign strings have sum 0.

If we extend C to a maximal collection C' of sign strings with sum -1, 1, or 0, the resulting collection should arise as the union of the collections of face labels of three (ordinary) plabic graphs. We can build these plabic graphs with the techniques from [9]; they are shown in Figure 5.1.

The dashed lines in Figure 5.1 separate the part of the graph containing elements of C from the part of the graph containing elements in $C' \backslash C$.

Look at the edges coming up out of the white vertices along the dashed line in the first and second graph. In both graphs, the first edge is traveled upwards by 2's trip and downwards by 4's trip, and the second edge is traveled upwards by 1's trip and downwards by 2's trip.

Similarly, in both the first and third graphs, the edge coming out of the black vertex along the dashed line is traveled upwards by -1's trip and downwards by

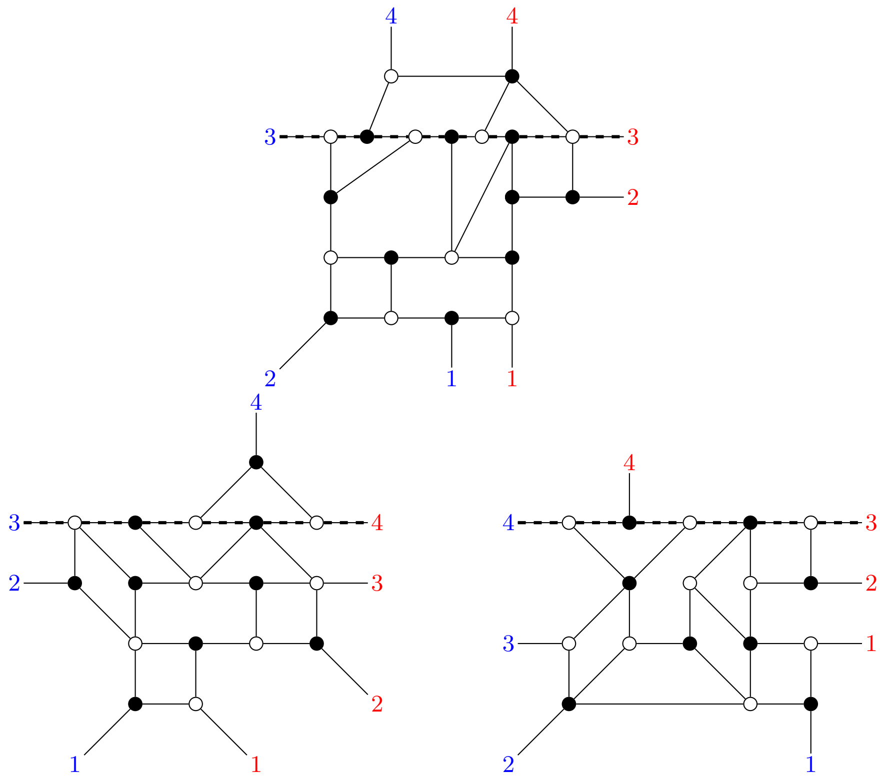

Figure 5.1: The three plabic graphs whose collections of face labels make up the maximal collection C'.

-4's trip.

Now imagine cutting all three graphs through the special edges just discussed, keeping the bottom portion of the first graph and the top portions of the second and third graphs. The face labels that we keep are precisely the elements of C, though we have now broken all of our trips. We can now glue the pieces together along these special edges; since the edges we are gluing together agree about which trips to use, gluing the pieces together fixes the trips so that the face labels from C that we started with are now the face labels of the glued graph. The glued graph is shown in Figure 5.2

Before we prove Theorem 5.5, we need some results that will eventually help us with the cutting and gluing.

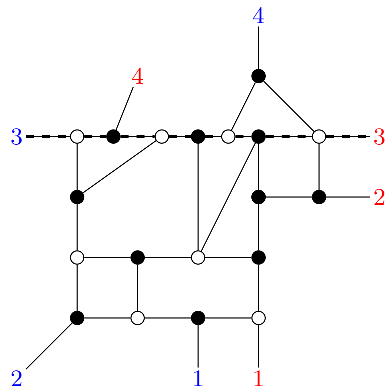

Figure 5.2: The result of gluing the three plabic graphs of Figure 5.1.

Threads in Plabic Graphs

Recall from Section 3.3 that to prove that a certain graph is a mixed plabic graph, it is important that the seam is a thread in the graph. To guarantee that cutting and gluing our three plabic graphs will leave us with a mixed plabic graph, then, we will need to know how the threads in these plabic graphs behave.

Our three constituent plabic graphs are ordinary plabic graphs with (unmixed) Grassmann permutations as their trip permutations, but some of the boundary edges are labeled with positive indices and some with negative indices. As such, we can label the faces according to the rules of mixed plabic graphs, where negative indices contribute to faces to the right of their trips rather than to the left, and we can also define the threading as edges used in one direction by a positive trip and in one direction by a negative trip.

We will need the following proposition:

Proposition 5.19. *Let P be a plabic graph whose trip permutation is the Grassmann permutation π^k_{a+b}, but whose boundary edges are labeled clockwise by the indices $-a, -(a-1), \dots, -1, 1, 2, \dots, b$ rather than $1, 2, \dots, a+b$. Then the threads of P divide the faces into strips according to the number of negative indices in each face* *label.*

Proof. It easily follows from definition that when crossing a thread, the number of positive indices and the number of negative indices in each face label either both go up by one or down by one. Using the terminology of Section 3.3, we wee that to prove the proposition, it suffices to show that the threads of P are simple and parallel. By Proposition 3.7, it suffices just to show that one plabic graph with this trip permutation has simple and parallel threads. Hence without loss of generality, we may assume that P is the quadrilateral plabic graph for π^k_{a+b} of Lemma 3.15, with the boundary label 1 in the bottom left corner.

By symmetry, we can assume $b \geq a$. We prove this proposition by induction on $\min(k, a)$.

The proposition is clear in the two base cases, when either $k = 1$ (or even $k = 0$) or $a = 0$.

Let P' be the quadrilateral plabic graph for π^{k-1}_{a+b-2} with labels from $-(a-1)$ to $b-1$. By the inductive hypothesis, the proposition is true for P'. But we can build P from P' by attaching an L-shaped shell (similar to Figure 3.14) where the threads remain simple and parallel. Hence P also has simple and parallel threads, as desired.

□

Elements in C' Near the Seam

The previous subsection verified that if we cut our three plabic graphs where we want to, then all three cuts occur along threads. In order to make sure we can glue the three cut parts together along a single seam, we need to study the elements in C' near the seam.

If s is a sign string with k minuses and k pluses, define s_+ to consist of the pluses of s along with minuses in all of the last m positions, and likewise define s_- to consist of the minuses of s along with pluses in the first n positions. For example, if $s = \texttt{++0+-00-0} \in C^4_{4,5}$, then

$$s = \texttt{++0+-00-0}$$

$$s_+ = \texttt{++0+-----}$$

$$s_- = \texttt{++++-00-0}.$$

Proposition 5.20. *Let s be an element of C'_0 with k pluses and k minuses. Then $s_+ \in C_+$ and $s_- \in C_-$.*

Proof. We show that any sign string in C weakly separated from s is also weakly separated from s_+ (and, by symmetry, s_-). Then since $s_+ \in C^k_{n,m}$ and C is maximal in $C^k_{n,m}$, we must have $s_+ \in C^k_{n,m}$.

So suppose some $t \in C$ is not weakly separated from s_+; we will show that t is also not weakly separated from s, completing the proof by contradiction. We consider each of the three types of sign strings t could be.

The sign string t cannot be dual Plücker as all dual Plücker sign strings are weakly separated from all Plücker sign strings, and s_+ is a Plücker sign string.

If t is Plücker, then $s_+ \ominus t$ is zero in all the last m positions, so it must contain a forbidden pattern in the first n positions. Since s_+ and s agree in those positions, we see that $s \ominus t$ also contains this pattern and hence t also fails to be weakly separated from s.

Finally, if t is mixed, again consider the last m positions of $s_+ \ominus t$. Since s_+ consists of all minuses there and t has strictly fewer than k minuses, $s_+ \ominus t$ has a mixture of minuses and zeros, with at least one minus. Then for $s_+ \ominus t$ to fail to be segregated, we must have a `+-+` substring in the first n positions. Therefore, $s \ominus t$

also has a +−+ substring there. But since s has k minuses in the last m positions and t strictly fewer, $s \ominus t$ also has at least one minus in the last m positions, completing a +−+− pattern in $s \ominus t$ and proving s is not weakly separated from t. $\square$

Proposition 5.21. *Let $s \to t$ be an edge in C_0' such that s and t each have k pluses and k minuses. Suppose further that there is some $u \in C_0'$ left of $s \to t$ with $k-1$ pluses and $k-1$ minuses (and hence in C_0). If $s_+ \neq t_+$, then there is some $v \in C_+'$ left of $s_+ \to t_+$ with $k-1$ pluses and $m-1$ minuses. If $s_- \neq t_-$, then there is some $v \in C_-'$ left of $s_- \to t_-$ with $n-1$ pluses and $k-1$ minuses.*

Proof. We prove the result when $s_+ \neq t_+$; the situation for $s_- \neq t_-$ is analogous.

Since $s_+ \neq t_+$, s and t differ in the location of one plus and have the same minuses. Cyclically shift our indexing so that s has its extra plus at position 1 and t at position x. Pick some $u \in C_0'$ left of $s \to t$; since u has fewer pluses and minuses than s and t themselves, then the m minuses in our support must lie in the interval $(1, z)$. Let $[x, y]$ be this interval, and let α be the location of the minuses of s missing from u. With the positions defined so far, we have $1 < x \leq \beta \leq y = x + m - 1 < z$.

Let v_1 be the sign string in C' which agrees with s_+ and t_+ except for one plus located at the greatest possible position α in $[1, x)$. Such a v_1 exists, as s_+ itself satisfies the requirements with $\alpha = 1$. Similarly, let v_2 be the sign string in c' which agrees with s_+ and t_+ except for one plus located at the least possible position γ in $(y, z]$. v_2 exists because t_+ satisfies the requirements with $\gamma = z$.

The sign strings v_1 and v_2 are adjacent and $v_1 \to v_2$ is not on the left boundary, so by Lemma 5.13, there is some $v \in C'$ left of $v_1 \to v_2$.

Note that $u \ominus v_1$ has a minus at position α and some pluses in $[x, y]$ (in particular, one at β). Now v cannot be negatively left of $v_1 \to v_2$: in that case $v \ominus v_1$ has only a minus at $\delta \in (\gamma, \beta)$ and a plus at γ, forming a −+−+ in $u \ominus v_1$ at $\alpha, \beta, \gamma, \delta$. This contradicts that v_1 and v are both in C' and hence weakly separated.

Therefore, v must be positively left of $v_1 \to v_2$. Therefore, $v \ominus s_+$ consists of a minus at 1 and a plus somewhere in $[x, y]$. We conclude that v has one fewer plus and one fewer minus than s_+ and hence has the desired $k-1$ pluses and $m-1$ minuses. $\square$

5.5 Proof of Theorem 5.5

Proof. As previously discussed, it suffices to show that given a maximal collection of weakly separated sign strings $C \subset C^k_{n,m}$, we can construct a mixed plabic graph whose collection of face labels is precisely C.

As described in the previous section, construct $C' = C'_0 \cup C'_+ \cup C'_-$ and their corresponding plabic graphs p'_0, p'_+, and p'_-, as guaranteed by Corollary 5.16. By Proposition 5.19, these plabic graphs have a threading which separates its face labels into strips corresponding to the number of minuses in each face label. In particular, p'_0 has a thread T_0 separating the face labels with less than k minuses from those with at least k minuses; i.e., T separates the faces labeled by elements of C from those labeled by elements of C'. Likewise, p'_+ and p'_- have threads T_+ and T_- separating face labels in C from face labels in C'. Let p_0, p_+, and p_- be the plabic graphs formed from p'_0, p'_+, and p'_- by throwing out everything on the "wrong" side of the corresponding thread, keeping only the faces with labels in C. The trips in these new plabic graphs may bear little resemblance to the original trips, but this will be fixed once we glue them together.

We orient all of our plabic graphs so that positive indices travel left to right along threads, and negative indices right to left. It is helpful to think of p_0 as being the "bottom" part of p'_0 and p_+ and p_- as the "top" parts of p'_+ and p'_-.

We claim that Proposition 5.21 allows us to glue together p_0, p_+, and p_- by superimposing T_0, T_+, and T_-, forming a mixed plabic graph which restores the correct face labels (and hence the correct trip permutation). We will focus on gluing T_+ to T_0; attaching T_- to T_0 is analogous.

Incident to T_0 on the p_0 side is a strip of faces labeled by sign strings with $k-1$ pluses and $k-1$ minuses, and on the other is a strip labeled by sign strings with k pluses and k minuses. Let this latter sequence of labels, from right to left, be $s_1 \to s_2 \to \cdots \to s_j$ for some j. By Proposition 5.20, $s_{1,+}, \ldots, s_{j,+}$ are all in C_+, possibly with repeats. Whenever $s_{i,+} \neq s_{i+1,+}$, however, they are adjacent. Hence we have the corresponding strip

$$s_{1,+} = \cdots = s_{j_1,+} \to s_{j_1+1,+} = \cdots = s_{j_2,+} \to \cdots \to s_{j_q+1,+} = \cdots = s_{j,+}$$

in p'_+. By Proposition 5.21, each adjacent pair of labels in this strip has an element with $k-1$ pluses and $m-1$ minuses to its left, and hence each face in this strip must be incident to the thread T_+.

Consider any two such adjacent pairs of labels $s_{i,+} \to s_{i+1,+}$. Since these two face labels differ by some positive indices a and b, the edge between them in p'_+ must be part of a's trip in one direction and b's in the other. Hence this edge must connect to T_+ at a white vertex so that a and b travel along T_+ in the correct direction. Meanwhile, in p'_0, $s_i \to s_{i+1}$ differ by the same positive indices, so the edge between them is also part of both a's trip and b's trip, so this edge is also attached to a white vertex in T_0. Moreover, whenever we have an edge leaving p_0 from a white vertex on T_0, it must be traveled by positive indices (since T_0 is a thread), hence it separates some $s_i \to s_{i+1}$ with different positive indices and hence distinct $s_{i,+}$ and $s_{i+1,+}$. Hence we have a bijection between edges leaving p_0 from white vertices on T_0 and edges incident to T_+ in p_+, and this bijection preserves trips—the same trips use the edge in p_0' as in p_+'. Therefore, when we glue together p_0 and p_+ along corresponding white vertices of T_0 and T_+, all positive trips in the result behave exactly as they did in p_0' and p_+', and hence the positive indices of each face are restored to their desired values.

Similarly, we glue together p_0 and p_- along black vertices, preserving negative trips and restoring the negative indices of each face label.

The result is a mixed plabic graph whose collection of face labels is precisely C, completing the proof. □

CHAPTER 6

Conjectures and Open Questions

6.1 Total Positivity

We quickly review some of the results of [11].

The *totally nonnegative Grassmannian* consists of all points in the Grassmannian where all Plücker coordinates can be chosen to be nonnegative. The totally positive Grassmannian is similarly defined as points where all Plücker coordinates can be chosen to be positive.

Any plabic graph corresponding to the Grassmann permutation can be used to parametrize points in the totally positive Grassmannian as follows. Use the basic rules (M2) and (M3) (monochromatic edge expansion and vertex removal) to transform the plabic graph into a trivalent graph. Then choose a *perfect orientation* of the plabic graph; i.e., orient it so that the outdegree of each white vertex is 2 and the outdegree of each black vertex is 1. Every perfect orientation will result in exactly k sources and $n-k$ sinks along the boundary. Label the sources $1, 2, \ldots, k$ and the sinks $k+1, k+2, \ldots, n$.

If we assign a positive weight to each edge, then we define M_{ij} to be the weighted sum of paths from the source i to the boundary vertex j (source or sink), with an appropriate choice of sign. Then all maximal minors of the $k \times n$ matrix (M_{ij}) are positive, so this matrix corresponds to a point in the totally positive Grassmannian. Moreover, all points in the totally positive Grassmannian arise this way.

Just as the Grassmannian $\mathrm{Gr}(k,n)$ is the variety $\mathrm{Spec}(R_n^k)$, we can define the *mixed Grassmannian* to be $\mathrm{Spec}(R_{n,m}^k)$. Based on the cluster structure in Theorem 4.7, we define the *totally positive mixed Grassmannian* to be the subvariety of the mixed Grassmannian on which all the special invariants (Plücker coordinates, dual Plücker coordinates, and mixed invariants) are positive.

Outside of special cases, no satisfactory intrinsic description of the mixed Grassmannian is known, but perhaps mixed plabic graphs will provide parameterizations of points in the totally positive mixed Grassmannian. It is possible to perfectly orient a mixed plabic graph so that every negative index is a source and every positive index is a sink. Computing the weighted path sums M_{ij} amounts to computing the pairing between the ith covector and the jth vector. To check whether this provides a parameterization of the totally positive mixed Grassmannian, it suffices to verify that all resulting minors of (M_{ij}) of size at most k are positive, and that all points in the totally positive mixed Grassmannian arise in this way.

Are there other trip permutations whose mixed plabic graphs parameterize points in the *totally nonnegative mixed Grassmannian*, where all special invariants are required to be nonnegative rather than strictly positive? Is there a notion of a mixed positroid?

6.2 Generalizing Mixed Plabic Graphs

Recall that, to construct a mixed plabic graph corresponding to a maximal collection of weakly separated sets in $\mathcal{C}_{n,m}^k$, we constructed three separate (unmixed) plabic graphs and glued them along threads. Can we glue together more than three plabic graphs to get more complicated generalizations of mixed plabic graphs? In this section, we describe one way to glue five plabic graphs together.

In Section 5.3, we discussed maximal collections of pairwise weakly separated sign strings in $C^k_{n,m}$, sign strings corresponding to Plücker coordinates, dual Plücker coordinates, and mixed invariants. We now consider a more general collection of sign strings, $C^{k_1,k_2}_{n,m}$, where $0 \le k_1 < k_2 \le n, m$, consisting of all sign strings supported by $\texttt{+}^n\texttt{-}^m$ of any of the following five types:

1. "Big" Plücker sign strings, with exactly k_2 pluses and all m minuses;
2. "Big" dual Plücker sign strings, with exactly k_2 minuses and all n pluses;
3. Mixed sign strings, with k pluses and k minuses, where $k_1 < k < k_2$;
4. "Small" Plücker sign strings, with exactly k_1 pluses no minuses;
5. "Small" dual Plücker sign strings, with exactly k_1 minuses and no pluses.

Note that the sums of each type of sign strings are, respectively, $k_2 - m$, $n - k_2$, 0, k_1, and $-k_1$. We will assume for simplicity that these sums are all distinct (i.e., $k_1 > 0$, $k_2 < m, n$, and $k_1 + k_2 \neq m, n$), though what follows will hold regardless.

Let us consider a maximal collection C of pairwise weakly separated sign strings in $C^{k_1,k_2}_{n,m}$. Following the arguments in Section 5.3, we can extend C' to a maximal collection of pairwise weakly separated sign strings supported by $\texttt{+}^n\texttt{-}^m$. By Corollary 5.16, C' is the union of the collections of face labels of several plabic graphs. The five plabic graphs corresponding to the five sums occurring in C are the plabic graphs we wish to glue together. The gluing process should be analogous to that in Section 5.5, allowing us to construct a sort of "doubly-mixed" plabic graph for any such maximal collection C, with one seam connected the mixed part to the big Plücker and big dual Plücker parts, and another seam connected the mixed part to the small Plücker and small dual Plücker parts (see Figure 6.1).

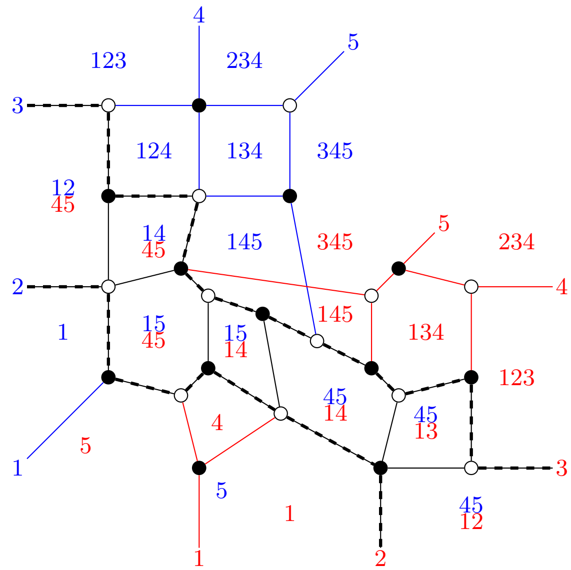

Figure 6.1: A doubly-mixed plabic graph with faces labeled by subsets of $[5] \cup [-5]$. Blue labels are positive; red labels are negative. Doubly-mixed plabic graphs now have two seams, indicated by the dashed lines.

A proof similar to that of Theorem 3.11 should still go through for these doubly-mixed plabic graphs. Up to checking all these details, then, we have all but proved the following conjecture.

Conjecture 6.1. *All maximal collections of weakly separated sign strings in* $C_{n,m}^{k_1,k_2}$ *have the same cardinality.*

We saw that sign strings in $C_{n,m}^{k}$ corresponded to special invariants in a cluster structure on the mixed Grassmannian. Is there a similar interpretation of the sign strings in $C_{n,m}^{k_1,k_2}$ as special elements in some algebra? If so, then these doubly-mixed plabic graphs are likely to correspond to special clusters in a cluster structure on this algebra.

6.3 Generalizing Weak Separation

Mixed weak separation treats sets that decompose into two pieces, the positive and the negative part. Is it possible to generalize further, namely, to consider sets that have three or more components? In this section, we sketch some basic attempts towards such a generalization.

Definition 6.2. Let $n = (n_1, n_2, \ldots, n_\ell)$ be a vector of positive integers, $\ell \geq 2$. A *set of type* (k, n) is a tuple of sets $S = (S_1, S_2, \ldots, S_\ell)$ where either

1. $S_i \subset [n_i]$ and $|S_1| = |S_2| = \cdots = |S_\ell| < k$ ("mixed" sets of type (k, n)) or
2. $|S_i| = k$ for some i, and $S_j = \emptyset$ for $j \neq i$ ("Plücker" sets of type (k, n)).

When considering sets of type (k, n), we will use the notation $S_i' := \{x - n_i - 1 | x \in S_i\}$.

For example, sets of type $(k, (n, m))$ are in bijection with $\mathcal{C}_{n,m}^k$ from Section 5.1. The set $S = (S_1, S_2)$ of type $(k, (n, m))$ corresponds to $S_1' \cup S_2 \subset [n] \cup [-m]$ in $C_{n,m}^k$.

Definition 6.3. Let S and T be two sets of type $n = (n_1, n_2, \ldots, n_\ell)$. Then S and T are *weakly separated* if, for all pairs $1 \leq i < j \leq \ell$, $S_i' \cup S_j$ is weakly separated from $T_i' \cup T_j$.

Note that by the bijection above, this definition agrees with the notion of weak separation in $C_{n,m}^k$.

It is *not* true that maximal collections of pairwise weakly separated sets of type (k, n) have the same cardinality. For example, computer-generated data suggests that most maximal collections of pairwise weakly separated sets of type $(4, (5, 4, 4))$ have cardinality 29, but some (about 2% of the data generated) have cardinality 28.

While it does not seem that a purity result holds for most choices of k and n, the data suggest the following conjectures.

Conjecture 6.4. *All maximal collections of pairwise weakly separated sets of type $(3, n)$ have the same cardinality.*

Notice that in Conjecture 6.4, the length ℓ of the vector n is unconstrained; we require only that the size of each Plücker set is 3, and that each mixed set $S = (S_1, S_2, \ldots, \mathcal{S}_\ell)$ has $|S_1| = |S_2| = \cdots = |S_\ell| < 3$.

Conjecture 6.5. *Let $n = (n_1, n_2, \ldots, n_\ell)$ and $N = n_1 + n_2 + \cdots + n_\ell$. Then the largest cardinality of a collection of pairwise weakly separated sets of type (k, n) is*

$$kN - (k-1)\left(\frac{\ell}{2}k + \ell - 1\right) + 1.$$

In particular, Conjecture 6.5 would imply that if sets of type (k, n) satisfy a purity result, then this is the common cardinality.

Example 6.6. According to Theorem 5.5, maximal collections of pairwise weakly separated subsets of $\mathcal{C}^k_{n,m}$ have cardinality $k(n + m - k) + 2$. This matches Conjecture 6.5, which predicts that such collections have cardinality

$$k(n+m) - (k-1)(k+1) + 1 = k(n+m) - k^2 + 2 = k(n+m-k) + 2.$$

Example 6.7. By Conjecture 6.4, we expect maximal collections of pairwise weakly separated sets of type $(3, (7, 6, 5))$ all to have the same cardinality. All such maximal collections in our computer-generated data have cardinality 42. This agrees with Conjecture 6.5, which predicts that such collections have cardinality

$$3(5+6+7) - (3-1)\left(\frac{3}{2}\cdot 3 + 3 - 1\right) + 1 = 54 - 2 \cdot \frac{13}{2} + 1 = 42.$$

Example 6.8. As mentioned above, sets of type $(4, (5, 4, 4))$ do *not* satisfy a purity condition. Still, most maximal collections appear to have cardinality 29, agreeing with Conjecture 6.5, which predicts a cardinality of

$$4(5+4+4) - (4-1)\left(\frac{3}{2}\cdot 4 + 3 - 1\right) + 1 = 52 - 3 \cdot 8 + 1 = 29.$$

BIBLIOGRAPHY

[1] V. Danilov, A. Karzanov, and G. Koshevoy, On maximal weakly separated set-systems. *J. Algebraic Combin.* 32 (2010), no. 4, 497–531.

[2] S. Fomin and P. Pylyavskyy, Tensor diagrams and cluster algebras, `arXiv:1210.1888`.

[3] S. Fomin and A. Zelevinsky, Cluster algebras I: Foundations, *J. Amer. Math. Soc.* 15 (2002), no. 2, 497–529.

[4] C. Geiss, B. Leclerc, and J. Schröer, Partial flag varieties and preprojective algebras. *Ann. Inst. Fourier (Grenoble)* 58 (2008), no. 3, 825–876.

[5] M. Gekhtman, M. Shapiro, and A. Vainshtein, *Cluster algebras and Poisson geometry*, Mathematical Surveys and Monographs, 167. Amer. Math. Soc., Providence, RI, 2010.

[6] J. Grabowski and S. Launois, Graded quantum cluster algebras and an application to quantum Grassmannians, `arXiv:1301.2133`.

[7] G. Kuperberg, Spiders for rank 2 Lie algebras. *Comm. Math. Phys.* 180 (1996), no. 1, 109–151.

[8] B. Leclerc and A. Zelevinsky, Quasicommuting families of quantum Plücker coordinates, in *Kirillov's seminar on representation theory*, Amer. Math. Soc. Transl. Ser. 2, 181, 85-108, Providence, RI, 1998.

[9] S. Oh, A. Postnikov, and D. Speyer, Plabic graphs and weak separation, `arXiv:1109.4434`.

[10] V. L. Popov and E. B. Vinberg, Invariant theory, in *Algebraic geometry IV, Encyclopaedia of Mathematical Sciences, vol. 55*, Springer-Verlag, Berlin, 1994, 123284.

[11] A. Postnikov, Total positivity, Grassmannians, and networks, `arXiv:math/0609764`.

[12] J. Scott, Grassmannians and cluster algebras, *Proc. London Math. Soc.* 92 (2006), no. 2, 345–380.

[13] J. Scott, Quasi-commuting families of quantum minors, *J. Algebra* 290 (2005), no. 1, 204–220.

[14] H. Weyl, The classical groups. Their invariants and representations, Princeton University Press, Princeton, NJ, 1997 ed.