Flips in Reduced Plabic Graphs

SPUR Final Project, Summer 2018

Julian Wellman
Mentor: Alexey Balitskiy
Project Suggested by Alexander Postnikov and Pavel Galashin
August 1st, 2018

Abstract

Planar bicolored (plabic) graphs are combinatorial objects introduced by Postnikov to give parameterizations of the positroid cells of the totally nonnegative Grassmannian $\mathrm{Gr}^{\geq 0}(n,k)$. Any two plabic graphs for the same positroid cell can be related by a sequence of certain moves. The *flip graph* has plabic graphs as vertices and has edges connecting the plabic graphs which are related by a single move. A recent result of Galashin shows that plabic graphs can be seen as cross-sections of zonotopal tilings for the cyclic zonotope $Z(n,3)$. Taking this perspective, we show that the fundamental group of the flip graph is generated by cycles of length 4, 5, and 10, and use this result to prove a related conjecture of Dylan Thurston about triple crossing diagrams. We also investigate the diameter of the flip graph for the top cell of $\mathrm{Gr}^{\geq 0}(2k,k)$ and of a particular subgraph corresponding to double wiring diagrams. In both cases we prove that the diameter is at least $\frac{1}{2}k(k-1)^2$, and conjecture that this is exact.

1 Introduction

A *flip graph* for our purposes is the graph whose vertices form the set of all diagrams of a particular class, and whose edges correspond to *flips* in these diagrams, which are mutations which transform one diagram into a similar diagram with one small thing changed. There are many natural questions to ask about a flips graph. Is it connected? What do its cycles look like (what is the fundamental group)? What is its diameter, and which pairs of vertices achieve that diameter?

Possibly the most famous example is that of triangulations of an n-gon, whose flip graph forms the 1-skeleton of the *Stasheff associahedron*. Although not novel, a nice corollary of our result is that the fundamental group of the 1-skeleton of the associahedron is generated by cycles of length four and five. Another famous flip graph has *domino tilings* of a planar region as its vertices; in [9] it is proved that the flip graph is connected (provided that the region is simply connected) through a *height function* on tilings, which also gives formula for computing the distance between tilings. Dylan Thurston [8] introduced triple crossing diagrams, which are a generalization of domino tilings, proved that the flip graph is connected, and made a conjecture about the fundamental group of the flip graph. One of the results of this paper is a proof of that conjecture.

We will study several different flip graphs, whose objects and flips are as follows:

- Fine zonotopal tilings of the cyclic zonotope $Z(n, d)$, with flips corresponding to switching between the two tilings of $Z(d + 1, d)$.
- Reduced trivalent plabic graphs for a given connectivity, with flips corresponding to the moves (M1)–(M3) in Figure 3.
- Reduced trivalent plabic graphs for a given connectivity, considered modulo the moves (M1) and (M3). The flips are only given by the square move (M2).
- Triple crossing diagrams for a given connectivity, with flips being $2 \leftrightarrow 2$ moves (see Figure 4).
- Double wiring diagrams on k blue and k red wires as in Figure 5, with flips given by the 2-moves and 3-moves in Figure 6.

The cyclic zonotope $Z(n, d)$ is the Minkowski sum of n vectors in the plane $x = 1$ in $\mathbb{R}^d$ in conxev position, and its fine tilings correspond to projections of the n-dimensional hypercube into $\mathbb{R}^d$. For a complete definition see section 2. In the case $d = 2$, these are the well-studied rhombus tilings of the $2n$-gon, which correspond to (single) wiring diagrams. Wiring diagrams are ways to write the completely inverted permutation as a product of elementary transpositions in S_n, and the flip operation is the *Coxeter move*.

Our main theorem regards generating sets for the fundamental group of the flip graphs when considered as a 1-complex. Though we later phrase our theorems as proving that a 2-complex made out of the flip graph with certain 2-cells glued is simply connected, here we will simply state the sizes of the cycles which generate the fundamental group.

Theorem 1.1. *The fundamental groups for our flips graphs on the following objects are generated by cycles with sizes as follows*

1. *Fine zonotopal tilings of $Z(n,d)$, by cycles of sizes 4 and $2d+4$.*

2. *Reduced trivalent plabic graphs, by cycles of sizes 4, 5 (two types), and 10 (two types).*

3. *Reduced trivalent plabic graphs considered modulo black and white trivalent moves, by cycles of sizes 4 and 5.*

4. *Triple crossing diagrams, by cycles of sizes 4, 5, and 10.*

The first part is the subject of Section 2, and is a result of using Ziegler's results on the higher Bruhat order poset [10] to generalize the proof for $d = 2$ given by Henriques and Speyer [4]. The second result is completely new to our knowledge, and uses the previous result together with the relationship between fine zonotopal tilings and plabic graphs shown by Galashin [3]. This relationship between their flips more firmly established in Section 3, and the result is proven in Section 4. The last two are corollaries of the second, with the fourth result proving a conjecture of Dylan Thurston [8, Conjecture 21], see Section 5. We do not have a result about the fundamental group for double wiring diagrams, but at first glance it appears cycles of lengths 4, 5, and 8 might suffice.

In the case $d = 3$, the fine zonotopal tiling flip graph is generated by squares and decagons. The vertices of the decagons correspond to fine zonotopal tilings of $Z(5,3)$, one of which is shown in Figure 1. The plabic graph cycles of length 5 and 10 appear in the cross-sections of the tilings as shown.

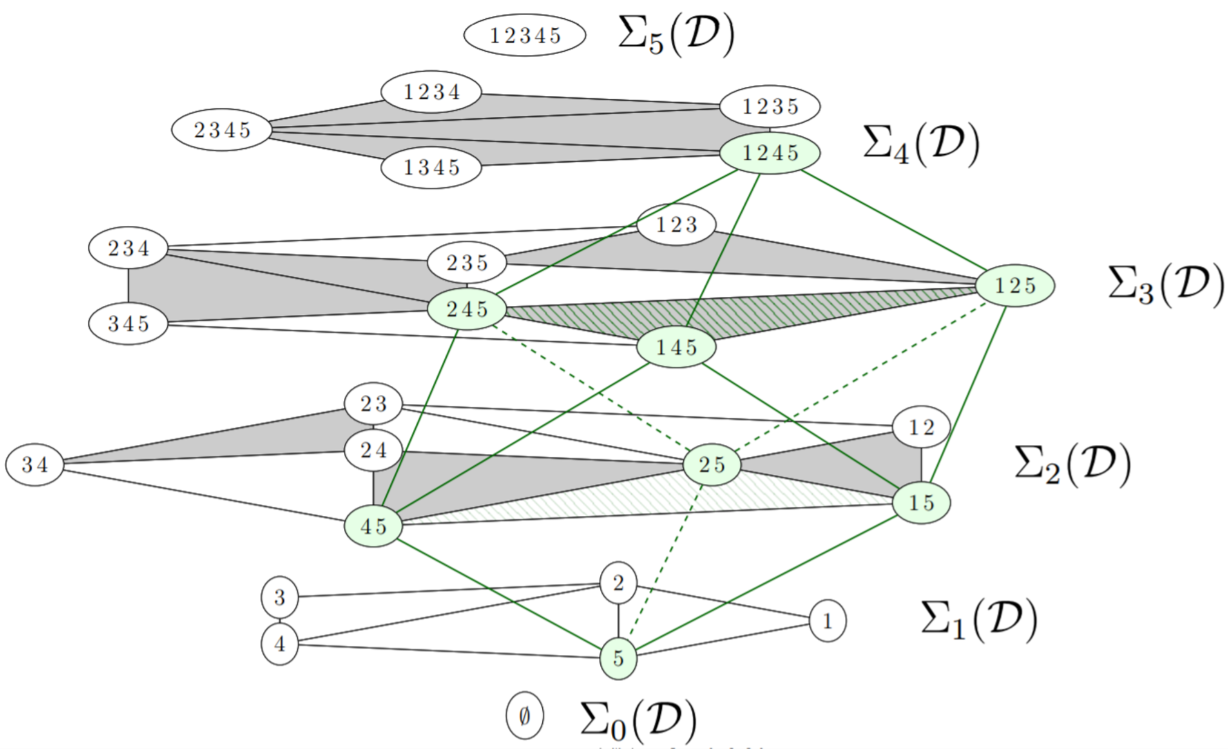

Figure 1: A fine zonotopal tiling of $Z(5,3)$ with the tile $\tau_{(\{5\},\{3\})}$ highlighted in green. The four nontrivial cross-sections shown are planar duals to plabic graphs, and exhibit four of the five types of 2-cells in $X_{\sigma(n,k)}$ (see Theorem 4.5). Figure from Galashin [3]

We then turn to double wiring diagrams, or slightly more generally, reduced trivalent plabic graphs with connectivity $\sigma^{(2k,k)}$, up to moves (M1) and (M3). Here we investigate the diameter of the flip graph. The plabic graphs each give a parameterization of the top cell of the nonnegative Grassmannian $\mathrm{Gr}^{\geq 0}(2k,k)$ [6], and the square moves (M2) give relations between the parameterizations. The diameter then gives a bound on how far the relation might be. Double wiring diagrams were introduced by Fomin and Zelevinsky [2] to give a criterion for a totally positive matrix. The moves between double wiring diagrams give relations between these criteria. [2, Conjecture 19] describes the form these relations might take on, and the conjecture is checked for $k \leq 4$ in [1]. Through personal communication, we learn that Miriam Farber conjectures that the plabic graphs with complementary face labels are exactly the pairs of vertices in the square flip graph at the maximum possible distance, specifically $\frac{1}{2}k(k-1)^2$. Pavel Galashin has independently found an (unpublished) proof that the diameter of the square flip graph is at least $\frac{1}{2}k(k-1)^2$. In Section 6, we prove this lower bound for both plabic graphs and double wiring diagrams, and agree with and extend Farber's conjecture to both cases.

Remark 1. *This manuscript is not a completed version. In particular, there is a small hole which was discovered extremely recently in the proof of Theorem 4.5, which the author has not yet fixed. In its current state, the result is only proven for connectivities $\sigma^{(n,k)}$, rather than any decorated permutation $\pi^{:}$.*

2 Cycles for Zonotopal Tilings

Definition 2.1. *Let $v_1, \ldots, v_n \in \mathbb{R}^d$ be any collection of distinct vectors on the dimension-d moment curve parameterized by $(1, t, t^2, \ldots, t^{d-1})$ The cyclic zonotope $Z(n,d)$ consists of all points which can be written as $\sum_{i=1}^{n} c_i v_i$ for some $\{c_i\}_{i=1}^n \in [0,1]^n$.*

Following the definition used by Galashin [3], we define tilings of the cyclic zonotope as collections of signed subsets. A pair $X = (X^+, X^-)$ of disjoint subsets of $[n]$ is a signed subset of $[n]$, and we also define $X^0 := [n] \setminus (X^+ \sqcup X^-)$. Then the signed subsets are exactly the strings in $\{+,-,0\}^n$. For X a signed subset, the tile τ_X consists of all points which can be written as $\sum_{i \in X^+} v_i + \sum_{j \in X^0} c_j v_j$ for some $\{c_j\}_{j \in X^0} \in [0,1]^{|X^0|}$.

Definition 2.2. *A collection Δ of signed subsets of $[n]$ is called a fine zonotopal tiling of $Z(n,d)$ provided that*

1. *$Z(n,d) = \bigcup_{X \in \Delta} \tau_X$,*
2. *Whenever $\tau_X \cap \tau_Y \neq \emptyset$ for $X, Y \in \Delta$, there exists $Z \in \Delta$ such that $\tau_X \cap \tau_Y = \tau_Z$ is a face of both τ_X and τ_Y, and*
3. *For all $X \in \Delta$, we have $|X^0| \leq d$.*

When the third condition fails, Δ is a zonotopal tiling but is not *fine*. Fine zonotopal tilings Δ of $Z(n,d)$ can be related to each other through a series of mutations. Geometrically, these mutations consist of finding a tiling copy of $Z(d+1,d)$ inside Δ, which has only two fine tilings, and flipping the way it is tiled. We will use the combinatorial definition in terms of signed subsets. Suppose that $S \in \binom{[n]}{d+1}$ has elements $i_1 < \cdots < i_{d+1}$. For $i \in S$, let X_i be the unique signed subset in Δ for which $X_i^0 = S \setminus \{i\}$. Let $S_i^+ := X_i^+ \setminus S$ and $s_i := 1_{X_i^+}(i)$.

Definition 2.3. *A flip of the set S is available in a fine zonotopal tiling Δ if $S_i^+ = S_j^+$ for all $i, j \in S$ and $s_{i_\ell} \neq s_{i_{\ell+1}}$ for all $\ell \in [n = d]$. Performing the flip of the set S results in a new fine zonotopal tiling Δ for which all the signed subsets are identical except that all the values of s_i have changed.*

The *flip graph* is the graph which has fine zonotopal tilings as vertices and edges connecting those tilings which are related by a single flip. It is a fact that any two fine zonotopal tilings of $Z(n,d)$ can be related by a series of flips, so the flip graph is connected, as we will see.

Henriques and Speyer ([4, Proposition 3.14]) prove that the fundamental group of the flip graph of $Z(n,2)$ as a 1-complex is generated by 4-cycles and 8-cycles, where the 4-cycles correspond to pairs of commuting flips and the 8-cycles correspond to copies of $Z(4,2)$. In this section we generalize this result to any dimension using Ziegler's [10] results on the *higher Bruhat order*. Ziegler [10] shows (with different language) that the flip graph for $Z(n,d)$ is isomorphic to the Hasse diagram for the higher Bruhat order graded poset $B(n,k)$ for $k = d$. We will not bother to define the higher Bruhat order, rather, we will state the relevant results about it in the language of fine zonotopal tilings. Flips in zonotopal tilings correspond to covering relations in $B(n,d)$, and the functional ϕ used in [4, Proposition 3.14] on tilings can be related to the rank function on $B(n,2)$.

Theorem 2.4 ([10, Theorem 4.1])**.** *The edges of the flip graph for $Z(n,d)$ form the Hasse diagram for a graded poset with unique minimal and maximal elements Δ_{min} and Δ_{max} at ranks 0 and $\binom{n}{d+1}$. The set of minimal-length paths of flips between Δ_{min} and Δ_{max} modulo commutation of unrelated flips is in natural bijection with the elements of $Z(n,d+1)$, such that flips in tilings of $Z(n,d+1)$ swap the order in which $d+2$ flips occur in the corresponding path.*

It follows from the above that $Z(d+2,d+1)$ has only two fine zonotopal tilings, so $Z(d+2,d)$ has only two paths from $\Delta_{\min}$ to $\Delta_{\max}$ up to commutation. There are also no pairs of commuting flips in tilings of $Z(d+2,d)$, so its flip graph must be a single $(2d+4)$-cycle. We are now ready to characterize the cycles in the flip graph for zonotopal tilings.

Theorem 2.5. *Let $C(n,d)$ be the 2-complex formed by the flip graph for $Z(n,d)$ with the following 2-cells glued:*

- *Quadrilaterals, wherever there is a cycle of length four corresponding to commuting pairs of flips*
- *$(2d+4)$-gons, wherever there is a cycle of length $(2d+4)$ whose vertices are all refinements of a particular zonotopal tiling which is fine except for a single signed subset which creates a tile isomorphic to $Z(d+2,d)$.*

Then $C(n,d)$ is simply connected.

Proof. We will use a technique similar to the proof in [4, Proposition 3.14], and use results about the higher Bruhat order as a black box to generalize to higher dimensions.

Let $\gamma = S_1 S_2 \cdots S_m$, where each S_i is a flip which turns tiling Δ_i into Δ_{i+1} and $\Delta_1 = \Delta_{m+1}$, be a loop in the flip graph for $Z(n,d)$ which connects the tilings $\Delta_1, \Delta_2, \ldots, \Delta_{m+1} = \Delta_1$. It suffices to show that γ can be continuously deformed to a point in $C(n,d)$. All we know is that the squares and the cycles corresponding to the $(2d+4)$-gon from copies of $Z(d+2,d)$ are nullhomotopic, so our only tool is to replace paths in γ with their complement in a square or $(2d+4)$-gon.

First suppose that γ is a cycle of length $2\binom{n}{d+1}$ that includes $\Delta_{\min}$ and $\Delta_{\max}$. Since γ connects the minimal and maximal elements twice in the shortest possible time, it can be divided into two parts, α and β, each of which is a series of monotonic in terms of rank flips in $Z(n,d)$. Then by Theorem 2.4, α and β are each representative elements of some equivalence classes of paths between $\Delta_{\min}$ and $\Delta_{\max}$ given by fine zonotopal tilings A and B of $Z(n,d+1)$, respectively. The flip graph for $Z(n,d+1)$ is connected, so there exists a sequence of flips to transform A into B. Along the way, commutation of flips in α is required to get the right representative element of A, to allow the flips in $Z(n,d+1)$ to be realized as $(2d+4)$-gons in $C(n,d)$. The flips in $Z(n,d+1)$ involve $d+2$ tiles in a copy of $Z(d+2,d+1)$, which appear as $d+2$ flips in α, all inside a copy of $Z(d+2,d)$. Therefore commutation of flips moves α over a quadrilateral, while flips of A involve moves α over a $(2d+4)$-gon. At each step, a continuous deformation of α occurs, eventually transforming it to β, at which point γ is trivial because it is $\beta\beta^{-1}$.

Now suppose γ is any arbitrary cycle as before. Then for each vertex Δ_i in γ, draw a path δ_i of length $\binom{n}{d+1}$ between $\Delta_{\min}$ and $\Delta_{\max}$ which goes though Δ_i, using Theorem 2.4. Let's say that $\delta_i = \delta_i^- \delta_i^+$, where δ_i^+ connects $\Delta_{\min}$ to Δ_i, and then δ_i^- connects Δ_i to $\Delta_{\max}$, both in the shortest possible time. Suppose that the loops $S_i \delta_{i+1}^+ (\delta_i^+)^{-1}$ are all deformable to point. Then after a continuous deformation we could compute to conclude the result

$$[\gamma] = \prod_{i=1}^{m} \delta_i^+ (\delta_{i+1}^+)^{-1} = \delta_1^+ \left(\prod_{i=2}^{m} (\delta_i^+)^{-1} \delta_i^+ \right) (\delta_{m+1}^+)^{-1} = \delta_1^+ (\delta_{m+1}^+)^{-1} = 1.$$

Each flip S_i is either an upward flip or a downward flip, depending on whether Δ_{i+1} has a higher or lower rank than Δ_i when seen in the higher Bruhat order. If it is an upward flip, then $\delta_i^- (\delta_i^-)^{-1} S_i \delta_{i+1}^+ (\delta_i^+)^{-1}$ is a cycle of minimal length which includes $\Delta_{\min}$ and $\Delta_{\max}$, and so is trivial in $C(n,d)$ as we showed in the previous paragraph. If it is a downward flip, then $S_i (\delta_i^-)^{-1} \delta_i^- \delta_{i+1}^+ (\delta_i^+)^{-1}$ is similarly trivial. In either case, the new loop is certainly homotopic to $S_i \delta_{i+1}^+ (\delta_i^+)^{-1}$, completing the proof. $\square$

This result will be essential in Section 4 where we prove an analogous result for plabic graphs.

3 Plabic Moves in Zonotopal Tilings

Let G be an embedding of a planar graph in a disk with each vertex colored white or black (adjacent vertices need not be different colors). Also add n black *boundary vertices* $b_1, \ldots, b_n$

in clockwise order outside of the disk, each with a single edge to one of the vertices of G. This configuration is called a *plabic graph* and we refer to it by G (see Figure 2).

Definition 3.1. *A strand s_i in a plabic graph G is a path which starts at b_i, and proceeds along the edges of G until it reaches some boundary vertex b_j, according to the rules of the road; when s_i reaches a white (resp. black) vertex v through edge e, it makes a sharp left (resp. right) turn. That is, if the edges of v are shown in a circle, then s_i should traverse the next edge clockwise (resp. counterclockwise) of e. The strand permutation of G is the permutation $\pi_G \in S_n$ such that if s_i ends at b_j then $\pi_G(i) = j$.*

We will only deal with a special class of plabic graph. A *bad double crossing* is when two distinct strands both traverse edge e_1 followed by edge e_2.

Definition 3.2 (cf. [6, Theorem 13.2])**.** *A plabic graph G is reduced if and only if*

- *For any edge e between non-boundary vertices, exactly two distinct strands s_i and s_j traverse e.*
- *G does not contain any bad double crossings.*
- *When $\pi_G(i) = i$, the vertex b_i is connected to a single isolated vertex of G.*

The *decorated strand permutation* $\pi^{\cdot}_G$ for G a reduced plabic graph is identical to π_G except that the fixed points of π_G are *decorated* (black) if the single isolated vertex they are connected to is black, otherwise they are *undecorated* (white).

Postnikov described how the *boundary measurements* for reduced plabic graphs provide parameterizations for the positroid cells $S^{\geq 0}_{\mathcal{M}} \subset \mathrm{Gr}^{\geq 0}(n,k)$ [6, Thm. 12.7]. The positroid cell each plabic graph parameterizes depends only on its decorated permutation. The cyclic permutation which sends i to $i+k$ (modulo n) corresponds to the top cell of $\mathrm{Gr}^{\geq 0}(n,k)$ and so is of special interest; we refer to the permutation by $\sigma^{(n,k)}$.

Pavel Galashin [3] shows that the k-th cross-section, $1 \leq k \leq n-1$, of fine tilings of the three dimensional cyclic zonotope correspond to 3-valent reduced plabic graphs with connectivity $\sigma^{(n,k)}$. Let Δ be a fine zonotopal tiling of $Z(n,3)$. The cross-section Σ_k of Δ with the plane $x = k$ in $\mathbb{R}^3$ is a triangulation of an n-gon, possibly with some interior vertices. The vertices of Σ_k are labeled by strings in $\{+,-\}^n$ with exactly k '+' symbols, or equivalently, by elements of $\binom{[n]}{k}$. For any triangle in Σ_k, either the union of the labels of the vertices has $k+1$ elements, or the intersection of the labels of the vertices has $k-1$ elements, depending on the location of the triangle has a cross-section of a single parallelepiped tile. In the first case consider the triangle to be *black*, in the second case consider it *white*. Let G_k be the planar dual to the triangulation Σ_k, and color the vertices of G_k according to the color of the triangle to which it belonged.

Theorem 3.3 (Galashin [3])**.** *G_k is a 3-valent reduced plabic graph with strand connectivity $\sigma^{(n,k)}$. Further, for any 3-valent reduced plabic graph G with strand connectivity $\sigma^{(n,k)}$, there exists a fine zonotopal tiling of $Z(n,3)$ for which $G_k = G$.*

Dually, Σ_k is a triangulation of a *plabic tiling*, but we will stay in the language of plabic graphs. The vertices of Σ_k appear in the faces of G_k, and so we will refer to their labels as the *face labels* of G_k.

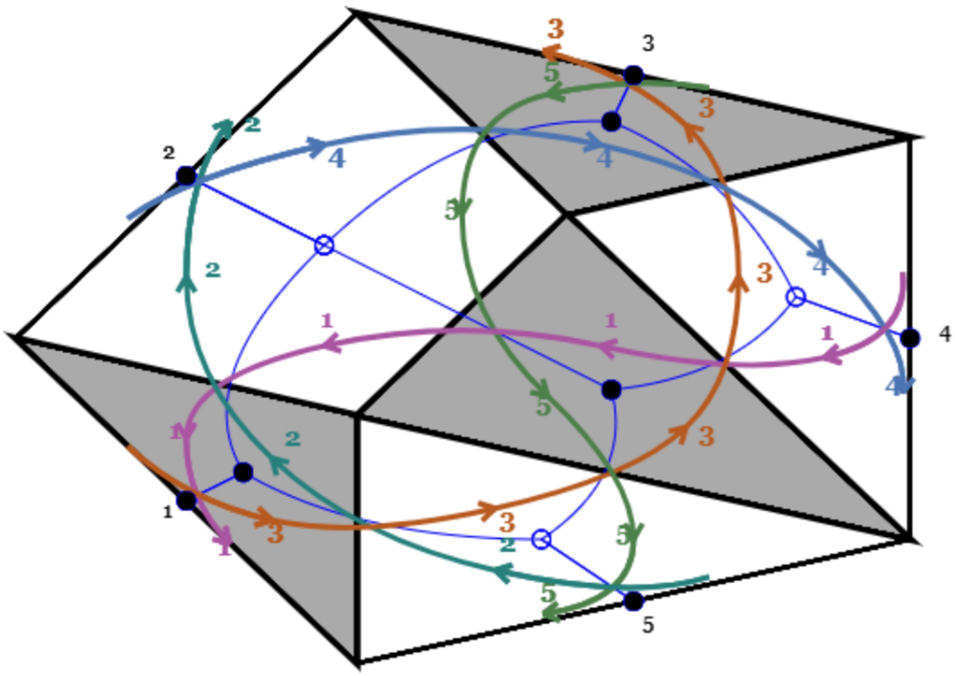

Figure 2: A plabic graph (blue) with connectivity $\sigma^{(5,2)}$. The five strands and the graph's planar dual are also drawn. See `http://math.mit.edu/~galashin/plabic.html`

When Postnikov [6] introduced plabic graphs, he gave some moves to relate them. One can check that the moves in Figure 3 preserve the strand connectivity and whether the plabic graph is reduced.

Theorem 3.4 (Postnikov [6]). *Any two reduced trivalent plabic graphs with the same connectivity can be related by a sequence of the moves (M1), (M2), and (M3) in Figure 3.*

We will primarily deal with trivalent plabic graphs and so only use (M1)–(M3), but the contraction/uncontraction moves will be relevant for triple crossing diagrams. We would like to see how these trivalent plabic moves relate to the three-dimensional cyclic zonotopal flips. Galashin [3] observed that a zonotopal flip at height k performs a square move in G_k, a white trivalent move in G_{k-1}, and a black trivalent move in G_{k+1}.

Lemma 3.5. *For any zonotopal tiling Δ, the available flips are in bijective correspondence with the available square moves in the plabic graphs $\{G_k\}_{k=1}^n$.*

Proof. A proof does not completely appear in [3], so we will include one here. Take any available flip $S \in \binom{[n]}{4}$ for Δ, say $S = \{a, b, c, d\}$ with $a < b < c < d$. Let $k := |S_a^+| + 2$. Then the intersection $\tau_{X_a} \cap \tau_{X_b} \cap \tau_{X_c} \cap \tau_{X_d} =: v$ is a vertex in Δ which is in the cross-section Σ_k. Further, the cross-sections of the tiles τ_{X_i} at level k are triangles in Σ_k which include v as a vertex. The color of the triangle corresponding to X_i is determined by the whether the plane $x = k$ cuts the tile at a lower or higher part, and so depends only on the value of s_i. Then the triangles from X_a, X_c are of one color and X_b, X_d have the other color, by Definition 2.3. Finally, the intersections $\tau_{X_a} \cap \tau_{X_b}, \tau_{X_b} \cap \tau_{X_c}, \tau_{X_c} \cap \tau_{X_d}, \tau_{X_d} \cap \tau_{X_a}$ all appear as

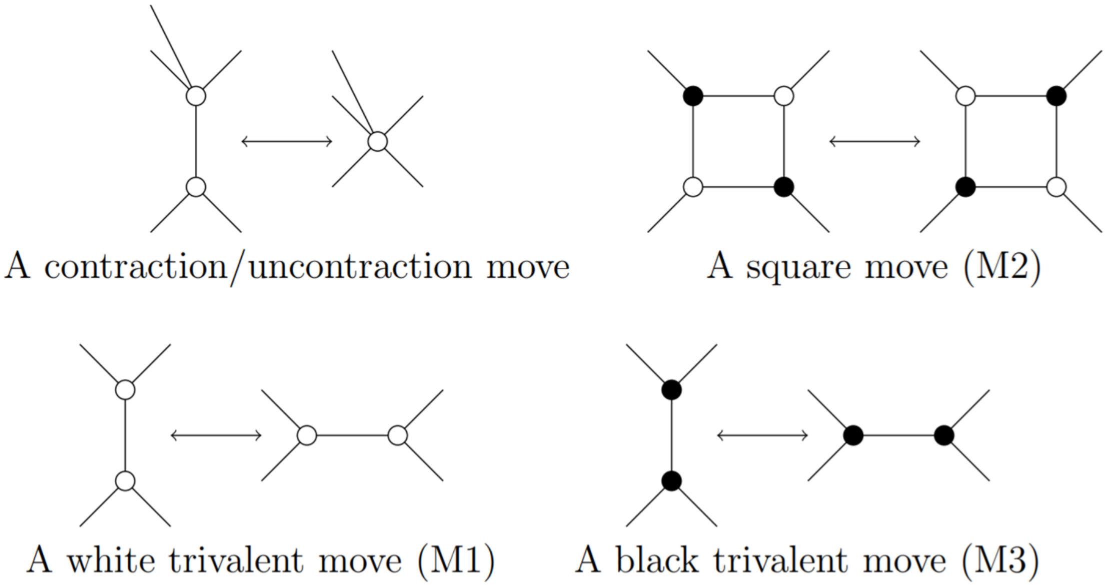

Figure 3: Moves in plabic graphs, from Galashin [3]

edges connected to v in Σ_k, because each of these tiles is a quadrilateral with two vertices at height k, one of which is v. Therefore G_k has a square move pattern formed by the vertices in the four vertices from the four tiles. Performing this flip performs this square move and no other square moves in any other layer.

It now suffices to invert this map. That is, take any available square move in any layer G_k, and recover the unique flip which performs that square move. Well, the square move is formed by four triangles in Σ_k, whose five vertices, when considered as strings in $\{+,-\}^n$, agree in all but four coordinates. This can be seen by noting that all five strings have exactly k '+' symbols, and that when two vertices are adjacent they can only differ in two coordinates. These four coordinates $a < b < c < d$ form our set S, and if the flip corresponding to S is available, then it must correspond to this square move in the map described in the previous paragraph. It then suffices to check that S satisfies the conditions in Definition 2.3. Indeed, $S_a^+ = S_b^+ = S_c^+ = S_d^+$, because the vertices agree on all coordinates outside of S. Now, of the four outer vertices, two are white and two are black, so two of a, b, c, d will have $s_i = 1$ and two will have $s_i = 0$. Moreover, these colors are oriented in a cyclically alternating fashion, and they also correspond to the signed subsets X_a, X_b, X_c, X_d in a cyclic fashion. Therefore we must have $X_a = X_c \neq X_b = X_d$, so we can conclude that S is an available flip in Δ. □

We would like to know when the other two plabic moves can be performed as well. A white or black trivalent move depends on the existence of a square move in a neighboring layer, so the following result about the relationship between the graphs G_k is helpful.

Lemma 3.6 (Galashin [3]). *Let Σ_k be a colored and labeled triangulation for some tiling Δ. Then Σ_{k+1} is fixed up to the triangulation of the white regions and S_{k-1} is fixed up to the triangulation of the black regions.*

Proof. By Galashin's [3, Corollary 4.4], the vertex labels of Σ_{k+1} and Σ_{k-1} are completely determined by Σ_k. The white triangles in S_k cut a tile of Δ which is cut by a black triangle in Σ_{k+1}, and all black triangles in Σ_{k+1} correspond to a white triangle in Σ_k. Similarly, the black triangles in Σ_k give the white triangles in Σ_{k+1}. Therefore all the white and black regions are determined in both Σ_{k+1} and Σ_{k-1}, and indeed all that is left is the triangulation of the white regions in Σ_{k+1} and the black regions in Σ_{k-1}. □

We would like to perform the plabic moves in each layer by doing flips in the tiling which cause them. Unfortunately the appropriate flip isn't always available, but luckily we can set it up without changing the relevant layer. Let Δ be a zonotopal tiling and G_k be a plabic graph formed by a cross-section of Δ.

Lemma 3.7. *Suppose M is a possible black (resp. white) trivalent move in G_k. Then there exists a finite sequence of flips $(S_1, S_2, \ldots, S_m)$ in Δ, such that G_ℓ is unchanged by each of the first $m-1$ mutations for any ℓ at least (resp. at most) k, but the move M occurs on the last mutation.*

Proof. Complementing all of the labels of the vertices doesn't change the structure of the available flips but does change the colors of all of the regions, so it suffices to prove the result when M is a black trivalent move. We proceed by induction on k. When $k \leq 2$, there are no legal black trivalent moves, so the claim holds vacuously. Now, the black trivalent move corresponds to two black triangles in Σ_k, which by Lemma 3.6 creates two white triangles in Σ_{k-1}, which are forced to border two black regions. If the black regions are triangulated such that a square move is legal using the white triangles, then perform the corresponding flip and we're done. Otherwise, there exists a sequence of triangulation flips in the black regions which would make the square move legal. By the inductive hypothesis, each of these flips can be done through a finite sequence of mutations, each of which (except the last) leave G_ℓ unchanged for all $\ell \geq k-1$. The last flip in each sequence performs a black trivalent move in Σ_{k-1}, so also leaves S_k unchanged. Therefore we can set up the square move in Σ_{k-1} without changing Σ_k at all, so the induction is complete. □

4 Main Result on Cycles

For a given decorated permutation $\pi^:$, any two trivalent reduced plabic graphs with connectivity $\pi^:$ can be related by a sequence of the moves (M1)–(M3). The *flip graph* $F_{\pi^:}$ is the graph whose vertices are trivalent reduced plabic graphs with connectivity $\pi^:$ and whose edges connect plabic graphs related by a move. Cycles in the flip graph correspond to sequences of moves which leave the plabic graph unchanged. To understand these cycles we would like to include them in a cycle of flips in a three-dimensional zonotopal tiling, which we understand much better. The cross-sections of zonotopal tilings only correspond to the connectivities $\sigma^{(n,k)}$, so we would like to extend any decorated permutation to one of these. First we need to better understand the relationship between plabic graphs and positroid cells.

Definition 4.1 ([6, Definition 16.1]). *A Grassmann necklace $\mathcal{I} = (I_1, I_2, \ldots, I_n)$ is a sequence of subsets of $[n]$ of the same size such that for all i there exists j such that $I_{i+1} = (I_i \setminus \{i\}) \cup \{j\}$, where as an index i is considered modulo n.*

Grassmann necklaces are in bijection with decorated permutations via *juggling patterns* of period n and throws of height at most n, and all of these are in bijection with positroid cells. Then we can state a useful result of Oh, Postnikov, and Speyer (for our purposes it does not matter what it means for a collection to be weakly separated).

Theorem 4.2 ([5, Theorems 1.3 and 1.5])**.** *Suppose $\mathcal{M} \subset \binom{[n]}{k}$ is a positroid, $\mathcal{I}$ is the corresponding Grassmann necklace, and $\pi^:$ the corresponding decorated permutation. Then the maximal by inclusion weakly separated collections inside $\mathcal{M}$ all contain $\mathcal{I}$ and are exactly the collections of face labels for the plabic graphs with connectivity $\pi^:$.*

Lemma 4.3. *Let $\pi^:$ be a decorated permutation on n elements. Then $\mathcal{F}_{\pi^:}$ is a subgraph of $\mathcal{F}_{\sigma^{(n,k)}}$ for some $k \in [n]$ depending on $\pi^:$.*

Proof. Fix $\pi^:$, and let $\mathcal{I}$ and $\mathcal{M}$ be the corresponding Grassmann necklace and positroid. Let G be any plabic graph with connectivity $\pi^:$. Then by Theorem 4.2, the face labels of G contain the sets in $\mathcal{I}$ and form a weakly separated collection $\mathcal{C} \subset \mathcal{M} \subset \binom{[n]}{k}$ for some $k \in [n]$. Extend $\mathcal{C}$ to a maximal by inclusion weakly separated collection $\mathcal{C}' \subset \binom{[n]}{k}$ in an arbitrary fashion. Then since $\binom{[n]}{k}$ is the positroid cell for the connectivity $\sigma^{(n,k)}$, we can conclude from Theorem 4.2 that $\mathcal{C}'$ forms the collection of face labels for some plabic graph G' with connectivity $\sigma^{(n,k)}$. Now, $\mathcal{C}'$ contains $\mathcal{I}$, which must form the boundary regions for some subgraph of G' which has the same face labels as G. Then any move in G can be realized as a move in G' in the region enclosed by the Grassmann necklace $\mathcal{I}$. Because $\mathcal{F}_{\{\pi^:\}}$ is connected, all edges in $\mathcal{F}_{\pi^:}$ can be realized by performing moves in $\mathcal{F}_{\sigma^{(n,k)}}$ without changing anything in G' outside of the region enclosed by $\mathcal{I}$, which completes the proof. □

Remark: The previous proof could also have been formulated in the language of *Grassmannian graphs*, where the graph G would be seen to have *helicity* k [7]. We will not use the machinery introduced for this proof again, so it was kept to a minimum where possible.

Now we only need to consider cycles in $F_{\sigma^{(n,k)}}$, and can use the cycles in $Z(n,3)$. In order to properly embed cycles as cyclic zonotopal flips, we need to extend Lemma 3.7.

Lemma 4.4. *Let Δ and Δ' be two fine zonotopal tilings of $Z(n,3)$ which are identical on G_k for some fixed k. Then there exists a series of flips, none of which alter G_k, which transform Δ into Δ'.*

Proof. It suffices to find such a sequence of moves which make Δ match Δ' on Σ_{k+1} (without ever changing G_k or any lower layer) and Σ_{k-1} (without ever changing G_k or any higher layer). Once this is done we can recursively match all of the layers to transform Δ into Δ'. By Lemma 3.6, Δ and Δ' already agree up to white (resp. black) triangulation on Σ_{k+1} (resp. Σ_{k-1}). By the flip connectivity of triangulations, there exists a sequence of white (resp. black) trivalent flips in G_{k+1} (resp. G_{k-1}) which transform Δ to completely match Δ' on Σ_{k+1} (resp. Σ_{k-1}). By Lemma 3.7, for each of these flips there exists a finite sequence of flips which perform only this move in G_{k+1} (resp. G_{k-1}, none of which change G_ℓ for any $\ell \leq k$ (resp. $\ell \geq k$). Therefore all of these triangulation moves can be performed without ever changing G_k or any lower (resp. higher) layer, as desired. □

We are now ready to prove our main result.

Theorem 4.5. *Let $\mathbb{X}_{\pi^:}$ be the 2-complex given by the flip graph of trivalent reduced plabic graphs with connectivity $\pi^:$, with the following 2-cells glued to it:*

- *A quadrilateral, wherever there is a 4-cycle generated by two moves occurring in separate parts of a plabic graph,*
- *A pentagon, wherever there is a 5-cycle generated by five white or five black trivalent moves, such that all flips take place in a subgraph which forms a plabic graph with connectivity $\sigma^{(5,1)}$ or $\sigma^{(5,4)}$,*
- *A decagon, wherever there is a 10-cycle consisting of 5 plabic moves alternating with 5 white or 5 black trivalent moves, such that all flips in the cycle take place in a subgraph which forms a plabic graph with connectivity $\sigma^{(5,2)}$ or $\sigma^{(5,3)}$.*

Then $\mathbb{X}_{\pi^:}$ is simply connected for all decorated permutations $\pi^:$.

Proof. Fix any decorated permutation $\pi^:$, and let $\gamma = M_1 M_2 \cdots M_m$ be a loop in $\mathbb{X}_{\pi}^:$ connecting the plabic graphs $G^1_{\pi^:}, G^2_{\pi^:}, \ldots, G^{m+1}_{\pi^:} = G^1_{\pi^:}$. By Lemma 4.3, there exists $k \in [n]$ and plabic graphs $G^1_k, G^2_k, \ldots, G^{m+1}_k = G^1_k$ with connectivity $\sigma^{(n,k)}$ such that γ is a cycle in $\mathbb{X}_{\sigma^{(n,k)}}$ as well. We will show that can contract γ to a point in $\mathbb{X}_{\sigma^{(n,k)}}$, but unfortunately this does not imply that it can be contracted in $\mathbb{X}_{\pi^:}$. This proof is therefore incomplete, and the reader should see Remark1.

Now we construct a loop $Z(\gamma)$ in $C(n,3)$ such that the flips in $Z(\gamma)$ cause exactly the moves $M_1, M_2, \ldots, M_m$ to occur in G_k, in that order. For $i = 0, 1, \ldots, m-1$, there exists Δ_i whose cross-section at height k is exactly G^{i+1}_k. By Lemma 3.5 (if M_{i+1} is a square move) and Lemma 3.7 (if M_{i+1} is a black or white trivalent move), there exists a sequence of moves starting from Δ_i which performs only the move M_{i+1} in G^{i+1}_k. The resulting tiling Δ'_i from this sequence of moves is identical to Δ_{i+1} at height k, so by Lemma 4.4 there exists another sequence flips, none of which cause a move in G_k, which turns Δ'_i into Δ_{i+1}, where $i+1$ is considered modulo m. Concatenating all these sequences of moves results in our loop $Z(\gamma)$ with the desired properties.

By Theorem 2.5, the loop $Z(\gamma)$ is contractible to a point by moving it across the 2-cells in $C(n,k)$. We will show that these 2-cells correspond to 2-cells in $X_{\sigma^{(n,k)}}$ nicely, so that we can also contract γ to a point.

The quadrilaterals in $C(n,3)$ are formed by two commuting flips in $Z(n,k)$, which result in either two moves in separate parts of G_k (a quadrilateral in $\mathbb{X}_{\sigma^{(n,k)}}$), one move being performed twice in G_k (an edge in $\mathbb{X}_{\sigma^{(n,k)}}$), or no moves in G_k (a point in $\mathbb{X}_{\sigma^{(n,k)}}$). In all cases, when $Z(\gamma)$ is moved across a quadrilateral, the image of the quadrilateral in $X_{\sigma^{(n,k)}}$ is a vertex, edge, or 2-cell which γ can also be moved across.

The only other 2-cells in $C(n,3)$ are decagons whose vertices correspond to the ten refinements of an instance of $Z(5,3)$ inside $Z(n,3)$. Depending on where the plane $x = k$ intersects the copy of $Z(5,3)$, one of five things could happen in G_k as the ten flips in the decagon are performed (see Figure 1).

1. If $x = k$ does not intersect the copy of $Z(5,3)$ or only touches the top or bottom vertex, no moves occur in G_k and the image of the decagon in $X_{\sigma^{(k,n)}}$ is a vertex.

2. If $x = k$ intersects the copy of $Z(5,3)$ at relative height 1, then five white trivalent moves occur in a subgraph of G_k with connectivity $\sigma^{(5,1)}$. The image of the decagon in $X_{\sigma(k,n)}$ is a pentagon.

3. If $x = k$ intersects the copy of $Z(5,3)$ at relative height 2, then five square moves and five white trivalent moves occur in a subgraph of G_k with connectivity $\sigma^{(5,2)}$. The image of the decagon in $X_{\sigma(k,n)}$ is another decagon.

4. If $x = k$ intersects the copy of $Z(5,3)$ at relative height 3, then five square moves and five black trivalent moves occur in a subgraph of G_k with connectivity $\sigma^{(5,3)}$. The image of the decagon in $X_{\sigma(k,n)}$ is another decagon.

5. If $x = k$ intersects the copy of $Z(5,3)$ at relative height 4, then five black trivalent moves occur in a subgraph of G_k with connectivity $\sigma^{(5,4)}$. The image of the decagon in $X_{\sigma(k,n)}$ is a pentagon.

In all cases, when $Z(\gamma)$ is moved across the decagon, the image of the decagon is a vertex or 2-cell in $X_{\sigma(n,k)}$ which γ can be moved across.

Finally, let $Z(\gamma)'$ be a deformation of $Z(\gamma)$ by moving it across a 2-cell. We have considered all possible 2-cells in $C(n,k)$ and shown that there always exists a cell in $X_{\sigma(n,k)}$ which γ can be moved across to create γ' such that $Z(\gamma') = Z(\gamma)'$. Therefore by contracting $Z(\gamma)$ to a point in $C(n,k)$ step-by-step while adjusting γ along the way, γ is also contracted to a point. We can conclude that $X_{\sigma(n,k)}$ is simply connected. $\square$

We can consider any two reduced trivalent plabic graphs to be equivalent if they can be related by only white and black trivalent moves. Then there is a *square flip graph*, whose vertices are equivalence classes of plabic graphs for each connectivity $\pi^{:}$, and whose edges connect equivalence classes of graphs which have a pair of representative elements related by a square move. By Theorem 3.4, the square flip graph is connected for every decorated permutation $\pi^{:}$. Our result can be restricted to the square flip graph as follows.

Corollary 4.6. *Let $\mathbb{Y}_{\pi^{:}}$ be the 2-complex given by the square flip graph for plabic graphs with connectivity $\pi^{:}$, with the following 2-cells glued to it:*

- *A quadrilateral, wherever there is a 4-cycle generated by two square moves occurring in separate parts of a plabic graph,*

- *A pentagon, wherever there is a 5-cycle generated by five square moves which take place in a subgraph which forms a plabic graph with connectivity $\sigma^{(5,2)}$ or $\sigma^{(5,3)}$.*

Then $\mathbb{Y}_{\pi^{:}}$ is simply connnected.

Proof. Any loop γ in $\mathbb{Y}_{\pi^{:}}$ can be extended to a loop γ' in $X_{\pi^{:}}$ by adding the necessary extra white and black trivalent moves. Contract γ' to a point step-by-step by moving it across the 2-cells in $\mathbb{X}_{\pi^{:}}$. Each 2-cell in $\mathbb{X}_{\pi^{:}}$ corresponds to either a point or a 2-cell in $\mathbb{Y}_{\pi^{:}}$, so γ may also be continuously deformed while maintaining the correspondence between γ and γ'. Then once γ' has been deformed to a point, so has γ. Therefore $\mathbb{Y}_{\pi^{:}}$ is simply connected. $\square$

5 Triple Crossing and Double Wiring Diagrams

Dylan Thurston [8] introduced *triple crossing diagrams* as a generalization of the domino tilings and their flip operation. Just as the space of domino tilings is flip-connected [9], so is the space of (minimal) triple crossing diagrams with a given connectivity [8]. We will consider only what Thurston [8] calls *minimal* triple crossing diagrams, defined as when introduced by Postnikov to study perfect orientations of plabic graphs [6].

Definition 5.1. *Consider a disk with boundary vertices labeled $b_1, b'_1, \ldots, b_n, b'_n$ in clockwise order. A triple crossing diagram with connectivity (strand permutation) $\pi \in S_n$ consists of n oriented strands drawn inside the disk which start at b_i and end at $b'_{\pi(i)}$ for each $i \in [n]$, satisfying the following properties*

1. *Wherever two strands intersect, exactly three distinct strands meet in a triple crossing.*
2. *When considered in cyclic order, the orientation of the six rays from any triple crossing alternates.*
3. *The diagram contains no bad double crossings, defined as when two distinct strands both arrive at triple crossing c_1 followed by triple crossing c_2.*

It follows from [8, Theorem 7] that this definition is equivalent to Thurston's definition for minimal triple crossing diagrams.

Similar to plabic graphs, triple crossing diagrams on n strands have a connectivity $\pi \in S_n$ given by the final positions of the strands. There is also a notion of flip in a triple crossing diagram, the $2 \leftrightarrow 2$ move, shown in Figure 4.

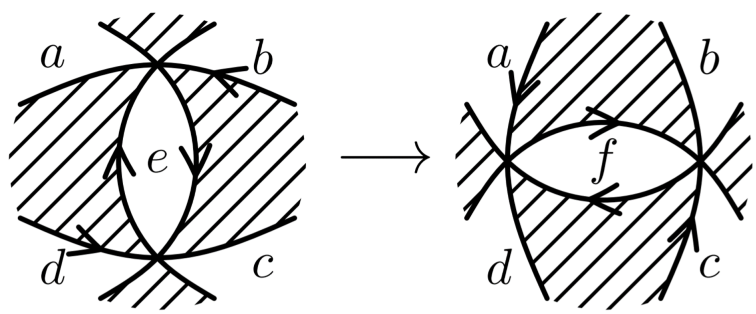

Figure 4: A $2 \leftrightarrow 2$ move in a triple crossing diagram with a clockwise interior region. Figure taken from Thurston [8]

Dylan Thurston [8, Theorem 5] proved that all minimal triple crossing diagrams with the same connectivity can be related by a series of $2 \leftrightarrow 2$ moves.

Postnikov gave the following correspondence between triple crossing diagrams and plabic graphs. For any triple crossing diagram D, the plabic graph $\phi(D)$ has white vertices corresponding to triple crossings in D, black vertices corresponding to regions bounded by counterclockwise-oriented strands that aren't in the middle of a possible $2 \leftrightarrow 2$ move, and edges corresponding to counterclockwise regions bordered by triple crossings and to triple crossing which could be involved in a $2 \leftrightarrow 2$ move together.

Lemma 5.2 ([6, Lemma 14.4]). *The map ϕ described above gives a bijection between triple crossing diagrams with strand connectivity π and reduced plabic graphs for the connectivity π (fixed points undecorated) with all white vertices trivalent and no edges with both endpoints black.*

Such plabic graphs can be considered to be trivalent plabic graphs where the configuration of the edges between black vertices is arbitrary (choose any sequence of uncontraction moves on the black vertices with degree more than three). We observe that the flips in the two contexts correspond nicely

Lemma 5.3. *Let D and D' be triple crossing diagrams related by a single $2 \leftrightarrow 2$ move in D. Then $\phi(D)$ and $\phi(D')$ are related by a square move and several black contraction/uncontraction moves if the interior region of the $2 \leftrightarrow 2$ move was oriented clockwise, otherwise they are related by a single white trivalent move. Conversely, if G and G' are reduced plabic graphs with all white vertices trivalent, and no edges with both endpoints black which are related by a single white trivalent move or a square move and several black contraction/uncontraction moves, then $\phi^{-1}(G)$ and $\phi^{-1}(G')$ are related by a single $2 \leftrightarrow 2$ move.*

Proof. Examine how ϕ transforms $2 \leftrightarrow 2$ moves and ϕ^{-1} transforms plabic moves locally. □

Dylan Thurston [8] conjectured the following, which we now prove as a theorem

Theorem 5.4. *Let $\mathbb{T}_\pi$ be the 2-complex whose vertices are given by triple crossing diagrams with connectivity π, edges given by $2 \leftrightarrow 2$ moves, and the 2-cells created in the following circumstances*

- *Quadrilaterals, where there are two commuting flips in different parts of the triple crossing diagram.*
- *Pentagons, where there are instances of the triple crossing diagram with connectivity $\sigma^{(5,1)}$ or $\sigma^{(5,3)}$.*
- *Decagons, where there are instances of the triple crossing diagram with connectivity $\sigma^{(5,2)}$.*

Then $\mathbb{T}_\pi$ is simply connected for all permutations π.

Proof. Let γ be a cycle $D_1, D_2, \ldots, D_{m+1} = D_1$ of triple diagrams with connectivity π related by $2 \leftrightarrow 2$ moves. Then let $\phi(\gamma)$ be the cycle $G_1, G_2, \ldots, G_{M+1} = G_1$ of plabic graphs in $\mathbb{X}_\pi$ constructed by Lemma 5.3 which contains $\phi(D_1), \phi(D_2), \ldots, \phi(D_{m+1}) = \phi(D_1)$ in that order. By Theorem 4.5, $\phi(\gamma)$ can be continuously deformed to a point in $\mathbb{X}_\pi$ by moving it across the cells in X_π. By construction of $\mathbb{T}_\pi$ and Lemma 5.3, the cells in X_π correspond to either points or cells in $\mathbb{T}_\pi$. In particular, if $\phi(\gamma)'$ is a deformation of $\phi(\gamma)$ from moving across a cell in X_π, then there is a (possibly trivial) cell in $\mathbb{T}_\pi$ which γ can be moved across to create γ' such that $\phi(\gamma') = \phi(\gamma)'$. Then after deforming $\phi(\gamma)$ to a point in $\mathbb{X}_\pi$ while doing the corresponding deformations to γ, the cycle γ must also have been deformed to a point. Therefore $\mathbb{T}_\pi$ is simply connected. □

Postnikov [6] considers a special class of triple crossing diagrams, called *monotone* triple crossing diagrams, and remarks that for the connectivity $\sigma^{(2k,k)}$ they are in bijection with *double wiring diagrams*, introduced by Fomin and Zelevinsky [2]. One can think of double wiring diagrams as two decompositions of the completely inverted permutation as products of elementary transpositions interlaced, considered modulo some commutation rules, but we will just consider them as strand diagrams. Consider the numbers $1, \ldots, k$ to be *blue* and the numbers $k+1, \ldots, 2k$ to be *red*.

Definition 5.5. *A strand configuration for a reduced plabic graph G with connectivity $\sigma^{(2k,k)}$ is a double wiring diagram on k red and blue strands if strands i and j have exactly one intersection whenever i and j are the same color.*

We depict a double wiring diagram as having k strands of each color which run horizontally except to cross an adjacent strand, when they move up or down a level (see figure 5).

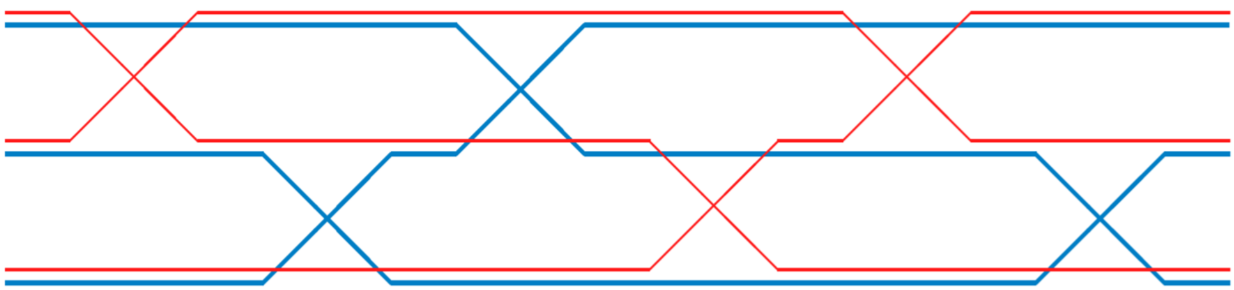

Figure 5: A double wiring diagram for $k = 3$, adapted from [2].

For a double wiring diagram D on k strands of each color, let $G(D)$ be the plabic graph created by replacing blue crossings with a black vertex with an edge down to a white vertex and red crossings with a white vertex down to a black vertex. All vertices lie on one of the k rows and adjacent vertices in a row are connected. Then $G(D)$ is a reduced trivalent plabic graph with connectivity $\sigma^{(2k,k)}$, which has D as its strand diagram.

There are two types of moves on double wiring diagrams; the 3-move, which corresponds to the Coxeter move in the reduced word representation of either single wiring diagram, and the 2-move, which passes two crossings of different colors through each other (see Figure 6).

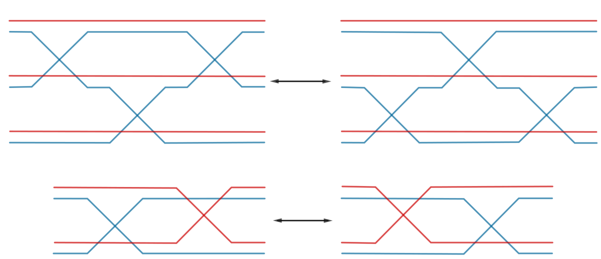

Figure 6: Moves in double wiring diagrams. The 3-move can also be performed with red.

Lemma 5.6. *Suppose D and D' are double wiring diagrams related by a 3-move or a 2-move. Then $G(D)$ and $G(D')$ are related by a square move, possibly with some additional trivalent moves.*

Proof. Examine how the map G transforms 3-moves and 2-moves locally. □

Unfortunately, the converse is false; many square moves cannot be realized in double wiring diagrams, as many reduced plabic graphs do not correspond to a wiring diagram. For the case $k = 3$ however, the flip graphs for double wiring diagrams and the square flip graph for $\sigma^{(6,3)}$ are identical (see [1, Figure 13] and [2, Figure 10]).

6 Square Flip Graph Diameter

Consider the reduced plabic graphs with strand connectivity $\sigma^{(2k,k)}$. The square flip graph is connected, and we wish to investigate its diameter. Paths between vertices of the square flip graph relate the different parameterizations for the top cell of the totally nonnegative Grassmannian $\mathrm{Gr}^{\geq 0}(2k, k)$, so the diameter gives a bound on how complicated the relationship can be. In this section we prove a lower bound on the diameter by constructing an example, and conjecture that this bound is tight.

Any plabic graph can be described by its alternating strand diagram [6, Section 14]. We will label the strands $1, 2, \ldots, k, 1', 2', \ldots, k'$ in counterclockwise order around the disk. Call the primed strands *red* and the unprimed strands *blue*. We also provide a specific ordering on the strands, given by $1 < 1' < 2 < 2' < \cdots < k < k'$.

The strands can be pushed to the horizontal extremes without changing any of the structure, so that all the blue strands are oriented from left to right and the red strands are oriented from right to left. Each strand cuts the disk into a *top* and *bottom* halves relative to that strand. Also, due to the connectivity of the graph being $\sigma^{(2k,k)}$, every pair of strands must intersect an odd number of times. Say that a point of intersection not involving a strand i (a *crossing*) lies *above* strand i if it is in the top half relative to i, otherwise say it is *below* strand i. Fix a strand i and a pair of other strands j, ℓ which are the same color.

Definition 6.1. *We call the pair $(i, \{j, \ell\})$ oriented if exactly one of the following is true:*

- *$j < i < \ell$ or $\ell < i < j$*
- *An odd number of the crossings of strands j and ℓ lie above strand i.*

Otherwise we call the pair $(i, \{j, \ell\})$ unoriented.

We aim to prove the following two facts about oriented pairs.

Lemma 6.2. *There exists a plabic graph for the connectivity $\sigma^{(2k,k)}$ such that for every triple of distinct strands i, j, ℓ with j, ℓ the same color, $(i, \{j, \ell\})$ is oriented. The plabic graph with complementary face labels has all such pairs unoriented.*

Lemma 6.3. *Any square move in a reduced plabic graph changes the total number of distinct oriented pairs $(i, \{j, \ell\})$ in the corresponding alternating strand diagram by at most four.*

Once these have been proven, we will be able to conclude the following lower bound.

Theorem 6.4. *The square flip graph for connectivity $\sigma^{(2k,k)}$ has diameter at least $\frac{1}{2}k(k-1)^2$.*

Proof. There are $2k(k-1)^2$ distinct pairs $(i,\{j,\ell\})$ with i,j,ℓ distinct strands and j,ℓ either both primed or both unprimed. The plabic graphs constructed in Lemma 6.2 then differ by $2k(k-1)^2$ in the total number of distinct oriented pairs. By Lemma 6.3, each square move can only reduce this difference by at most four, so it must take at least $\frac{1}{2}k(k-1)^2$ square moves to connect the two, as desired. □

We can also consider the flip graph whose vertices are double wiring diagrams on k red and blue strands and whose edges connect diagrams related by a single 2-move or 3-move. This is a subgraph of the square flip graph for connectivity $\sigma^{(2k,k)}$, but our result can be extended to it as well.

Corollary 6.5. *The double wiring flip graph on k red and blue strands has diameter at least $\frac{1}{2}k(k-1)^2$.*

Proof. The plabic graphs constructed in the proof of Lemma 6.2 will be of the form $G(D_1), G(D_2)$ for some double wiring diagrams D_1 and D_2. By Lemma 5.6, any sequence of double wiring diagram moves connecting D_1 and D_2 can be converted to a sequence of square moves connecting $G(D_1)$ and $G(D_2)$. Then by Theorem 6.4, this sequence must have length at least $\frac{1}{2}k(k-1)^2$. □

If we consider the completely oriented and completely unoriented plabic graphs to be a sort of minimal and maximal elements, one might hope that we could draw a path of length $\frac{1}{2}k(k-1)^2$ connecting them which passes through any other particular vertex of the square flip graph, and thus prove that the diameter is exactly $\frac{1}{2}k(k-1)^2$. Unfortunately, this is not the case, but nevertheless we still conjecture that our lower bound is tight.

Conjecture 6.6 (Miriam Farber)**.** *Both the square flip graph for connectivity $\sigma^{(2k,k)}$ and the double wiring diagram flip graph on k red and blue strands have diameter exactly $\frac{1}{2}k(k-1)^2$. Moreover, the antipodal pairs in these graphs are exactly the plabic graphs with complementary face labels and the double wiring diagrams which are 180° rotations of each other.*

Interestingly, the number $2k(k-1)^2$ is not only the total number of possible oriented pairs, but it is also the total number of face labels in each plabic graph, or the total number of labels in the chamber minors of the double wiring diagrams (see [2]). We now move on to the proofs of our two lemmas.

Proof of Lemma 6.2. The desired plabic graphs are given by a particular double wiring diagram, which for $k=4$ is depicted in figure 7. In general, if the rows are numbered 1 through $k-1$ from bottom to top, row x consists of x blue intersections and x red intersections, and they alternate color from left to right, with a blue crossing appearing first. This completely determines the double wiring diagram, because the relative orderings of the rows is forced in order to meet Definition 5.5. Call this diagram D_1. Take any triple of distinct strands i,j,ℓ. Each pair among the three intersect at exactly one point (check this). Further, the three intersections form a triangle with exactly one strand running above the intersection

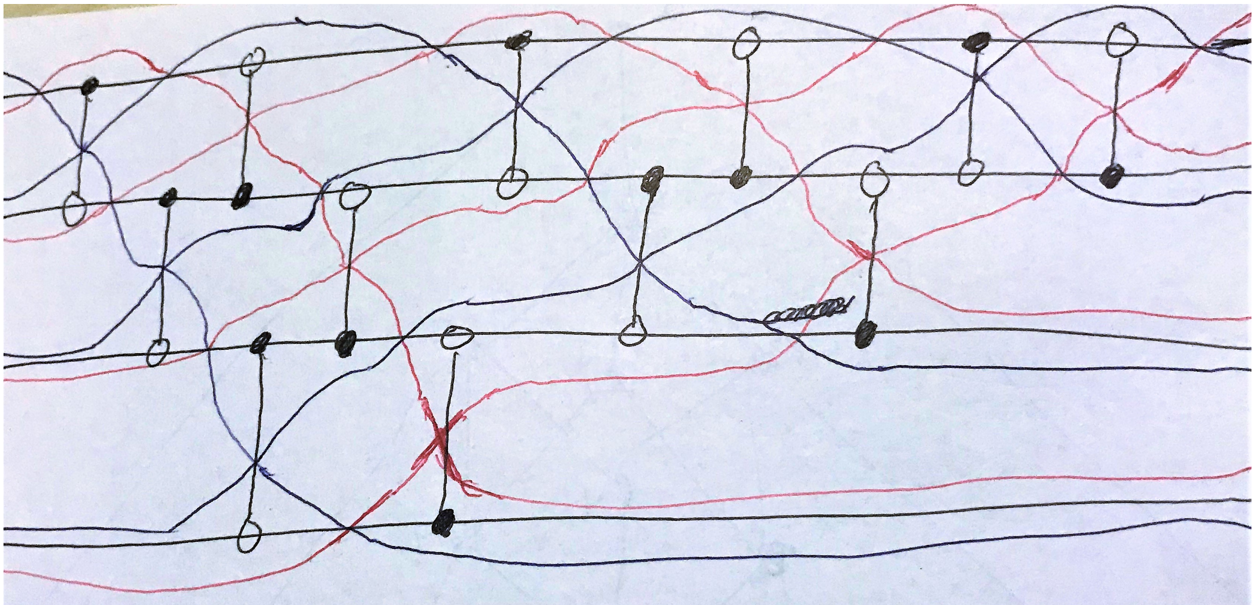

Figure 7: Double wiring diagram D_1 for $k = 4$ with all pairs oriented. The corresponding plabic graph $G(D_1)$ is shown in black behind the double wiring diagram. In from top to bottom starting from the upper-left, the endpoints of the strands should be labeled $1, 1', \ldots, 4, 4'$

of the other two, and in the ordering on the far left, this strand must have started between the other two. Since the strands on the far left are labeled from top to bottom in ascending order according to our ordering $1 < 1' < \cdots < k < k'$, it follows by checking Definition 6.1 that every pair $(i, \{j, \ell\})$ must be oriented.

Any double wiring diagram can be rotated $180°$ in order to get a new double wiring diagram, because every pair of same-colored strands still cross. This is equivalent to reversing the order of the occurrence of all of the inversions. Moreover, if an odd number of crossings of strands j and ℓ lied above i before rotating, then after rotating and even number of these crossings lie above i. Therefore if D_2 is D_1 rotated by $180°$, then none of its pairs are oriented, because the status of whether i is between j and ℓ is unchanged while the second bullet in Definition 6.1 is always changed. The plabic graphs $G(D_1)$ and $G(D_2)$ has strand identical to those in the wiring diagrams, so their pairs have the same orientations. □

Remark: The choice we made that $i < i'$ was arbitrary; by switching all the colors of the crossings in the double wiring diagrams D_1 and D_2 we get a pair with every pair oriented, when we instead consider the ordering where $i' < i$.

Proof of Lemma 6.3. Any square move involves exactly four strands, call them i, j, ℓ, m as labeled in the Figure 8 in counterclockwise order. The orientation of any pair $(x, \{y, z\})$ with $\{x, y, z\} \subset \{i, j, \ell, m\}$ distinct and y, z the same color is changed by the square move, while any other pair has its orientation unchanged (check this). Choose an arbitrary strand, say ℓ. We will prove a slightly stronger claim than needed: the square move changes the number of oriented pairs $(\ell, \{x, y\})$ with $x, y \in \{i, j, m\}$ distinct and the same color by exactly one. We are only considering the proof for the strand ℓ, but this choice was arbitrary. Once this

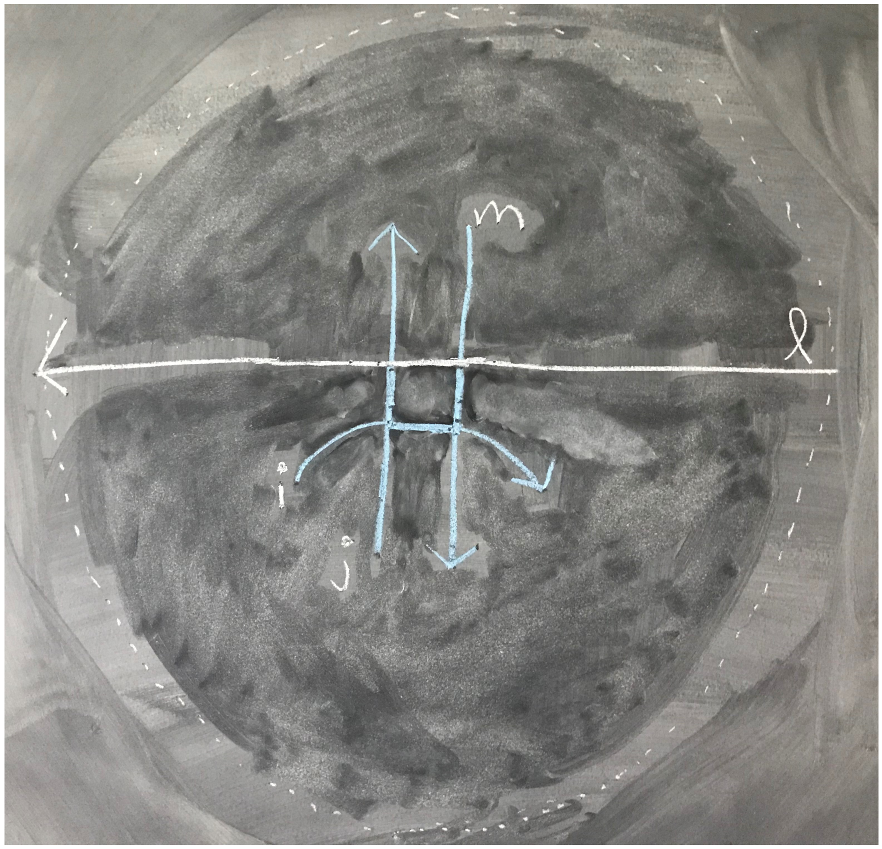

Figure 8: An arrangement of strands in a square move. The start points of i, j, ℓ, m must be in that cyclic (counterclockwise) order.

has been done, we can conclude that the total number of oriented pairs changes by either $0, 2$, or 4 from the square move (in fact it is always 0 or 4).

Well, first suppose that there exists both strands of both colors in $\{i, j, m\}$. Then there is actually only one such pair that can change. It does change, so the claim that the number changes by exactly one holds. Then assume that i, j, m are all blue (we could have chosen red, it doesn't matter). We will show that the three pairs $(\ell, \{i, j\}), (\ell, \{j, m\})$, and $(\ell, \{i, m\})$ cannot have the same orientation. Then since the orientation of all three changes, the number of oriented pairs will change by exactly one, as desired.

Note that the sides of the interior square must not be crossed by any strand for this to be a legal square move. Further, in the diagram of the square move the source endpoints are in the cyclic order i, j, ℓ, m, and for each pair $(i, j), (j, \ell), (\ell, m), (m, i)$ it is impossible for the strands to have crossed previously without creating a bad double crossing. Therefore the start points of i, j, ℓ, m are in exactly that cyclic (counterclockwise) order.

Let's look at the case where ℓ is red. Then since ℓ came from the right and i, j, m came from the left, and the ordering increases going down the leftmost endpoints, and due to the cyclic order of the endpoints i, j, ℓ, m, we can conclude that $m < i < j$.

Call the side of ℓ which contains the depicted crossings of i, j and i, k the "left" side and the other the "right" side. It is unknown whether the left is the same as the top or the bottom half of ℓ. Since j, m must end in a position with j above m, they must cross an odd number of times after crossing ℓ. These crossings must occur to the right of ℓ, because neither j nor m may cross ℓ at a point to their respective lefts.

Also, there cannot be an odd number crossings between i, j or i, m which occur to the right of ℓ without causing either a bad double crossing or an endpoint of a strand to end up on the wrong side (check this). Then regardless of whether the right of ℓ is top or bottom, for all three pairs to have the same orientation, either ℓ needs to be between i, m and i, j but not j, m, or ℓ needs to be between j, m but not i, j or i, m. But $m < i < j$, so neither of those are possible. We can conclude that the claim holds when ℓ is red (i.e., not the same color as i, j, m), so now we suppose that ℓ is blue.

We still know the cyclic ordering of i, j, ℓ, m, so the total ordering on the strands is a cyclic permutation of $i < j < \ell < m$. The strands are all the same color, so all four endpoints must appear in an arc on the disk connecting two adjacent start points of the strands. If we fix the start points of the strands and know which arc the endpoints lie in, then there is a unique way up to adding pairs of extra crossings to wrap the endpoints into the appropriate quadrant. Consider the cases

1. Arc (i, j). Then $j < \ell < m < i$, and (i, j) is the only pair to cross an odd number of times to the right of ℓ.

2. Arc (j, ℓ). Then $\ell < m < i < j$ and (j, m) is the only pair to cross an odd number of times to the right of ℓ.

3. Arc (ℓ, m). Then $m < i < j < \ell$ and (j, m) is the only pair to cross an odd number of times to the right of ℓ.

4. Arc (m, i). Then $i < j < \ell < m$ and (m, i) is the only pair to cross an odd number of times to the right of ℓ.

The unique pair always being on the right (rather than left) is a result of putting the square on the left of ℓ; we could have put the square on the other side and gotten the opposite results, but it makes no difference. In each case, whether 'right' is considered to be above or below ℓ, the three pairs do not have the same orientation. Therefore in all cases the net number of oriented pairs involving ℓ changes by exactly one when the square move is performed, as desired. □

7 Acknowledgements

This work was done in the MIT Summer Program for Undergraduate Research (SPUR) program, 2018. I thank my mentor Alexey Balitskiy for many helpful discussions and teaching me about many of these topics, as well as noticing when my claims were false. I'd also like to thank Alexander Postnikov and Pavel Galashin for suggesting such an interesting and colorful project. Finally I thank Ankur Moitra and Davesh Maulik for their thoughtful comments and for organizing SPUR.

References

[1] Patrick Dukes and Joe Rusinko. Commutation classes of double wiring diagrams. *Involve, a Journal of Mathematics*, 5(2):207–218, 2013.

[2] Sergey Fomin and Andrei Zelevinsky. Total positivity: tests and parametrizations. *The Mathematical Intelligencer*, 22(1):23–33, 2000.

[3] Pavel Galashin. Plabic graphs and zonotopal tilings. *To appear in Proceedings of the London Mathematical Society*, 2017.

[4] Andre Henriques and David E Speyer. The multidimensional cube recurrence. *Advances in Mathematics*, 223:1107–1136, 2010.

[5] Suho Oh, Alexander Postnikov, and David E Speyer. Weak separation and plabic graphs. *Proceedings of the London Mathematical Society*, 110(3):721–754, 2015.

[6] Alexander Postnikov. Total positivity, grassmannians, and networks. *arXiv preprint math/0609764*, 2006.

[7] Alexander Postnikov. Positive grassmannian and polyhedral subdivisions. *arXiv preprint arXiv:1806.05307*, 2018.

[8] Dylan P Thurston. From dominoes to hexagons. In *Proceedings of the 2014 Maui and 2015 Qinhuangdao Conferences in Honour of Vaughan FR Jones' 60th Birthday*, pages 399–414. Centre for Mathematics and its Applications, Mathematical Sciences Institute, The Australian National University, 2017.

[9] William P Thurston. Conway's tiling groups. *The American Mathematical Monthly*, 97(8):757–773, 1990.

[10] Günter M Ziegler. Higher bruhat orders and cyclic hyperplane arrangements. *Topology*, 32:259–279, 1993.