GROWTH DIAGRAMS FROM POLYGONS IN THE AFFINE GRASSMANNIAN

A Dissertation

Presented to the Faculty of the Graduate School

of Cornell University

in Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy

by

Tair Akhmejanov

August 2018

GROWTH DIAGRAMS FROM POLYGONS IN THE AFFINE GRASSMANNIAN

Tair Akhmejanov, Ph.D.

Cornell University 2018

We introduce growth diagrams arising from the geometry of the affine Grassmannian for GL_m. These affine growth diagrams are in bijection with the $c_{\vec{\lambda}}$ many components of the polygon space $\mathrm{Poly}(\vec{\lambda})$ for $\vec{\lambda}$ a sequence of minuscule weights and $c_{\vec{\lambda}}$ the Littlewood–Richardson coefficient. Unlike Fomin growth diagrams, they are infinite periodic on a staircase shape, and each vertex is labeled by a dominant weight of GL_m. Letting m go to infinity, a dominant weight can be viewed as a pair of partitions, and we recover the RSK correspondence and Fomin growth diagrams within affine growth diagrams. The main combinatorial tool used in the proofs is the n-hive of Knutson–Tao–Woodward. The local growth rule satisfied by the diagrams previously appeared in van Leeuwen's work on Littelmann paths, so our results can be viewed as a geometric interpretation of this combinatorial rule. Similar diagrams appeared in the work of Speyer on osculating flags.

BIOGRAPHICAL SKETCH

Tair Akhmejanov was born in Aktau, Kazakhstan on June 25th, 1990, later growing up in Waltham, Massachusetts. He graduated from the Boston College Honors Program in 2012 with a B.S. in Mathematics and a B.S. in Computer Science. He then completed Part III of the Mathematical Tripos (M.A.St.) at Wolfson College, Cambridge in 2013. He began his Ph.D. at Cornell in 2013 and completed his dissertation under the supervision of Professor Allen Knutson.

To Atashka and Azheka.

ACKNOWLEDGEMENTS

First and foremost, I would like to thank my advisor Allen Knutson for all of the time spent teaching me mathematics, as well as his patience during that process. I am fortunate to have had as a resource his constant willingness and energy to discuss mathematics, and have benefitted not only from his wealth of knowledge, but also from his intuition for good mathematics as a whole.

I would also like to thank my committee members Mike Stillman and Dexter Kozen. In particular, I would like to thank Mike Stillman for many helpful discussions during my first year at Cornell. Thanks to all of my fellow graduate students who made graduate school a great experience and from whom I have learned a great deal, and to the entire Cornell Mathematics Department for a welcoming environment in which to pursue a PhD.

I am grateful to Ed Swartz for the 2011 summer REU at Cornell, my first exposure to doing research mathematics. I would like to thank the faculty at Boston College for giving me a solid undergraduate foundation, in particular Ben Howard, Robert Meyerhoff, and Mark Reeder among many others. Reaching even further back to St. Mark's, I am especially indebted to James Tanton for introducing me to the world of abstract math and the joy of mathematics.

On a personal level, I owe a great deal to my family for supporting me through this process. A special thanks goes to Anne Guenther for her amazing, unwavering emotional support. Last, but not least, I would like to thank all of my friends for always providing a fun outlet outside of the world of math.

TABLE OF CONTENTS

LIST OF FIGURES

CHAPTER 1

INTRODUCTION

The classical Robinson–Schensted correspondence (RS correspondence) is a bijection between permutations and pairs of same-shape standard Young tableaux. It was first written down in [28] by Robinson in an attempt to prove the Littlewood–Richardson rule [24]. Schensted later rediscovered the bijection in [31], with the aim of understanding the longest increasing subsequences of a permutation, a point of view that was extended by Greene in [12]. Knuth further generalized the correspondence in [20] to a weight-preserving bijection between matrices of natural numbers and same-shape semistandard Young tableaux, giving the full Robinson–Schensted–Knuth correspondence (RSK correspondence).

Later in [6], Fomin realized the RS correspondence in terms of growth diagrams where standard Young tableaux are interpreted as saturated chains of partitions. These diagrams are square arrays of vertices, each labeled by a partition, where three of the vertex labels of a unit square determine the label of the fourth vertex via Fomin's local rules. Finally, Roby showed in [28] how to extend the growth diagram perspective to include the full RSK correspondence and to prove other combinatorial bijections. The relevant combinatorics of the RS correspondence is reviewed in Section 2.1.

The combinatorics of the RS correspondence was realized geometrically by Steinberg [35, 36], Hesselink [17], and Spaltenstein [32], using the variety of flags preserved by a fixed nilpotent η, i.e. a Springer fiber. The components of this variety are indexed by standard Young tableaux of shape equal to the Jordan type of η. For two generic flags f and f' in components corresponding

to tableaux T and S respectively, their relative position is the permutation given by the RS correspondence. These results were clarified and extended by van Leeuwen in [40]. In particular, each partition of Fomin's growth diagram is the Jordan type of η restricted to $f_i \cap f'_j$. Hence, van Leeuwen completed this geometric interpretation of Fomin's combinatorial rules. See also the work of Rosso in [29] for further generalizations.

This thesis shows that the combinatorics of growth diagrams arises from the geometry of the affine Grassmannian for GL_m, denoted Gr. More specifically, we study the components of $\mathrm{Poly}(\lambda^1, \ldots, \lambda^n)$, the reducible variety of polygons in the affine Grassmannian with fixed distances given by fundamental or dual-fundamental weights, $\lambda^1, \ldots, \lambda^n$, of GL_m. That is, the configuration space of tuples $(L_0, \ldots, L_{n-1})$ of elements in Gr such that $d(L_{i-1}, L_i) = \lambda^i \pmod n$. Definitions are given in Section 2.2.

Fix a component X of $\mathrm{Poly}(\lambda^1, \ldots, \lambda^n)$. There is an open set in X on which for all i and j the distances $d(L_i, L_j)$ are constant. These are the generic distances of the component X. We show that the generic distances $d(L_i, L_j)$ of a component satisfy the local rule in Theorem 12, whereby three distances determine a fourth. This is the main theorem that allows us to construct affine growth diagrams, which we state in Theorem 12 of Section 3.1.2, but do not prove until Chapter 5. It is the same local rule that appeared in van Leeuwen's work on jeu de taquin in the context of the Littelmann path model [39]. See also [27, 42] for a connection to coboundary categories and crystals.

In Chapter 4, we use the local rule of Theorem 12 to construct affine growth diagrams as follows. Fontaine, Kamnitzer, and Kuperberg showed in [8] that it is enough to know the subset $\{d(L_0, L_i)\}_{i \in [n]}$ of generic distances to index the com-

ponents of $\mathrm{Poly}(\lambda^1, \ldots, \lambda^n)$. The remaining generic distances can be discovered from this initial data by repeatedly applying the local rule, and this is how the growth diagrams are constructed. We define affine growth diagrams in Definition 16 of Section 4.1 and show that they are in bijection with the components of the polygon space $\mathrm{Poly}(\lambda^1, \ldots, \lambda^n)$ in Theorem 17. From this, we note in Corollary 19 that affine growth diagrams are periodic with period n. Similar diagrams were constructed by Speyer in [33] from the geometry of osculating flags, where they are called cylindrical growth diagrams (see also [42]).

In Section 4.2, we show that there is a natural way to mark affine growth diagrams, thereby associating a fixed-point-free involution to each growth diagram. This recovers a bijection of Stanley–Sundaram between oscillating tableaux and fixed-point-free involutions, which is demonstrated in Section 4.3. We show that a special case of affine growth diagrams contains Fomin growth diagrams. This is Corollary 31 and Corollary 37. In this way, we rediscover the RS and RSK correspondences arising from the geometry of the affine Grassmannian (see Figure 4.15).

CHAPTER 2

BACKGROUND

2.1 Combinatorial background

In this section we give some basic combinatorial definitions to establish notation and to recall the RS correspondence and Fomin growth diagrams. The notation is based on that used in [28].

2.1.1 The RS correspondence and Greene's Theorem

A **partition** is a sequence of weakly decreasing non-negative integers $\lambda = (\lambda_1, \lambda_2, \ldots)$ such that only finitely many are nonzero. The **size** of a partition is $|\lambda| = \sum \lambda_i$. If n is the size of λ, then it is said that λ is a partition of n and denoted $\lambda \vdash n$. Partitions are depicted as **Young diagrams**, left-justified arrays of empty boxes with the number of boxes in each row weakly decreasing from top to bottom. We will not distinguish between a partition and its Young diagram. Let λ^t denote the transpose of the partition λ, i.e. λ^t_1 is the size of the first column of λ and so forth. The empty partition of all zeros is denoted by $\emptyset$. A **semistandard Young tableau** is a filling of a Young diagram with natural numbers such that the entries are weakly increasing left to right along rows and strictly increasing down columns. For a tableau T, the shape of T, denoted $\mathrm{sh}(T)$, is the underlying partition. The **content** of a semistandard tableau is the sequence $(n_1, n_2, \ldots)$ where n_i is the number of boxes filled with the number i. A **standard Young tableau** is a semistandard Young tableau with content (1,...,1), i.e. the shape is a partition of n and each of $1, \ldots, n$ shows up exactly once. Equivalently, a

standard tableau T may be viewed as a saturated chain of partitions in Young's lattice, $\emptyset = \lambda^0 \lessdot \lambda^1 \lessdot \lambda^2 \lessdot \cdots \lessdot \lambda^n$ where λ^i is the partition obtained from λ^{i-1} by adding the box with entry i in T. Here $\lambda \lessdot \mu$ denotes that μ covers λ, i.e. μ is obtained from λ by adding a single box.

The RS correspondence is a bijection between permutations in S_n and pairs of same-shape standard Young tableaux where the common shape is a partition of n. It will be denoted as $\pi \leftrightarrow (P(\pi), Q(\pi))$. The usual algorithmic ways to realize this bijection are Schensted insertion [31], the matrix-ball construction [9], Viennot's shadow construction [41], and Fomin growth diagrams [6]. Recall that Schensted originally used the insertion algorithm to prove that the length of the longest increasing subsequence of $\pi \in S_n$ is the length of the first row of the common shape of $P(\pi)$ and $Q(\pi)$ (the one-line notation of π is viewed as a sequence). Furthermore, the length of the longest decreasing sequence of π is the length of the first column.

In [12], Greene extended Schensted's result to disjoint unions of chains. Write a permutation π as a permutation matrix such that row i has a mark in column $\pi(i)$ where rows are ordered top to bottom and columns left to right. The associated poset of π is defined as follows. The elements are the marked squares of the permutation matrix where for a, b two distinct marked squares, $a < b$ as elements of the poset if b is southeast of a. Chains, respectively antichains, in this poset correspond to increasing subsequences, respectively decreasing subsequences, of π when written in one-line notation. For this reason we will let π denote both the permutation and the associated poset.

Let $i_k(\pi)$ be the largest size of the union of k disjoint chains of π, and let $d_k(\pi)$ be the largest size of the union of k disjoint antichains. The following is

due independently to Greene [13] and Fomin [5] (see the survey [3] and [28, Thm 2.4.1]).

Theorem 1 (Greene's Theorem)**.** *Let P be any poset with $i_k(P)$ and $d_k(P)$ defined as above. Let $\lambda_k(P) = i_k(P) - i_{k-1}(P)$ and $\mu_k(P) = d_k(P) - d_{k-1}(P)$. Then $\lambda(P) = (\lambda_1, \lambda_2, \ldots)$ and $\mu(P) = (\mu_1, \mu_2, \ldots)$ are partitions, and $\lambda^t = \mu$.*

The following is a generalization, due to Greene, of Schensted's result for the largest chain and antichain. It relates the previous theorem to RS.

Theorem 2 ([12],[30, p. 104])**.** *Let π correspond to the tableaux (P, Q) under the RS correspondence, and let $\lambda = sh(P) = sh(Q)$ be their common shape. Then*

$$i_k(\pi) = \lambda_1 + \cdots + \lambda_k$$

$$d_k(\pi) = \lambda^t{}_1 + \cdots + \lambda^t{}_k.$$

Consider changing the ordering of the elements of the poset corresponding to π so that now $a < b$ if b is northeast of a. This is equivalent to reflecting the permutation matrix of π about the horizontal line joining the midpoints of opposite sides of the square to get a new permutation matrix π_r, and considering the usual southeast dominant poset structure on π_r. Note that $\pi_r = \pi \cdot w_0$ where $w_0 \in S_n$ is the longest element, i.e. switches i and $n + 1 - i$. We could also reflect π about the vertical line joining the midpoints of opposite sides of the permutation matrix to get $\pi^r = w_0 \cdot \pi$. This is equivalent to considering the poset corresponding to π such that now $a < b$ if b is SW of a. Both operations, i.e. going from π to π_r or from π to π^r, exchange the roles of chains and antichains in the poset, hence following proposition holds.

Proposition 3. *If π corresponds to the partition λ under Theorem 1, then π^r and π_r correspond to λ^t.*

Although this proposition only determines the common shape of the pairs of tableaux corresponding to π^r and π_r under RS, the effects on the P and Q tableaux themselves are also known. More generally, we could ask how all of the symmetries of the square acting on the permutation matrix π, change the tableaux (P, Q) under the RS correspondence. If $\pi \leftrightarrow (P, Q)$, then

$$\pi^{-1} \leftrightarrow (Q, P) \tag{2.1}$$

$$\pi \cdot w_0 \leftrightarrow (P^t, ev(Q)^t) \tag{2.2}$$

$$w_0 \cdot \pi \leftrightarrow (ev(P)^t, Q^t) \tag{2.3}$$

$$w_0 \cdot \pi \cdot w_0 \leftrightarrow (ev(P), ev(Q)), \tag{2.4}$$

where $ev(P)$ is the evacuation tableau of P, i.e. the image of P under the Schützenberger involution. Since π^{-1} corresponds to reflecting π about the main diagonal, these operations generate all of D_4, the symmetries of the square. See [30] or [38] for their proofs, and note that the first is an easy consequence of the growth diagram perspective described below.

2.1.2 Fomin growth diagrams

Fomin [6] gave rules to build up partitions labeling the vertices of an $n \times n$ permutation matrix, such that the vertex labels along the east and south edges give the tableaux in the RS correspondence. The Fomin growth diagram is built up by applying a set of local rules to each unit square of the $n \times n$ matrix. Begin by labeling the vertices of the north and west sides of the matrix with the empty partition as in Figure 2.1. Let λ, μ, ν, ρ be partitions labeling the vertices of a unit

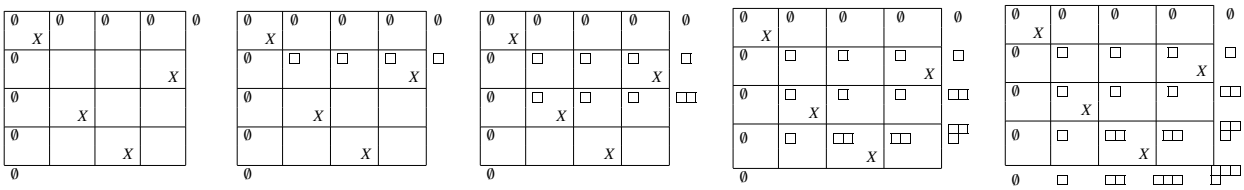

$$\pi = 1423 \leftrightarrow P(\pi) = \begin{array}{|c|c|c|}\hline 1 & 2 & 3 \\ \hline 4 \\ \cline{1-1}\end{array}, Q(\pi) = \begin{array}{|c|c|c|}\hline 1 & 2 & 4 \\ \hline 3 \\ \cline{1-1}\end{array}$$

Figure 2.1: An example of a Fomin growth diagram.

square as in the following figure.

The following Fomin growth rules determine ρ when λ, μ, ν are given. See Figure 2.1 for an example.

1. If $\mu \neq \nu$, then $\rho = \mu \cup \nu$.
2. If $\lambda \lessdot \mu = \nu$, then μ must have been obtained from λ by adding a box to λ_i for some i. Let ρ be obtained from $\mu = \nu$ by adding 1 to μ_{i+1}.
3. If $\lambda = \mu = \nu$, then $\rho = \lambda$ if the current square is unmarked, and $\rho = \lambda$ with 1 added to λ_1 if the current square is marked.

These rules give a deterministic procedure that assigns a partition to each vertex (in Figure 2.1 the diagram is filled in row by row, but of course the order doesn't matter). Reading left to right along the bottom row of vertex labels gives a saturated chain of partitions corresponding to the $P(\pi)$ tableau. Likewise, reading down the rightmost column of vertex labels gives the tableau $Q(\pi)$. Furthermore, the partition at a vertex v is the partition of Theorem 1 for the poset $\pi|_v$ given by the marked squares northwest of v. See [28, App.] for a proof based on Fomin's original in [6]. Thus, Fomin's growth rules build up

the partition of Greene's Theorem by introducing one element of the poset at a time.

In what follows we will need the transpose of the RS correspondence, $\pi \leftrightarrow (P(\pi)^T, Q(\pi)^T)$. In this case, Fomin's growth rules can be modified to build up $P(\pi)^T$ and $Q(\pi)^T$ directly. Simply replace the role played by the rows in rules (2) and (3) to be played by the columns. Namely, for rule (3) add a box in the first column, and for rule (2), if μ was obtained from λ by adding a box in column i, then add a box to column $i+1$ of μ to get ρ. Then each vertex of the growth diagram is labeled by a partition λ such that $d_k(\pi|_v) = \lambda_1 + \cdots + \lambda_k$ and $i_k(\pi|_v) = \lambda^t{}_1 + \cdots + \lambda^t{}_k$, i.e. rows correspond to antichains and columns to chains. With these roles reversed, we will refer to this as the transpose Greene's Theorem. Likewise, Schensted insertion, the matrix-ball construction, and Viennot's shadow construction can be done column-wise rather than row-wise to give $P(\pi)^T, Q(\pi)^T$. When convenient, we will refer to these as the transpose RS correspondence, or the transpose Fomin rules.

2.2 Geometric background

In this section we define the affine Grassmannian, state the geometric Satake correspondence, and define the polygon space for general weights. In the next chapter we specialize these weights to be fundamental and dual fundamental weights.

2.2.1 The affine Grassmannian

Set $G = GL_m$, as we will only be concerned with this case. Recall that the dominant weights of G are weakly decreasing sequences of m integers, $(\lambda_1, \ldots, \lambda_m)$, and that weights and coweights coincide. The dual of a weight λ is given by negating and reversing the sequence, and is denoted λ^*. Let V_λ denote the irreducible representation corresponding to λ. For any representation V, let V^G denote the space of invariants. Although G is not semisimple, we will say that the **minuscule weights** are the fundamental weights $\omega_i = (1, \ldots, 1, 0, \ldots, 0)$ and the dual fundamental weights $\omega_i^* = (0, \ldots, 0, -1, \ldots, -1)$ for $1 \leq i \leq m$ (where there are i many 1's and -1's respectively). Let $\emptyset$ denote the zero weight, $\rho = (m-1, m-2, \ldots, 1, 0)$ the half the sum of positive roots, and $|\lambda| = \lambda_1 + \cdots + \lambda_m$.

Let $\mathcal{O} = \mathbb{C}[[t]]$ be the ring of power series and $\mathcal{K} = \mathbb{C}((t))$ its field of fractions, the field of Laurent series. The **affine Grassmannian** is $Gr = G(\mathcal{K})/G(\mathcal{O})$, an ind-variety over $\mathbb{C}$. It can be seen as the space of $\mathcal{O}$-lattices in $\mathcal{K}^m$. An **$\mathcal{O}$-lattice in $\mathcal{K}^m$** is a free, finitely generated $\mathcal{O}$-submodule L in $\mathcal{K}^m$, that generates the entire vector space $\mathcal{K}^m$ as a $\mathcal{K}$-module, i.e. $\mathcal{K} \otimes L = \mathcal{K}^m$. For example, consider the standard basis vectors $e_1, \ldots, e_m$ of the vector space $\mathcal{K}^m$. The $\mathcal{O}$-module generated by $e_1, \ldots, e_m$ is the **base lattice** L_0. The space of $\mathcal{O}$-lattices has a transitive action of $G(\mathcal{K})$, and the stabilizer of L_0 is $G(\mathcal{O})$, hence gives an identification of $G(\mathcal{K})/G(\mathcal{O})$ with the space of $\mathcal{O}$-lattices. See the examples in [1, §2.1] for a more diagrammatic interpretation of the lattices in the affine Grassmannian.

The $G(\mathcal{O})$ orbits form a stratification of Gr. For each dominant (co)weight λ let t^λ be the corresponding torus fixed point. In terms of lattices this is the $\mathcal{O}$-module generated by $t^{\lambda_1}e_1, \ldots, t^{\lambda_m}e_m$. Let $Gr(\lambda) = G(\mathcal{O})t^\lambda$ be the orbit of $G(\mathcal{O})$ in Gr. These are all of the $G(\mathcal{O})$ orbits, one for each dominant weight. The

stratification is then $Gr = \sqcup Gr(\lambda)$, and $\overline{Gr(\lambda)} = \cup_{\mu \leq \lambda} Gr(\mu)$ are the closure relations, where $\mu \leq \lambda$ means $\lambda - \mu$ is a nonnegative integer combination of simple roots. As is easily seen, the orbits of $G(\mathcal{K})$ on $Gr \times Gr$ are in bijection with orbits of $G(\mathcal{O})$ on Gr. For $p, q \in Gr$ define the **distance** to be the dominant weight $\lambda = d(p, q)$ if (p, q) are in the same $G(\mathcal{K})$ orbit as (t^0, t^λ). Note that $d(q, p) = \lambda^*$.

2.2.2 Geometric Satake

The geometric Satake correspondence of Lusztig [25], Ginzburg [10], Beilinson–Drinfeld [2], and Mirković–Vilonen [26] gives an equivalence of tensor categories between the category of representations of the Langlands dual $G^\vee$ and the category of $G(\mathcal{O})$−equivariant perverse sheaves on the affine Grassmannian for G. In our case, $G = GL_m$ is its own Langlands dual.

Theorem 4 (Geometric Satake)**.** *The tensor category of equivariant perverse sheaves on the affine Grassmannian Gr with respect to the above stratification is equivalent to the tensor category of representations of G. Under this equivalence, for a dominant weight λ of G, $IH(\overline{Gr(\lambda)}) \cong V_\lambda$, where IH denotes intersection homology.*

For a sequence of dominant weights $\vec{\lambda}$ consider the twisted product of orbit closures,

$$\overline{Gr(\lambda^1)} \tilde{\times} \cdots \tilde{\times} \overline{Gr(\lambda^n)} := \left\{ (g_1, \ldots, g_n) \mid d(t^0, g_1) \leq \lambda^1, d(g_i, g_{i+1}) \leq \lambda^{i+1}, 1 \leq i \leq n-1 \right\}, \tag{2.5}$$

a sort of "path space" in Gr. Let $c_{\vec{\lambda}} = \dim(V_{\lambda^1} \otimes \cdots \otimes V_{\lambda^n})^G$ be **the Littlewood–Richardson coefficient**. Theorem 4 gives a way to construct the invariant space

$(V_{\lambda^1} \otimes \cdots \otimes V_{\lambda^n})^G$ geometrically from the following convolution morphism.

$$m_{\vec{\lambda}} : \overline{Gr(\lambda^1)} \tilde{\times} \overline{Gr(\lambda^2)} \tilde{\times} \cdots \tilde{\times} \overline{Gr(\lambda^n)} \to Gr$$
$$(g_1, \ldots, g_n) \mapsto g_n$$

The **polygon space** is the projective variety $\mathrm{Poly}(\vec{\lambda}) = m_{\vec{\lambda}}^{-1}(t^0)$, i.e. the path space "closed up". The geometric Satake correspondence yields the following theorem. See [18, Lem 4.6] for a proof.

Theorem 5. *The invariant space is isomorphic to the top homology of Poly($\vec{\lambda}$).*

$$(V_{\lambda^1} \otimes \cdots \otimes V_{\lambda^n})^G \cong H_{top}\left(Poly(\vec{\lambda})\right)$$

Haines showed in [14, Prop. 1.8] that for GL_m, $\mathrm{Poly}(\vec{\lambda})$ is equidimensional of dimension $\langle \lambda^1 + \cdots \lambda^n, \rho \rangle$, so as a corollary, the number of components of $\mathrm{Poly}(\vec{\lambda})$ is equal to $c_{\vec{\lambda}}$. In the next chapter, we restrict the λ^i to be minuscule, which is the only case we will care about.

Remark 6. Note that when the λ^i are minuscule, the only setting we will be interested in, $c_{\vec{\lambda}}$ is equal to a Kostka number: Let $\{\lambda^1, \ldots, \lambda^n\} = \{\omega_{i_1}, \ldots, \omega_{i_k}, \omega^*_{j_1}, \ldots, \omega^*_{j_l}\}$ be the minuscule weights and let $\Lambda_a = V_{\omega_a}, \Lambda^*_a = V_{\omega^*_a}$ be the corresponding irreducible representations. Then we must have $\sum i_a = \sum j_a$ for $c_{\vec{\lambda}}$ to be nonzero, and

$$\begin{aligned} c_{\vec{\lambda}} &= \dim\left(\Lambda_{i_1} \otimes \cdots \otimes \Lambda_{i_k} \otimes \Lambda^*_{j_1} \otimes \cdots \otimes \Lambda^*_{j_l}\right)^{GL_m} \\ &= \dim\left(\Lambda_{i_1} \otimes \cdots \otimes \Lambda_{i_k} \otimes \Lambda_{m-j_1} \otimes \cdots \otimes \Lambda_{m-j_l} \otimes V_{\det}^{*\otimes l}\right)^{GL_m} \\ &= \dim \mathrm{Hom}_{GL_m}\left(V_{\det}^{\otimes l}, \Lambda_{i_1} \otimes \cdots \otimes \Lambda_{i_k} \otimes \Lambda_{m-j_1} \otimes \cdots \otimes \Lambda_{m-j_l}\right) \end{aligned}$$

Thus, by the Pieri rule (and tranposing) $c_{\vec{\lambda}}$ is equal to the number of semistandard tableaux of shape $l \times m$ and content $(i_1, \ldots, i_k, m - j_1, \ldots, m - j_l)$.

CHAPTER 3

STATEMENT OF THE LOCAL RULE

In this chapter, we define the polygon space for a sequence of minuscule weights. Then we recall the Fontaine, Kamnitzer, and Kuperberg [8] method of indexing components of the polygon space with the generic distances along the fan triangulation. We then show that the components can be indexed using the generic distances along the edges of any fixed extroverted triangulation, thereby generalizing the result in [8]. Finally, we state the local rule satisfied by the generic distances, which allows us to figure out all of the generic distances of a component starting from those along any one extroverted triangulation.

3.1 The polygon space

We are interested in studying the combinatorics of the components of the polygon space, $\mathrm{Poly}(\vec{\lambda}) = m_{\vec{\lambda}}^{-1}(t^0)$, when the λ^i are minuscule. The components of $\mathrm{Poly}(\vec{\lambda})$ will be in bijection with the affine growth diagrams to be defined in Chapter 4. For λ dominant minuscule, there are no dominant weights less than λ, so $\overline{Gr(\lambda)} = Gr(\lambda)$ and the inequalities in 2.5 become equalities. The **polygon space for a sequence of minuscule weights $\vec{\lambda} = (\lambda^1, \ldots, \lambda^n)$** is the variety of n-gons,

$$\mathrm{Poly}(\vec{\lambda}) = \left\{ (g_1, \ldots, g_{n-1}) \in Gr^{n-1} \mid t^0 \xrightarrow{\lambda^1} g_1 \xrightarrow{\lambda^2} g_2 \xrightarrow{\lambda^3} \cdots \xrightarrow{\lambda^{n-1}} g_{n-1} \xrightarrow{\lambda^n} t^0 \right\}. \quad (3.1)$$

Recall that tensoring any irreducible representation V_μ with a minuscule representation V_λ is multiplicity free, that is, the irreducible representations arising in the direct sum decomposition of $V_\mu \otimes V_\lambda$ all occur with multiplicity 1. In the minuscule case, the Littlewood–Richardson coefficient is also equal to

the number of **minuscule paths for $\vec{\lambda}$**, namely sequences of dominant weights $\vec{\mu} = (\mu^0 = \emptyset, \mu^1, \ldots, \mu^n = \emptyset)$ such that $\mu^k - \mu^{k-1} = w \cdot \lambda^k$ for some element w in the Weyl group. In terms of representation theory this corresponds to first choosing the irreducible representation V_{μ^1} in the decomposition of $V_{\lambda^1} \otimes V_{\lambda^2}$, followed by V_{μ^2} in $V_{\mu^1} \otimes V_{\lambda^3}$ and so forth. At each step, λ^i being minuscule implies that the tensor product is multiplicity free and is given by the Pieri rule.

Since the Littlewood–Richardson coefficient is equal to both the number of components of Poly($\vec{\lambda}$) and the number of minuscule paths for $\vec{\lambda}$, there should be a bijection between the two, hence giving a way to index components with minuscule paths. Such a bijection is given in [8], which we recall in the next subsection. We then show that there are many different ways to index the components, one for each extroverted triangulation of the n-gon, and that there is a local rule that allows to transition from one indexing convention to another, which gives rise to the growth diagrams in Chapter 4.

3.1.1 Indexing components

In [8], Fontaine, Kamnitzer, and Kuperberg gave a way to index the components of Poly($\vec{\lambda}$) with minuscule paths. They define the fan configuration space

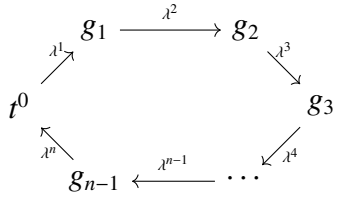

Figure 3.1: A point of Poly($\vec{\lambda}$)

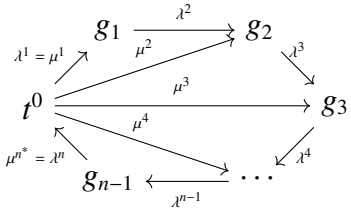

Figure 3.2: The fan triangulation

$Q(\vec{\lambda}, \vec{\mu})$, which imposes the additional conditions $d(t^0, g_i) = \mu^i$ on $\mathrm{Poly}(\vec{\lambda})$ (see Figure 3.2). Note that this forces $\mu^0 = \emptyset, \mu^1 = \lambda^1$, $\mu^{n-1} = \lambda_n^*$ and $\mu^n = \emptyset$. They proved the following theorem, which they noted is implicit in the work of Haines [14].

Theorem 7 ([8, Thm 4.1])**.** *Each component X of $\mathrm{Poly}(\vec{\lambda})$ has an open dense set X^o such that for all $0 \leq i \leq n$ the distances $d(t^0, g_i) = \mu^i$ are constant on X^o. The distances $\vec{\mu} = (\mu^0, \mu^1, \ldots, \mu^{n-1}, \mu^n)$ form a minuscule path for $\vec{\lambda}$, and this is a bijection between components and minuscule paths.*

Example 8. Let $n = 4$ and consider $\vec{\lambda} = (\omega_1, \omega_1, \omega_1^*, \omega_1^*)$ with the corresponding polygon space $\mathrm{Poly}(\vec{\lambda})$. There are two minuscule paths, $(\emptyset, (1), (2), (1), \emptyset)$ and $(\emptyset, (1), (11), (1), \emptyset)$, hence the polygon space consists of two irreducible components, one with generic distance $d(t^0, g_2) = (2)$, and the other with generic distance $d(t^0, g_2) = (11)$. Here by weight (11) we mean the dominant weight $(1, 1, 0, \ldots, 0)$. We represent the two components with the pictures below, where the corresponding minuscule paths can be read off along the edges of the fan triangulation.

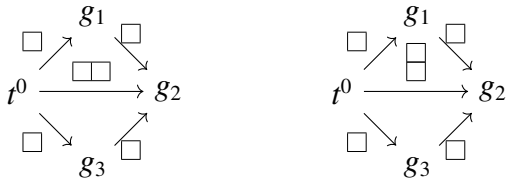

The proof of Theorem 7 relies on the following lemma. Let

$T_e(\mu, \lambda, \nu) =$ 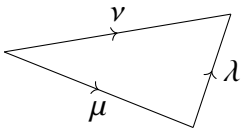

be a triangle with minuscule weight λ, based at the edge e of length μ. For two points $h_1, h_2 \in Gr$ such that $d(h_1, h_2) = \mu$, the edge-based configuration space $Q(T_e(\mu, \lambda, \nu))$ is $\{g \in Gr \mid d(h_2, g) = \lambda, d(h_1, g) = \nu\}$. For example, if $(h_1, h_2) = (t^{-\mu}, t^0)$, then $Q(T_e(\mu, \lambda, \nu))$ is a subvariety of $Gr(\lambda)$.

Lemma 9 ([8, Lem 4.2])**.** *For λ a minuscule weight, the edge-based configuration space $Q(T_e(\mu, \lambda, \nu))$ is non-empty if and only if there exists $w \in W$ such that $\mu + w\lambda = \nu$. If it is non-empty, then it is smooth and has complex dimension $\langle \nu - \mu + \lambda, \rho \rangle$.*

We reproduce the proof of Theorem 4.1 from [8] to see that it generalizes easily to triangulations other than the fan triangulation.

Proof of Theorem 7. The fan configuration space $Q(\vec{\lambda}, \vec{\mu})$ decomposes as a twisted product of triangle spaces as follows:

$$Q(\vec{\lambda}, \vec{\mu}) = Q(T_e(\mu^0, \lambda^1, \mu^1)) \tilde{\times} Q(T_e(\mu^1, \lambda^2, \mu^2)) \tilde{\times} \cdots \tilde{\times} Q(T_e(\mu^{n-1}, \lambda^n, \mu^n)).$$

Each factor contains a minuscule edge λ^i, so by the lemma, the ith factor is smooth of dimension $\langle \mu^i + \lambda^i - \mu^{i-1}, \rho \rangle$. Hence, $Q(\vec{\lambda}, \vec{\mu})$ is smooth irreducible of dimension the sum of the dimensions of the factors,

$$\dim Q(\vec{\lambda}, \vec{\mu}) = \sum_{i=1}^{n} \langle \mu^i + \lambda^i - \mu^{i-1}, \rho \rangle = \langle \lambda^1 + \cdots + \lambda^n, \rho \rangle = \dim Poly(\vec{\lambda})$$

Since $\mathrm{Poly}(\vec{\lambda})$ is partitioned into the equidimensional smooth subvarieties $Q(\vec{\lambda}, \vec{\mu})$, their closures must be the irreducible components of $\mathrm{Poly}(\vec{\lambda})$. □

The fan configuration space imposes conditions along the edges of the fan triangulation, but the reasoning of the proof can be applied to other triangulations. The key point of the proof is that each triangle contains an edge labeled by a minuscule weight. Hence, the same reasoning can be applied to any **extroverted triangulation**, a triangulation of the n-gon in which every triangle shares

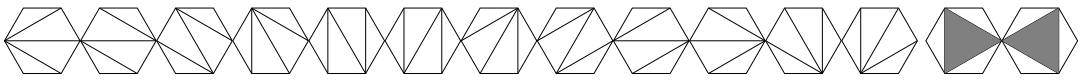

Figure 3.3: The 12 extroverted, and 2 non-extroverted, triangulations of the hexagon

an edge with the n-gon. See Figure 3.3 for the 12 extroverted triangulations of the hexagon.

To state the analogue of Theorem 7 for extroverted triangulation, we need some notation. Each extroverted triangulation contains two doubly extroverted triangles. Choose one of these triangles, say T, and let λ^i, λ^{i+1} be the clockwise labels of its two external edges in clockwise order. Set $\mu^1 = \lambda^i$ or $\mu^1 = \lambda^{i+1}$. Let μ^2 label the interior edge of T, oriented clockwise, that is, the tail and head coincide with the tail and head of the external edges of T oriented clockwise. Since each triangle of the extroverted triangulation contains an external edge, there is a triangle sharing an edge with T, having another internal edge, and one external edge. Its external edge is labeled by either λ^{i-1} or λ^{i+2}. It shares a single vertex with the internal edge labeled by μ^2. Label the new internal edge μ^3, oriented so that either the heads or tails coincide at the common vertex. Continue shelling on triangles and defining edge labels μ^i until a second doubly extroverted triangle completes the triangulation. This defines a bijection $\sigma : [n] \to [n]$ where $\lambda^{\sigma(i)}$ is the new external edge at step i, completing the triangle with edges μ^{i-1} and μ^i. For an extroverted triangulation τ, denote its interior edges by E_τ. The proof of the following theorem is a direct adaptation of the proof of Theorem 7.

Theorem 10. *Fix an extroverted triangulation τ. Each component X of Poly($\vec{\lambda}$) has an open dense set X^o_τ such that the distances along the edges of E_τ are constant on X^o_τ. This induces a bijection between the components of Poly($\vec{\lambda}$) and distances $\{\mu^i\}_{i\in[n]}$ satisfying $\mu^i - \mu^{i-1} = w \cdot \lambda^{\sigma(i)}$ for some $w \in W$ depending on i.*

Proof. The proof is identical to the one before, except for extroverted triangula-

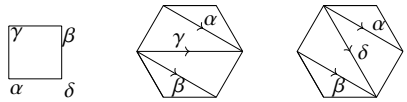

Figure 3.4: Dominant weights α, β, γ determine δ via one application of the local rule

tion τ with bijection σ, the decomposition is the following.

$$Q_\tau(\vec{\mu}, \vec{\lambda}) = Q(T_e(\mu^0, \lambda^{\sigma(1)}, \mu^1))\tilde{\times} \cdots \tilde{\times} Q(T_e(\mu^{n-1}, \lambda^{\sigma(n)}, \mu^n))$$

□

For a fixed component X, taking the intersection of the open dense sets X^o_τ over all of the extroverted triangulations τ gives an open dense set on which all the distances $d(g_i, g_j)$ are constant. Call these the **generic distances** for the component X.

Example 11. To continue Example 8, let $\vec{\lambda} = (\omega_1, \omega_1, \omega_1^*, \omega_1^*)$, where we saw that Poly($\vec{\lambda}$) has two irreducible components X_1 and X_2 corresponding to the minuscule paths $(\emptyset, (1), (2), (1), \emptyset)$ and $(\emptyset, (1), (11), (1), \emptyset)$ respectively. The square has only one other extroverted triangulation, and the generic value of $d(g_1, g_3)$ on X_1 is $\emptyset$, while the generic value of $d(g_1, g_3)$ on X_2 is $(1, 0, \ldots, 0, -1)$. Theorem 12 in the next subsection gives a way to calculate the generic distance $d(g_1, g_3)$ from the generic distance $d(t^0, g_2)$ of a component.

3.1.2 The local rule

We give a local rule that, given three generic distances, determines a fourth. Consider a quadrilateral with two external edges of the n-gon, such that the corresponding generic distances along the other two edges and one diagonal in

Poly($\vec{\lambda}$) are α, β, γ, as in Figure 3.4 where $n = 6$. We give a local rule that given these three generic distances, determines the fourth generic distance ρ, thereby performing a "quadrilateral flip."

Theorem 12 (Local Rule)**.** *For a generic point* $(g_1, \ldots, g_{n-1})$ *in a component of Poly*($\vec{\lambda}$)*, if the indicated distances of Figure 3.4 are given by* α, β, γ*, then* $\rho = sort(\alpha + \beta - \gamma)$*.*

Here dominant weights are added component-wise, that is $\alpha + \beta = (\alpha_1, \ldots, \alpha_n) + (\beta_1, \ldots, \beta_n) = (\alpha_1 + \beta_1, \ldots, \alpha_n + \beta_n)$. For three dominant weights α, β, γ, the expression $\alpha + \beta - \gamma$ need not be a weakly decreasing sequence, hence the sort operation orders the sequence of integers to make it a dominant weight. We leave the proof to Chapter 5.

Example 13. Continuing with Examples 8 and 11, recall that the two components X_1 and X_2 correspond to the minuscule paths $(\emptyset, (1), (2), (1), \emptyset)$ and $(\emptyset, (1), (11), (1), \emptyset)$, respectively. We can calculate the generic distance $d(g_1, g_3)$ on X_1 by calculating sort$((1) + (1) - (2)) = \emptyset$. For X_2, we get sort$((1) + (1) - (11)) =$ sort$(1, -1, 0, \ldots, 0) = (1, 0, \ldots, 0, -1)$.

Remark 14. As mentioned in the introduction, this rule previously appeared in [39] wherein van Leeuwen realized combinatorially the symmetry of the tensor product of irreducible representations in terms of Littelmann paths. For a reductive group G consider $V_\lambda \otimes V_\mu = \sum c^\nu_{\lambda,\mu} V_\nu$ and let $\pi : [0, 1] \to X \otimes_{\mathbb{Z}} \mathbb{Q}$ be a fixed piecewise-linear path in the space of rational weights such that $\pi(1) = \mu$ and the image of π is in the dominant Weyl chamber. For a simple root α, let e_α and f_α be the Littelmann root operators that act on such paths, and let B_π be the set of paths reachable from π by repeated applications of e_α and f_α. Then recall from [23] that $c^\nu_{\lambda,\mu}$ is equal to the number of paths in B_π that are λ-dominant and end at $\nu - \lambda$ (equivalently, these are paths in B_π that when translated by λ are dominant

and end at ν). One of Littelmann's results states that this does not depend on the choice of π. By the symmetry of the tensor product decomposition, the multiplicity $c^{\nu}_{\lambda,\mu} = c^{\nu}_{\mu,\lambda}$ can also be computed as the number of paths in B_σ that are μ-dominant and end at $\nu - \mu$ for a fixed dominant path σ such that $\sigma(1) = \lambda$. In [39], van Leeuwen writes down a bijection between such paths when the initial paths π and σ come from a natural subset of paths that are closed under e_α and f_α.

For example, for SL_m it is natural to consider not all piecewise-linear paths, but only those that are a concatenation of segments $\pi_{\epsilon_i}(t) = t\epsilon_i$ where $t \in [0, 1]$ and ϵ_i is the weight that returns the (i, i) diagonal entry ($i \in \{1, \ldots, m\}$). Such ϵ-paths are represented by semistandard Young tableaux where the column reading word gives the sequence of indices i for the segments ϵ_i in the path. Then the set of semistandard tableaux of shape μ is in bijection with $B_{\pi_c(\mu)}$ where $\pi_c(\mu)$ is the path given by the tableaux that contains entry i in each box of row i. The character of the irreducible representation V_μ is then $\sum_{p \in B_{\pi_c(\mu)}} e^{p(1)}$, as desired. Under this identification of paths with tableaux, van Leeuwen's bijection is between Littlewood-Richardson tableaux of skew-shape ν/λ with content μ and those of skew-shape ν/μ and content λ, and is realized by jeu de taquin. Jeu de taquin can be computed via a rectangular growth diagram with a local rule given in [39, §3] and [38] (see also the appendix of [34]). Note that this is the same rule given in [33].

For other groups, the analogous weights ϵ_i do not $\mathbb{N}$-span the weight lattice, hence the subset of paths considered must be enlarged to minuscule paths ([39, §4]), those paths that are a concatenation of segments that are straight paths to a (not necessarily dominant) minuscule weight. (In fact, this is not enough for

types G_2, F_4, E_8 and quasi-minuscule weights must be considered to cover these cases. A local growth rule is given in [39, §4] for this case as well.) For example, we can still consider SL_m, but now the segments are allowed to be straight line paths $\pi_\lambda(t) = t\lambda$ where

$$\lambda \in \{w \cdot \omega_j \mid \omega_j \text{ is a fundamental weight and } w \in W\}.$$

In the GL_m case, straight-line segments to weights in the orbits of ω_j and ω_j^* are allowed.

In this setting, $c^\nu_{\lambda,\mu} = c^\nu_{\mu,\lambda}$ is then both the number of λ-dominant minuscule paths to $\nu - \lambda$ in $B_{\pi(\mu)}$ ($\pi(\mu)$ some fixed minuscule path to μ) and the number of μ-dominant minuscule paths to $\nu - \mu$ in $B_{\pi(\lambda)}$ ($\pi(\lambda)$ some fixed minuscule path to λ). Viewing minuscule paths as a sequence of dominant weights van Leeuwen realized this bijection via a rectangular growth diagram with the same local rule as in Theorem 12. See also [27] and [42].

CHAPTER 4

GROWTH DIAGRAMS

In Section 4.1, we define affine growth diagrams, show that they are in bijection with the components of the polygon space in Theorem 17, and deduce that they are periodic in Corollary 19. In Section 4.2, we show that when all of the λ^i are ω_1 or ω_1^*, there is a combinatorial way to mark unit squares in the growth diagrams, which associates a fixed-point-free involution to an affine growth diagram. This is Proposition 25. In Section 4.3, we recover the Stanley–Sundaram bijection and Fomin growth diagrams. Section 4.4 deals with the Knuth versions, i.e. when λ^i are not restricted to ω_1 and ω_1^*.

4.1 Defining affine growth diagrams

As with Fomin growth diagrams, the local rule of Theorem 12 allows us to construct growth diagrams where α, β, γ label three vertices of a unit square and determine the fourth, southeast vertex label ρ, as indicated in Figure 3.4. Fix a sequence of minuscule weights $\vec{\lambda}$ and a minuscule path $\vec{\mu}$ for $\vec{\lambda}$. By Theorem 7, the μ^i are the generic distances $d(t^0, g_i) = \mu^i$ of a corresponding component X of $\mathrm{Poly}(\vec{\lambda})$. Repeated application of the local rule allows us to determine the distance along each interior edge, thereby determining all of the generic distances of X. By Theorem 10, we could also begin with the weights $\vec{\mu}$ along E_τ for any extroverted triangulation τ, not necessarily the fan triangulation.

Organize the information of the generic distances of a fixed component X of $\mathrm{Poly}(\vec{\lambda})$ into a diagram in the following manner. Consider an n-wide infinite staircase, such that each row contains $n-1$ unit squares and row i has a square

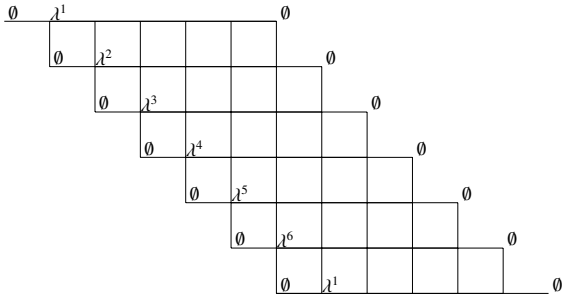

Figure 4.1: Empty affine growth diagram of type $\vec{\lambda} = (\lambda^1, \ldots, \lambda^6)$

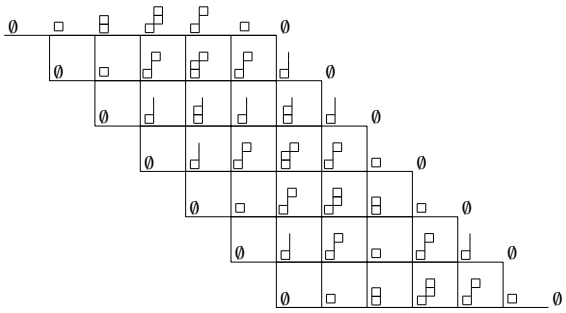

Figure 4.2: An affine growth diagram of type $\vec{\lambda} = (\omega_1, \omega_1, \omega_1^*, \omega_1^*, \omega_1, \omega_1^*)$

in columns $i+1, \ldots, i+n-1$. We call this the **standard position** of the staircase diagram. Let horizontal line i be the horizontal line segment below the ith row of unit squares. Similarly, let vertical line i be the vertical line segment to the right of the ith column of unit squares. Label the vertex at horizontal line i and vertical line j by the generic distance $d(g_i, g_j)$ of X, where indices are taken mod n. Recall a point in X is of the form $(g_0 = t^0, g_1, \ldots, g_{n-1}, g_n = t^0)$. Note that $d(g_i, g_i) = \emptyset$ and $d(g_i, g_{i+n}) = \emptyset$, so the diagonals are labeled by the zero weight.

The diagram is organized so that a path through the staircase from the left to the right diagonal, together with the weight labels along the path, corresponds to an extroverted triangulation τ, together with the generic distances of X along the edges E_τ. Taking a step through the staircase diagram arrives at a vertex labeled by the generic distance of X along the corresponding edge in the n-gon.

Oriented extroverted triangulations correspond to paths through the staircase diagram from the left diagonal to the right diagonal as follows. An oriented extroverted triangulation can be specified by first choosing an external edge i and orienting it clockwise (R) or counterclockwise (L). View the n-gon with edge i at the top. Then attach a doubly extroverted triangle to edge i ei-

ther to the right (R) or to the left (L), orienting the internal edge of this triangle towards the new external edge (note that this differs from the orientation in subsection 3.1.1 where internal edges were oriented to match edge i's orientation). Continue shelling on triangles to the right (R) or to the left (L), each time orienting the internal edge towards the new external edge, until a second doubly extroverted triangle is reached, completing the triangulation. In terms of paths through the staircase diagram, an R corresponds to a horizontal step, while an L corresponds to a vertical step. For example, the fan triangulation for $n = 6$ is encoded as (1R)(RRRR), and corresponds to the 0th horizontal line of the staircase diagram. The first external edge in the triangulation determines the row or column of unit squares across which the first step is taken in the staircase diagram (see Figure 4.3).

Under this correspondence, the vertices of the ith horizontal line of the staircase are labeled by the weights along the edges of the fan triangulation rotated i times. Proposition 15 will show that this labeling is consistent for every triangulation under the path-to-triangulation correspondence, keeping in mind the dual convention following vertical steps. Before proving it, we make some observations. When discussing vertices of a fixed horizontal line, index the vertices from 0 to n. Since the first edge (iR) of a fan triangulation is the external edge i read clockwise, vertex 1 of line $i - 1$ must be labeled by $\lambda^i = d(g_{i-1}, g_i)$ (see Figure 4.1). On the other hand, any oriented triangulation that begins with (iL) has weight λ^{i^*} along its first edge. Since (iR) and (iL) correspond to, respectively, a horizontal and a vertical step ending at the same vertex in the staircase diagram, this shows that it is necessary to take the dual of a weight following a vertical step. Also, notice that the last edge of ($i + 1$R)(R$\cdots$R) in the n-gon is the first edge of (iR)(R$\cdots$R) with the opposite orientation, so the $n - 1$ vertex of

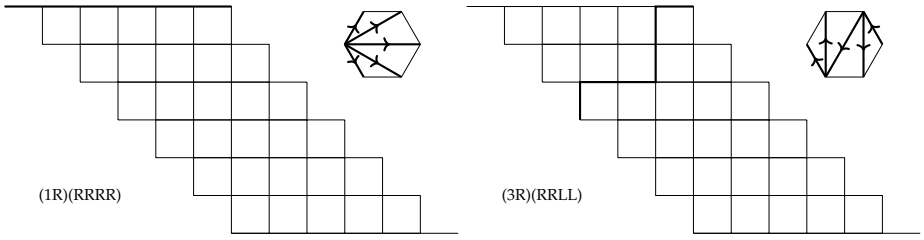

Figure 4.3: Extroverted triangulation to paths through staircase correspondence

horizontal line i is labeled by λ^{i^*} where the index i is taken mod n.

Proposition 15. *The above labeling of the vertices of a staircase diagram by the generic distances of a component X of Poly($\vec{\lambda}$) is consistent under the path-to-triangulation correspondence.*

Proof. Number the vertices and edges of the n-gon such that the ith edge connects vertex $i-1$ to vertex i clockwise. Let P be a path in the staircase from a SW vertex to a NE vertex, and suppose that the first step of the path ends at a vertex in the ith horizontal line. We have already mentioned in the previous paragraph that the first vertex label is λ^{i+1}, consistent with either a horizontal (i+1R) or vertical (i+1L) step. Choose a vertex x of P and suppose that there have been h horizontal steps and v vertical steps taken, not counting the first step. Then the current vertex is in horizontal line $i-v$ and vertical line $i+1+h$, hence labeled by $d(g_{i-v}, g_{i+1+h})$. But the h many horizontal steps correspond to h many R triangle shellings in the n-gon, arriving at vertex $i+1+h$, and the v many vertical steps correspond to v many L triangle shellings, arriving at vertex $i-v$. If the final step was a horizontal step, then the edge in the n-gon is oriented from vertex $i-v$ to vertex $i+1+h$, as desired, and vice versa for a vertical step. □

The previous proposition shows that the generic distances of X can be

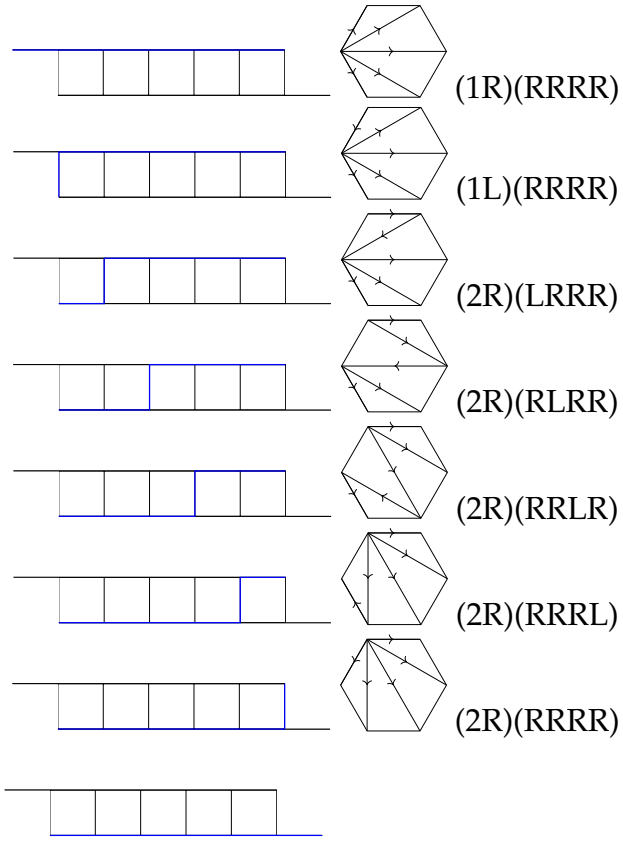

Figure 4.4: Rotation

recorded in a single diagram in a consistent way. However, knowing the minuscule path $\vec{\mu}$ corresponding to X from Theorem 7 allows us to build such a diagram using the local rule, thereby discovering all of the generic weights. Start by labeling the vertices of the first horizontal line of an empty staircase with $\vec{\mu}$. Apply the local rule repeatedly to discover the vertex labels of the second horizontal line (see Figures 4.4, 4.5, 4.6, 4.7).

Note that the local rule can be written as $\delta = \text{sort}(\alpha - (\gamma - \beta))$, which is interpreted as follows. By the Pieri rule, boxes were either only added to distinct rows of β to get γ, or boxes were only removed in distinct rows. Supposing that boxes were only added, record the distinct rows that increased. Then decrementing the same rows of α and sorting gives δ. Vice versa if boxes were origi-

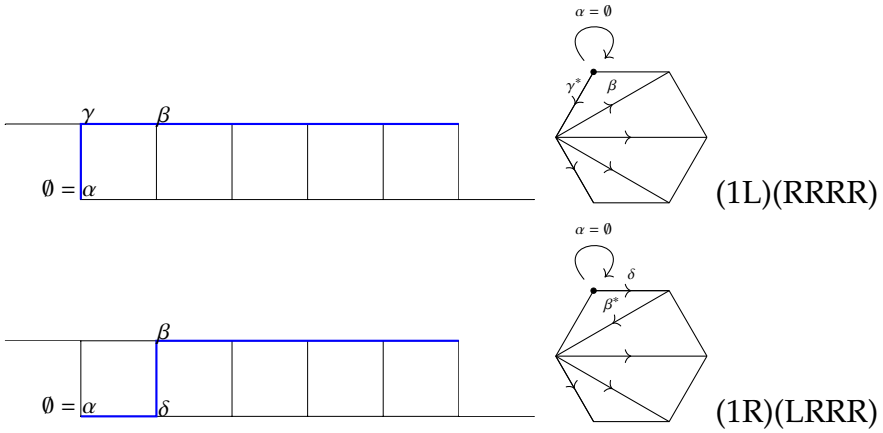

Figure 4.5: Interpreting first LR to RL switch as a tetrahedron move

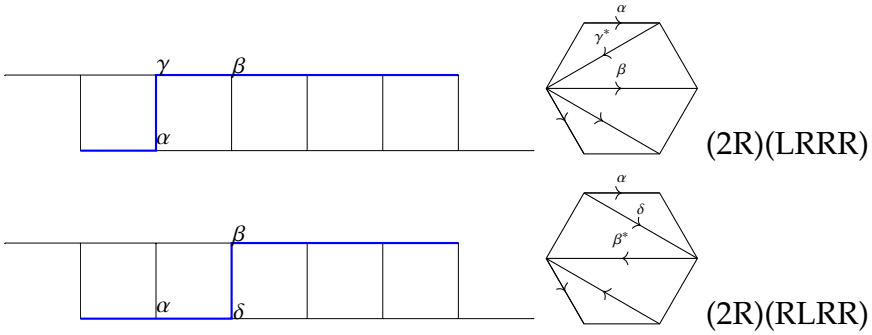

Figure 4.6: Tetrahedron move

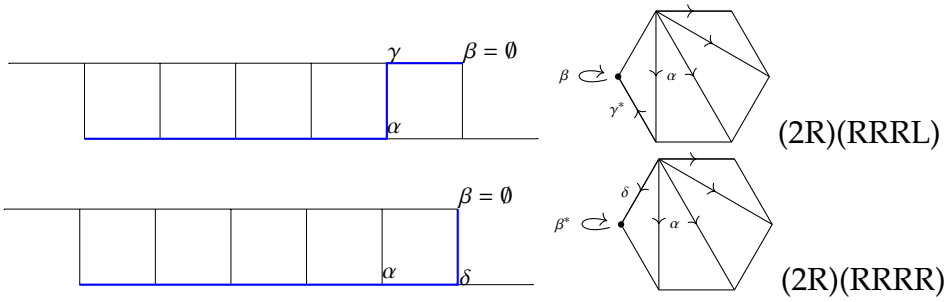

Figure 4.7: Interpreting last LR to RL switch as a tetrahedron move

nally removed from β to give γ. In particular, if $\gamma - \beta = w \cdot \omega_i$, then $\delta - \alpha = w' \cdot \omega_i$, hence the difference propagates down columns and rows. Consider the column of unit squares with μ^{i-1}, μ^i labeling the vertices across the top, so that $\mu^i - \mu^{i-1} = w \cdot \lambda^i$. Since this column eventually reaches a unit square with SW vertex a boundary vertex labeled by $\emptyset$, it's SE vertex has to be a dominant weight θ such that $\theta - \emptyset = w \cdot \lambda^i$, i.e. $\theta = \lambda^i$. This motivates the following definition (see Figures 4.9 and 4.18 for examples).

Definition 16. Start with an empty staircase with the left and right diagonal vertices labeled by the zero weight, $\emptyset$. Label the super-diagonal with $\vec{\lambda}$ periodically (see Figure 4.1 for an $n = 6$ example). An **affine growth diagram of type $\vec{\lambda}$** is a labeling of the remaining vertices by dominant weights such that the local rule is satisfied by every unit square.

Theorem 17. *The components of Poly($\vec{\lambda}$) are in bijection with affine growth diagrams of type $\vec{\lambda}$.*

Proof. The components of Poly($\vec{\lambda}$) are in bijection with minuscule paths for $\vec{\lambda}$, so we must show that affine growth diagrams of type $\vec{\lambda}$ are in bijection with minuscule paths $\vec{\mu}$ for $\vec{\lambda}$. As we just observed, labeling the top line of a staircase with a minuscule path $\vec{\mu}$ gives deterministically an affine growth diagram of type $\vec{\lambda}$ via the local rule. For the opposite direction we must show that any affine growth diagram of type $\vec{\lambda}$ has first line labeled by a minuscule path $\vec{\mu}$.

Consider an affine growth diagram of type $\vec{\lambda}$ with the 0th horizontal line labeled by $\mu^0, \ldots, \mu^n$. The difference across the first edge of horizontal line $i - 1$ is $\lambda^i - \emptyset = \lambda^i$. Since every unit square satisfies the local rule, the difference in line $i - 2$ across the edge in the same column is $w \cdot \lambda^i$ for some $w \in W$ (we have applied the local rule from southwest to northeast). This difference propagates

up column i (possibly with different w for each line), until it reaches the 0th horizontal line where we see that $\mu^i - \mu^{i-1} = w \cdot \lambda^i$ for some $w \in W$. □

Example 18. Let $\vec{\lambda} = (\omega_1, \omega_1, \omega_1, \omega_1^*, \omega_1^*, \omega_1^*)$. There are six minuscule paths corresponding to six components of Poly($\vec{\lambda}$). The corresponding six affine growth diagrams appear in Figure 4.15.

The proof of Theorem 17 shows that for any affine growth diagram of type $\vec{\lambda}$, the line with first vertex label λ^1 is labeled by a minuscule path for $\vec{\lambda}$. Thus, there is a component of Poly($\vec{\lambda}$) with these generic distances along the fan triangulation. Since rotating the fan triangulation n times gives back what one started with, the affine growth diagram must be n-periodic. This establishes the following combinatorially non-obvious fact. See also [15] and [16] for a discussion of periodicity of the bounded octahedron recurrence.

Corollary 19. *Affine growth diagrams of type* $(\lambda^1, \ldots, \lambda^n)$ *are n-periodic.*

Remark 20. For neighboring fan triangulations τ and τ' (as in Figure 4.8 where $n = 6$), Fontaine and Kamnitzer gave a combinatorial rule in [7] for determining $E_{\tau'}$ from E_τ, i.e. a rotation of a minuscule path. The local rule in Theorem 12 refines this rotation into smaller steps, each transforming an extroverted triangulation by a single quadrilateral flip. Hence, the rotation of Figure 4.8 is broken up into the three individual steps between the first four triangulations of Figure 3.3.

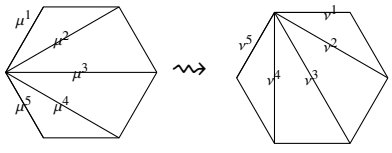

Figure 4.8: Rotation of minuscule path $\vec{\mu}$

Notice that (1R)(R$\cdots$R) and (nL)(L$\cdots$L) encode the same exact oriented triangulation, but the first edge of (1R)(R$\cdots$R) is the last edge of (nL)(L$\cdots$L). Since the nth vertical line is labeled by the weights of (nL)(L$\cdots$L), reading from top to bottom along this vertical line gives the dual of the weights when reading the 0th horizontal line from left to right. The same holds for horizontal line i and vertical line $i+n$, so affine growth diagrams have a dual symmetry:

Proposition 21. *Transposing an affine growth diagram and relabeling the resulting column i as $i+n$ in order to put into standard position gives the dual diagram.*

Proof. Taking the transpose results in a staircase diagram where row i has unit squares in columns $i-n+1,\ldots,i-1$. So adding n to the column indices gives an affine growth diagram in standard position. The resulting ith horizontal line was originally the $i+n$ vertical line, which are dual by the above discussion. □

The proposition says that transposing the left triangular region of Figure 4.9 and taking duals of the weights gives the right triangular region, so knowing the weights in one of these triangles determines the rest of the affine growth diagram. We call the left triangular region the **fundamental region**, and the right region the **dual fundamental region**. The markings in Figure 4.9 are explained in Section 4.2. We end this section with a remark.

Remark 22. In the SL_m setting, the local growth rule is the Bender–Knuth rule. Consider the weight lattice of SL_m rather than GL_m, which is $\mathbb{Z}^n/(1,\ldots,1)$. In this case, $\omega_i^* = \omega_{m-i}$, so the minuscule weights are $\omega_1,\ldots,\omega_{m-1}$. A minuscule path is equivalent to a semistandard Young tableau of rectangular shape, where here semistandard means strictly increasing along rows and weakly increasing down columns. For example, $\emptyset$, 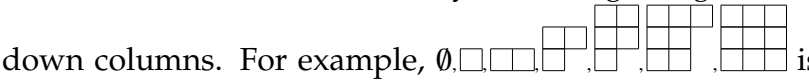 is a minuscule path of

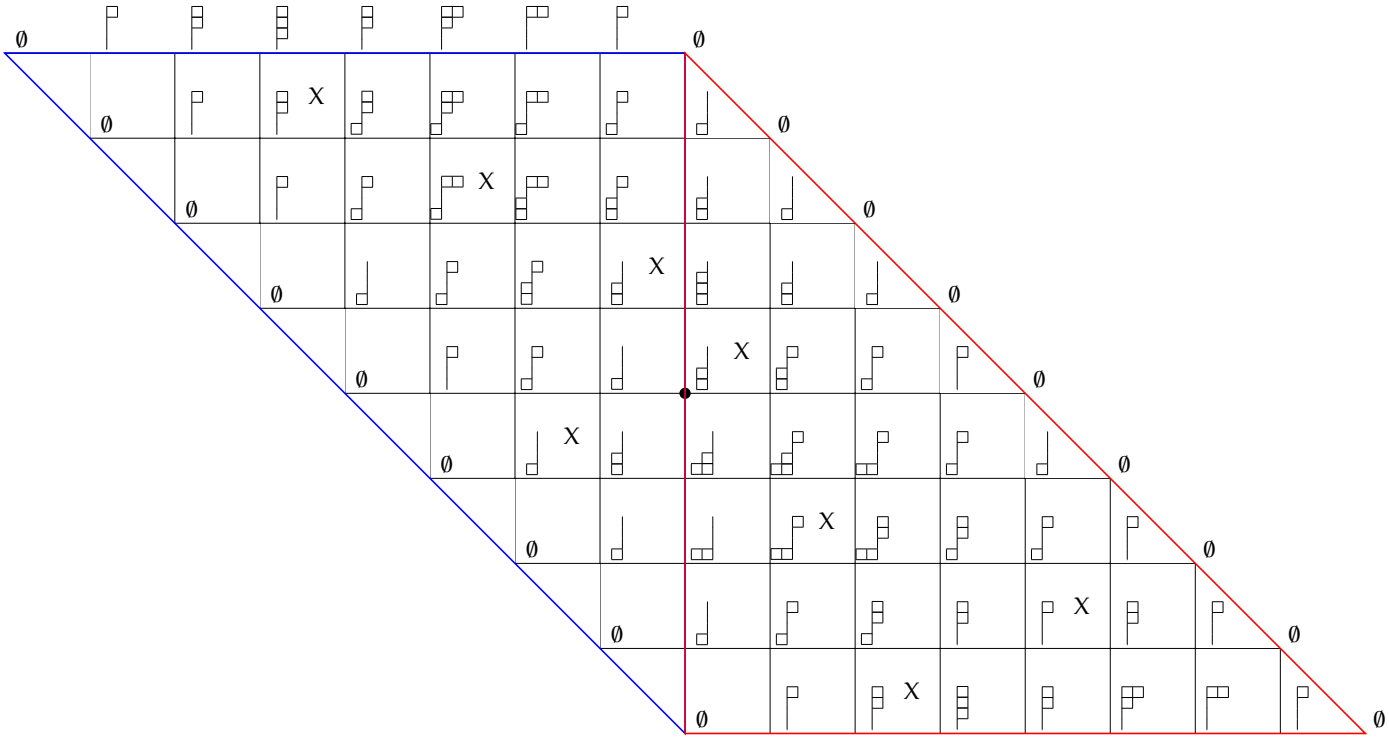

Figure 4.9: Affine growth diagram for $\vec{\lambda} = \left(\omega_1, \omega_1, \omega_1, \omega_1^*, \omega_1, \omega_1^*, \omega_1^*, \omega_1^*\right)$

type $(\omega_1, \omega_1, \omega_1, \omega_2, \omega_2, \omega_2)$, corresponding to tableau $\begin{array}{|c|c|c|}\hline 1&2&5\\\hline 3&4&6\\\hline 4&5&6\\\hline\end{array}$.

Recall the Bender–Knuth involution, σ_k, which acts on tableaux by exchanging the number of k's and $k+1$'s as follows. If an entry containing a k has an entry containing a $k+1$ directly to the right of it, these pair of entries are considered "locked." The remaining entries containing k or $k+1$ are considered free. In each column, these free entries are consecutive, and if column i contains a free entries with k and b free entries with $k+1$, edit column i so that it contains b free entries with k and a free entries with $k+1$, while preserving the semistandard

condition. Here is an example.

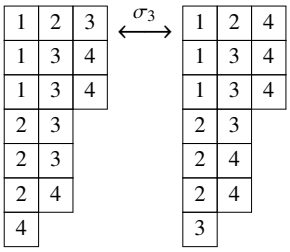

Bender–Knuth involutions show that the number of semistandard tableaux of a fixed shape λ and content remains the same when the content is permuted, hence the Schur function associated to λ is indeed a symmetric function.

The Bender–Knuth rule can be used as the local rule in a growth diagram as follows. Let T be a tableau and let $\lambda_i(T)$ be the shape of the portion of T containing boxes with entries from $\{1, \ldots, i\}$. Label the southwest corner of a unit square by $\lambda_{k-1}(T)$, the northwest corner by $\lambda_k(T)$, and the northeast corner by $\lambda_{k+1}(T)$. Define the southeast corner to be $\lambda_k(\sigma_k(T))$. The claim is that the Bender–Knuth rule and the local growth rule of Theorem 12 agree.

Proposition 23. *With the above notation,* $\lambda_k(\sigma_k(T)) = \mathit{sort}\,(\lambda_{k+1}(T) + \lambda_{k-1}(T) - \lambda_k(T))$.

Proof. Suppose the rule from Theorem 12 gives partition

$$\mu = \mathrm{sort}\,[\lambda_{k+1}(T) - (\lambda_k(T) - \lambda_{k-1}(T))]\,,$$

rewritten to emphasize that the set of distinct rows containing a k entry in T is $\lambda_k(T) - \lambda_{k-1}(T)$. The rule says to subtract a box from $\lambda_{k+1}(T)$ in each of those rows to get the new southeast partition. If a row contained a "locked" pair, then the length of that row in μ is the same as the length of the same row in $\lambda_k(T)$, which agrees with Bender–Knuth. If not, then the free k entries are deleted, and the sorting rule slides the free $k+1$ entries to get a valid partition μ such that $\lambda_k(\sigma_k(T)) = \mu$. □

It is known that the Bender–Knuth rule computes promotion of tableaux. See for example [4]. Fontaine and Kamnitzer already observed in [7] that their rotation of minuscule paths in the SL_m case is precisely promotion of rectangular tableaux. For SL_m the dominant minuscule weights are the fundamental weights $\omega_1, \ldots, \omega_{m-1}$, so a growth diagram of type $(\omega_{i_1}, \ldots, \omega_{i_n})$ for SL_m has $\emptyset$ along the western diagonal and the rectangular partition with m rows and $\sum i_j$ many boxes along the eastern diagonal. Hence, the horizontal lines are labeled by chains of partitions, which can be interpreted as row strict rectangular tableaux.

Let ∂ denote the promotion operation on tableaux computed via deflation, and ∂^* the dual promotion operation computed via inflation, so that $\partial \circ \partial^* = \mathrm{id} = \partial^* \circ \partial$. Suppose that tableau T labels line i of the affine growth diagram. Then line $i+1$ is labeled by $\partial(T)$ and line $i-1$ by $\partial^*(T)$. The n-periodicity of the growth diagram recovers the n-periodicity of (dual) promotion on rectangular tableaux, $\partial^n(T) = T = (\partial^*)^n(T)$, where the tableau is filled with numbers from $\{1, \ldots, n\}$. Likewise, traversing successive vertical lines to the left is given by promotion and to the right by dual promotion. Furthermore, the vertical line that begins where horizontal line i ends is labeled by the Schützenberger evacuation tableau $ev(T)$. The $n + i$ horizontal line is labeled by the dual evacuation procedure applied to $ev(T)$, which recovers the identity $ev^* \circ ev(T) = \partial^n(T)$, and the fact that $ev^* = ev$ in the rectangular case. We can also recover the identity $ev \circ \partial = \partial^* \circ ev$, as indicated in Figure 4.10.

Note that Speyer's diagrams appearing in [33] are the SL_m version with all λ^i equal to the first fundamental weight ω_1. In this special case, the sorting local rule (equivalently Bender–Knuth) specializes to the growth rule given in

[33], which is the jeu de taquin rule in [39, §3] and [34, Prop A.1.2.7]. Speyer's diagrams start with $\emptyset$ down the western diagonal and end with a rectangular partition $(n-d,\ldots,n-d)$ down the eastern diagonal for parameters d and n. It would be interesting to better understand the connection with the geometry arising in [33]. See also [27] and [42] for a connection to coboundary categories and crystals.

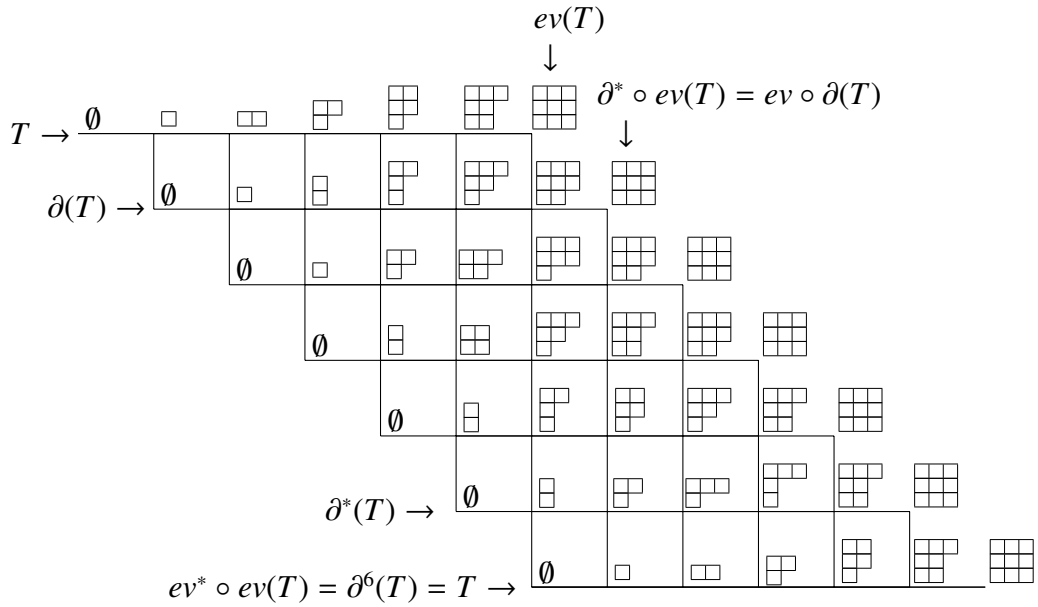

Figure 4.10: An SL_3 diagram labeled with the effects of promotion and evacuation

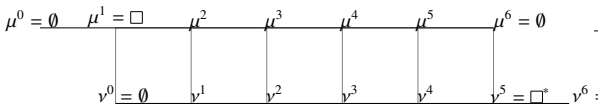

Figure 4.11: First row of a diagram

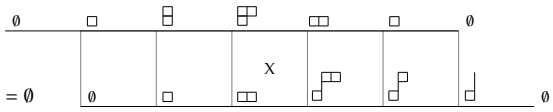

Figure 4.12: Marking the first row

4.2 Marking affine growth diagrams

In this section, suppose that all λ^i are either ω_1 or ω_1^*. In this case, there is a natural way to mark affine growth diagrams. Consider again starting with a minuscule path, the top line of an affine growth diagram, and determining the second line from repeated applications of the local rule (see Figure 4.11). Suppose that $\lambda^1 = \omega_1$, so that $\mu^1 = \omega_1$ as well. To compute ν^1, the local rule says to subtract one from the first row of μ^2 and sort to get ν^1. This can happen in two ways. If μ^2 had a box in the first row, then ν^1 has one fewer positive box than μ^2. If μ^2 did not have a box in the first row, then ν^1 has one more negative box than μ^2. In other words, going from μ^2 to ν^1 either removes a positive box, or adds a negative box. This relation between μ^i and ν^{i-1} continues to hold for all i. Since $\mu^1 = (1, 0, \ldots, 0)$, $\nu^0 = \emptyset$ and $\mu^6 = \emptyset$, $\nu^5 = (0, \ldots, 0, -1)$, the differences between μ^i and ν^{i-1} begin by being a removal of a positive box and finish by being an addition of a negative box, with a switch occurring somewhere in the middle. This switch is guaranteed to happen exactly once if $m \geq n/2$. Mark the unit square across which this switch occurs (see Figure 4.12). Similar reasoning applies if $\lambda^1 = \omega_1^*$.

Proposition 24. *Let each λ^i be ω_1 or ω_1^* and assume $m \geq n/2$. Then any affine growth diagram of type $\vec{\lambda}$ has exactly one mark in every row and every column, hence gives an affine permutation.*

Proof. Suppose $\mu^1 = \omega_1$. Then the local rule for each unit square implies that $\mu^i - \nu^{i-1} = w \cdot \omega_1$ for some $w \in W$, or in other words $|\mu^i| = |\nu^{i-1}| + 1$. Each dominant weight can be seen as a pair of partitions, the positive part α, and the negative part β. Hence, either the positive partition α decreased from μ^i to ν^{i-1}, or the negative partition β increased. Going from μ^1 to $\emptyset$ decreases α, and observing

the effects of traversing the subsequent vertical lines we see that for a while α decreases, then β increases. Once the switch occurs, β continues to increase along all subsequent vertical lines as long at $m \geq n/2$. Furthermore, a switch must occur at some point because the last vertical line connects $\emptyset$ to $\mu^{1^*} = \omega_1^*$.

Suppose we marked the diagram by the analogous rules for columns. By the same reasoning as for rows, a switch must occur exactly once as we proceed along a column. We claim that this switch for a column coincides with the switch for a row.

Suppose that not all switches coincide. Then without loss of generality there is a square where there is a row switch, but not a column switch. When traversing the north and south edges, either both traversals alter β or both alter α. On the other hand, exactly one of the west or east edges alters α and the other β. In either case, traversing first the north edge then the east edge cannot give the same dominant weight as first traversing the west edge and then the south edge, a contradiction. Thus the marks coincide whether filling the diagram along rows or columns. □

Let $\tilde{S_n}$ be the group of all bijections $f : \mathbb{Z} \to \mathbb{Z}$ such that $f(x + n) = f(x) + n$. These are called **affine permutations**. The previous proposition says that the marked squares of an affine growth diagram give an affine permutation. Such an f is uniquely determined by its values on $[n]$, written as $[f(1), \cdots, f(n)]$, and called the **window** of f. The example in Figure 4.9 gives the affine permutation $[4, 6, 8, 9, 7, 10, 13, 11]$. From the affine permutation we can get a classical permutation by reducing mod n. For example, reducing $[4, 6, 8, 9, 7, 10, 13, 11]$ mod 8 gives the permutation 46817253. This can be seen as taking the dual fundamental region of Figure 4.9 and placing it under the fundamental region to complete

an $n \times n$ square. The marked squares then give a permutation in the usual sense.

Recall that transposing a growth diagram gives the dual diagram (Proposition 21). Let shift_n denote the affine permutation $x \mapsto x + n$.

Proposition 25. *The affine permutation of an affine growth diagram composed with itself gives* shift_n*. The classical permutation associated to an affine growth diagram is a fixed-point-free involution.*

Proof. Let f denote the affine permutation. Marking the dual of an affine growth diagram gives the same marked squares. Since the dual growth diagram is the transpose shifted by n, it follows that $f \circ f(x) = x + n$. Reducing the window of $f \mod n$ gives a fixed-point-free involution. □

Let the ith **hook** be the portion of the ith column and ith row contained in the fundamental region. An equivalent criterion to the classical permutation being a fixed-point-free involution is that each hook has a marked square in either its horizontal portion or vertical portion, but not both. In the next section, we will show that there is a bijection between fixed-point-free involutions in S_n (n even) and affine growth diagrams such that the labels of line 0 only have positive parts.

4.3 Classical bijections and Fomin growth diagrams

Previously, we saw that there is a natural way to associate a fixed-point-free involution to an affine growth diagram. In this section, we define an inverse that takes a fixed-point-free involution written into an empty staircase diagram and

produces an affine growth diagram with 0th line an oscillating tableau. From this, we recover the Stanley–Sundaram bijection between oscillating tableaux and fixed-point-free involutions. We then show that Fomin growth diagrams appear within affine growth diagrams when $\vec{\lambda} = (\omega_1, \ldots, \omega_1, \omega_1^*, \ldots, \omega_1^*)$.

Definition 26. An **oscillating tableau** of length k is a sequence of partitions $(\mu^0 = \emptyset, \mu^1, \mu^2, \ldots, \mu^k = \emptyset)$ such that for all i, μ^i and μ^{i+1} differ by exactly one box.

Proposition 27. *For n even, there is a bijection between fixed-point-free involutions in S_n and length n oscillating tableaux.*

Proof. Use Fomin growth diagrams on a triangular region as in [28, Chp. 4]. Or, see [37, p. 60] for a proof attributed to Stanley. □

For $\lambda^1, \ldots, \lambda^n$ all equal to ω_1 or ω_1^*, Proposition 24 gives a map from affine growth diagrams to fixed-point-free involutions in S_n. Notice that a minuscule path consisting of weights with only nonnegative entries is precisely an oscillating tableau, so consider restricting this map to affine growth diagrams with the 0th line labeled by an oscillating tableau, and denote this map by ϕ (see the left diagram of Figure 4.18). We would like to define an inverse for ϕ, hence reproving Proposition 27 using affine growth diagrams. Furthermore, during the course of the proof, we will see that Fomin growth diagrams are contained within affine growth diagrams when $\vec{\lambda} = (\omega_1, \ldots, \omega_1, \omega_1^*, \ldots, \omega_1^*)$.

We will now define a map from fixed-point-free involutions to affine growth diagrams with 0th line an oscillating tableau. Let π be a fixed-point-free involution of $[n]$. For rows $1, \ldots, n$ fill an empty staircase diagram such that row i has a marked square in $\pi(i)$ if $i < \pi(i)$ and otherwise in $\pi(i) + n$. This coincides

with row i of the diagram having squares in columns $i+1, \ldots, i+n$. Extend the marking periodically, so that if square (i, j) is marked then so is $(i+n, j+n)$.

Given this marked staircase diagram we need to specify a dominant weight at each vertex such that the weights of the 0th line consist of nonnegative parts. Assume $m \geq n/2$, so that specifying a dominant weight is equivalent to giving the positive partition α and the negative partition β. Choose a vertex and extend a horizontal line to the east until it hits the last vertex of the line. Likewise, extend a vertical line to the south as far as possible. Then complete the rectangle with these two sides (see Figure 4.13). Let the positive partition α be the partition of transpose Greene's Theorem (Theorem 1) on the poset consisting of the markings in the subrectangle contained in the fundamental region of the staircase. As mentioned in Chapter 2, this can be done in any of the usual ways, such as performing transpose matrix-ball on this subrectangle, or applying the transpose Fomin rules to get the *final* partition.

Let the negative partition β be given by transpose Greene's Theorem on the poset of markings in the subrectangle contained in the dual fundamental region. By dual symmetry of affine growth diagrams, this is the same as the partition given by the corresponding rectangle in the fundamental region (see Figure 4.13). We claim that the local rule is satisfied by these weights in every unit square of the diagram.

The following lemma is the key to the proof that this map gives an affine growth diagram with top line an oscillating tableau. Let P_{11} be a poset containing P_{00} such that $P_{01} = P_{00} \cup e_1$, $P_{10} = P_{00} \cup e_2$, and $P_{11} = P_{00} \cup e_1 \cup e_2$ where $e_1 \neq e_2$. For $i, j = 0, 1$ let λ_{ij} be the partition given by transpose Greene's Theorem for the poset P_{ij}. For two partitions $\mu \subseteq \nu$, let ν/μ denote the skew diagram of the cells

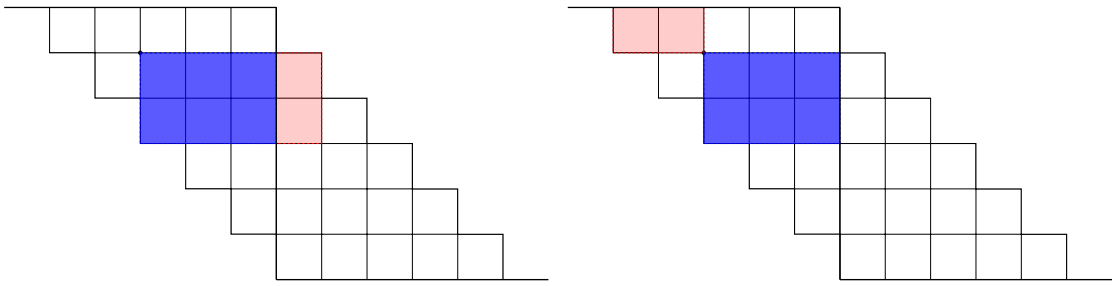

Figure 4.13: Rectangles for the reverse map, desired vertex indicated

in ν but not in μ. When referring to a cell (x, y) of a partition, x denotes the row and y the column, with indices increasing down the partition and to the right.

Lemma 28 ([28, A.3.1])**.** *Assume $\lambda_{01} = \lambda_{10}$. Let $a = (x_a, y_a)$ be the cell $\lambda_{01}/\lambda_{00}$ and $b = (x_b, y_b)$ be the cell $\lambda_{11}/\lambda_{01}$.*

- *If e_1 and e_2 are extremal elements of P_{11} of different types (one minimal, the other maximal), then $x_b = x_a$ or $x_b = x_a + 1$.*
- *If e_1 and e_2 are extremal elements of P_{11} of the same type (both minimal or both maximal), then $x_b \leq x_a$.*

In terms of growth diagrams, the local setup is the following.

$$\begin{array}{cc} \lambda_{00} & \lambda_{01} \\ \lambda_{10} & \lambda_{11} \end{array}$$

The lemma refers to rule (2) of the Fomin rules. Note that a Fomin growth diagram can only run into the second case of the lemma, i.e. adding an eastward column and southward row can only reveal two maximal marked squares. For affine growth diagrams, as we will see in the proof, the first case of the lemma can occur.

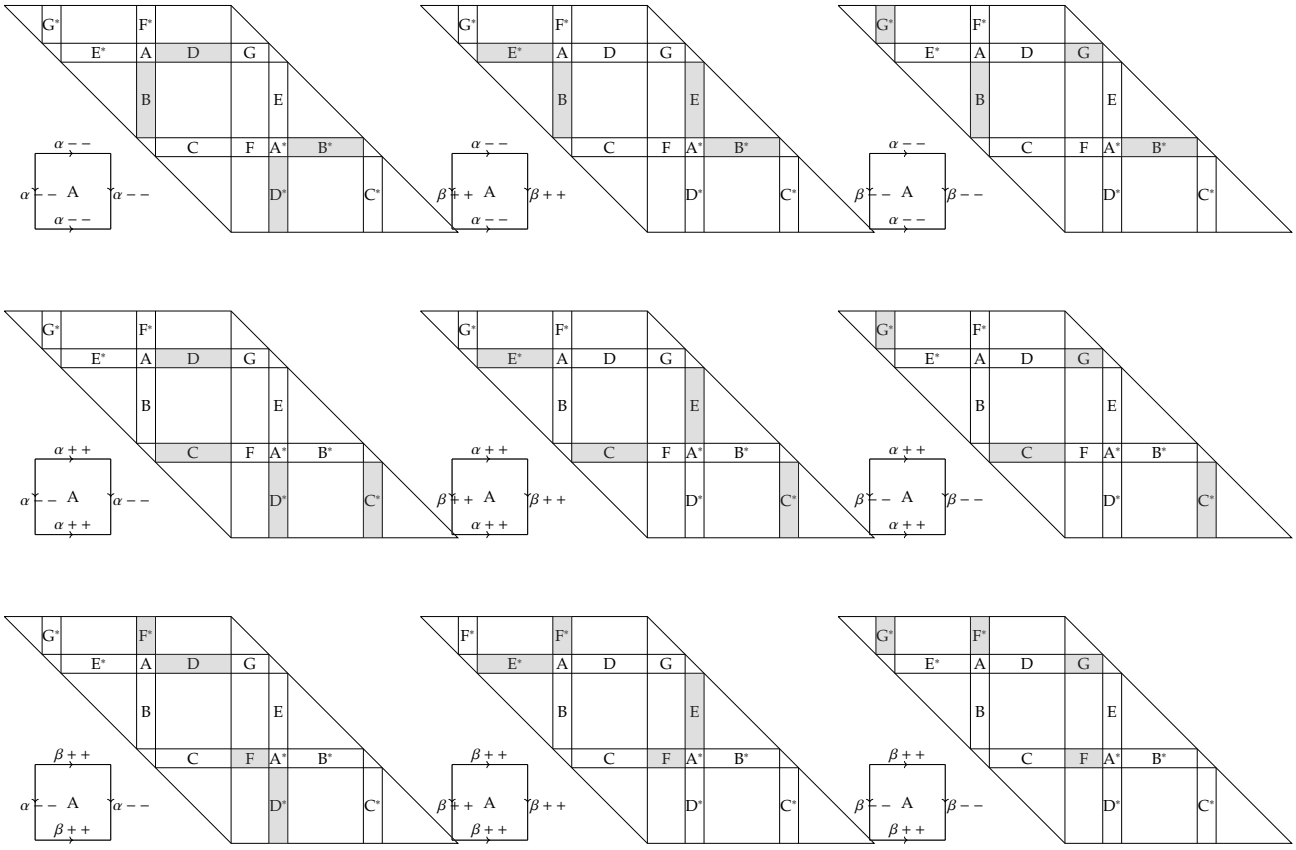

Figure 4.14: The 9 cases

Proposition 29. *Write a fixed-point-free involution π into a staircase diagram as described above. The above map produces an affine growth diagram $\psi(\pi)$ with 0th line labeled by an oscillating tableau. The fixed-point-free involution associated to this affine growth diagram by Proposition 25 is the original π, i.e. $\phi \circ \psi(\pi) = \pi$.*

Proof. We must show that each unit square is labeled by weights that satisfy the local growth rule. Choose a unit square A and consider the relevant regions in Figure 4.14. The fundamental region has two hooks that contain A, namely one consisting of G*, E*, A, D, and the other consisting of F*, A, B, C. Each of these hooks must have exactly one region with one marked square. Suppose first that square A is not marked, so that there are 9 possible ways of marking the hooks.

In each case, we show that the vertex labels of A satisfy the growth rule. There are 3 pairs of dual cases, so we will only have to analyze 6 cases, only two of which are interesting. The dual pairs are 1c with 2b, 1a with 3b, and 2a with 3c. By this we mean that if we know that the local rule is satisfied at square A, then by transpose symmetry we know that it is satisfied by square A^*, but the setup of square A^* of 1c is that of square A of 2b.

Let α and β denote the positive and negative partitions respectively of a dominant weight. When discussing traversing edges of A we will assume that the horizontal edges are traversed from W to E and vertical edges from N to S, unless specified otherwise. Note that traversing a single edge of square A results in a change in exactly one of α or β by exactly one box. For example, traversing a horizontal edge changes α's rectangle by eliminating region B and adding region C. Since a hook can have a mark in at most one of regions B or C, α can gain or lose at most one box. If neither region B nor C is marked then F^* is marked, hence β increases. Analyzing the vertical edges of square A is similar.

Throughout, let ν, μ, λ, ρ denote the positive partitions of the weight labels of the square A, and let $\nu', \mu', \lambda', \rho'$ denote the negative partitions.

1. **Region B is marked.**

In this case, since B is marked and C is not, traversing either horizontal edge results in an α decrease. The negative partition β does not change.

(a) Region D is marked; in this subcase, traversing the vertical edges results in an α decrease. Only the positive partition of the dominant weight is

changed across the edges of square A.

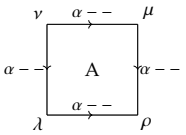

Thus, ν is one box larger than λ and μ, which are one box larger than ρ. If $\mu \neq \lambda$, then $\rho = \mu \cap \lambda$ and $\nu = \mu \cup \lambda$ simply by the properties of Young's lattice. This gives two different paths of length two in Young's lattice with common start and end points, so if ν and μ differ in row i, λ and ρ must differ in row i. Likewise, for some other row j when interchanging the roles of λ and μ. Thus, the local growth rule is satisfied when $\lambda \neq \mu$.

Now suppose that $\lambda = \mu$. To show that $\rho = \text{sort}(\lambda - (\nu - \mu))$ we appeal to the transpose Fomin local rule. Since region C is empty, it does not contribute anything to the Fomin growth diagram construction for μ and ρ. Thus, filling ν's rectangle southeast to northwest using Fomin's growth rules will give the correct label for all four vertices of A. Consider the effects of Fomin's transpose rule at square A. Suppose a box was added to ρ in column i to get $\mu = \lambda$. Then adding another box to column $i + 1$ gives a valid partition, and Fomin says that this is in fact the rule to get ν. From the perspective of the rows of the partitions, suppose the box in column $i + 1$, row j is removed from ν to get μ. Then removing a box from column i, is the same as removing box in column i and row j and sorting. Thus, the local rule is satisfied.

(b) Region E^* is marked; in this subcase, traversing the vertical edges results in a β increase.

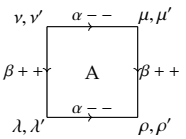

Therefore, $\nu = \lambda$ and $\mu = \rho$, so the difference of ν and μ equals the difference of λ and ρ, hence satisfying the local rule.

(c) Region G* is marked; in this subcase, traversing the vertical edges results in a β decrease. As in the previous subcase, $\nu = \lambda$, $\mu = \rho$ and $\nu' = \mu'$, $\lambda' = \rho'$, so the local growth rule is satisfied.

2. **Region C is marked**; in this case traversing the horizontal edges results in an α increase.

(a) Region D is marked; traversing vertical edges results in an α decrease. In this case, μ covers ν and ρ, which both cover λ. As in case (1a), if $\nu \neq \rho$, then $\mu = \nu \cup \rho$ and $\lambda = \mu \cap \rho$. Suppose that $\nu = \rho$. This is the other interesting case. Let P_{00} be the poset of the markings in λ's rectangle, and P_{10}, P_{10}, P_{11} those of ν, ρ, μ respectively. Thus, going from P_{00} to P_{11} introduces two extremal elements. However, unlike the usual Fomin scenario of case (1a), one of the elements is minimal and the other maximal (see Figure 4.14).

Let $a = (x_a, y_a)$ and $b = (x_b, y_b)$ be the cells of ν/λ and μ/ν respectively. By bullet point 1 of Lemma 28 we know that $x_b = x_a$ or $x_b = x_a + 1$. If $x_b = x_a$, then both boxes must have been added to the same row x_b. Thus, traversing the W vertical edge removes a box from row x_b, and likewise when traversing the E vertical edge, so the local growth rule is satisfied. Now suppose $x_b = x_a + 1$. In each of the four relevant rectangles, consider taking the dual order poset on the markings, i.e. the minimal elements are now to the northeast rather than southeast. By Proposition 3 the associated partitions are just the transposes of the original, i.e. $\lambda^t, \nu^t, \rho^t, \mu^t$. Then $\nu^t/\lambda^t = (y_a, x_a)$ and $\mu^t/\nu^t = (y_b, x_b)$. However, the two extremal elements are still of different types. Thus, by another application of Lemma 28, $y_b = y_a$ or $y_b = y_a + 1$. But $x_b = x_a + 1$ and $y_b = y_a + 1$ cannot simultaneously hold,

so we conclude that $y_b = y_a$. The situation is then $x_b = x_a + 1$ and $y_b = y_a$, meaning that the two boxes were added in the same column, in row x_a then row $x_a + 1$. Thus, traversing the W vertical edge removes a box from row x_a, and traversing the E vertical edge also removes a box from row x_a followed by a sort. Traversing the N edge adds a box to row $x_a + 1$, and traversing S edge adds a box to row $x_a + 1$ followed by a sort. Thus, the local growth rule is satisfied.

The remaining cases are dual, or similar to the previous ones.

Now suppose that square A is marked. Then ν covers $\rho = \mu = \lambda$, and ρ' covers $\nu' = \mu' = \lambda'$. Traversing the N and W edges results in an α decrease, and traversing the S and E edges results in a β increase.

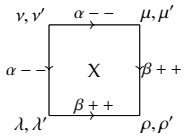

Since the X in square A dominates all of the other marks in ν's rectangle, it must be part of the longest chain. Therefore, ν/λ, is the last box in the first column, and likewise ρ'/λ' is the last box in the first column. Therefore, traversing the W edge removes a box in the last nonzero row. Removing a box from the same row of μ results in a new negative box and sorting, which coincides with adding a box to the bottom of the first column of ρ'. The same is true for the horizontal edges, so the local growth rule is satisfied.

From analyzing all of the cases we see that a switch from editing α to editing β or vice versa only occurs in this final case when the square is marked by π. Thus the marked squares according to Proposition 25 coincide with the original marked squares of π. □

Proposition 29 showed that $\phi \circ \psi$ is the identity map on fixed-point-free involutions. We now prove that the maps ϕ and ψ are two-sided inverses.

Proposition 30. *The composition $\psi \circ \phi$ is the identity on affine growth diagrams whose 0th line is labeled by an oscillating tableau.*

Proof. To prove that this composition is the identity map we must show that filling the diagram (and the X's) according to the local rule of Theorem 12 gives vertex labels that agree with transpose Greene's Theorem on the corresponding rectangles, or equivalently with the transpose Fomin rules. By dual symmetry we only need to check this in the fundamental region. For the remainder of the proof, whenever we refer to Greene's Theorem or Fomin's rule we mean the transpose versions.

As before, let α denote the positive partition of a weight, and β the negative partition. We will show that the local rule of Theorem 12 gives the same result as the Fomin growth rules for α when filling the fundamental region from the SE to the NW and for β when filling from the NW to the SE. When passing over the horizontal and vertical lines that delineate the fundamental region the roles of α and β are flipped. Let's start by constructing the β vertex labels as we fill the diagram from NW to SE. The same argument holds for α when filling in the diagram from the SE to NW.

The first line has $\beta = \emptyset$ at every vertex, which trivially agrees with Greene's Theorem on the empty rectangle. Consider the base case of the second horizontal line. Traversing the first vertical edge in the S direction decreases α. By the the local rule, this continues until α decreasing switches to β increasing at the marked square (since the fundamental region contains the entire first row, this

must happen at some point), but then this means that β is a single box for the rest of the second line. This coincides with Greene's on the β rectangle.

For the general step, suppose by induction that the β vertex labels of the $k-1$ line agree with Greene's Theorem on their corresponding β rectangles. There are two cases, either $\pi(k-1) > k-1$ or $\pi(k-1) < k-1$. In terms of hook regions, the first case corresponds to region D, A, or E* being marked, and the second case to region G* (see Figure 4.14).

Consider the first case, so that the row in the fundamental region contains a mark. Assume by induction that the first $i-1$, β vertex labels of line k agree with Greene's Theorem within their respective rectangles. Consider the square A whose SE vertex is the ith vertex of the kth line, and let the negative partition labels of the four vertices be ν, μ, ρ, λ as follows.

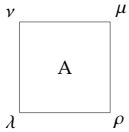

There are three cases depending on where the marked square is relative to A. Since column G* is unmarked, it doesn't contribute anything to the posets associated to ν and μ's rectangles. Thus, we can remove column G* and run the Fomin rules on ρ's rectangle to still get the correct vertex labels, λ, μ, ν, ρ at square A.

1. Region D is marked.

   (a) Region F* is not marked.

   By induction ν, μ, λ are given by Fomin's rules, so traversing the W vertical edge and the N horizontal edge doesn't change β. Then by the local rule ρ must also equal $\nu = \mu = \lambda$, which coincides with Fomin.

(b) Region F* is marked.

By induction traversing the W vertical edge doesn't change β, so $\nu = \lambda$, but traversing the N edge increases β. The local rule implies $\rho = \mu$, which coincides with Fomin's rules.

2. Region A is marked.

By induction traversing the W and N edges doesn't change β. Since the square is marked, the local rule implies that traversing the S and E edges increases β in the first column by 1. This is precisely Fomin's growth rule, so ρ agrees with Greene's Theorem.

3. Region E* is marked.

(a) Region F* is unmarked.

Traversing the N edge doesn't change β, so $\rho = \lambda$, which coincides with Fomin's rules.

(b) Region F* is marked.

By induction traversing the N and W edges causes β increases. By the local rule, ρ covers μ and λ, which cover ν. The same must hold when calculating the vertex labels by Fomin's rules. Thus, if $\mu \neq \lambda$, then in both cases we must have $\rho = \mu \cup \lambda$ and $\nu = \mu \cap \lambda$.

Now suppose that $\mu = \lambda$. Suppose a box was added in cell (i, j) to get μ from ν. Then by the local rule a box is added in row i to μ, which is then sorted to get ρ. With or without sorting, the box must have been added in column $j + 1$. This coincides with Fomin's growth rule, which says that if a box was added in column j of ν to get μ, then a box is added to column $j + 1$ to μ to get ρ.

Now suppose $\pi(k - 1) < k - 1$, so that G* is marked. By induction traversing the

W edge causes a β decrease.

1. Region F* is not marked.

By induction traversing the N edge doesn't change β, so $\nu = \mu$. The local rule implies $\lambda = \rho$, which coincides with Greene's Theorem because the F* column is empty.

2. Region F* is marked.

By induction traversing the N edge causes a β increase. Thus by the local rule, μ covers ν and ρ, which cover λ. As before, if $\rho \neq \nu$, then it must be that $\mu = \nu \cup \rho$, $\lambda = \nu \cap \rho$ for both the affine growth rule and Greene's Theorem.

Suppose that $\rho = \nu$. Let $a = (x_a, y_a)$ be the cell of ρ/λ and $b = (x_b, y_b)$ the cell of μ/ρ. We claim that either $(x_b, y_b) = (x_a + 1, y_a)$ or $(x_b, y_b) = (x_a, y_a + 1)$. The local rule says to add a box in row x_b to λ and sort to get ρ. If no sorting is needed then $x_a = x_b$ and $y_b = y_a + 1$. If sorting is needed then $y_a = y_b$, which forces $x_b = x_a + 1$.

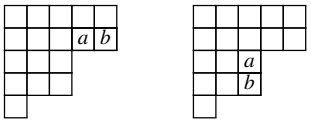

Consider the posets P_a, P_b, P_c, P_d of the rectangles associated to the vertices of square A, and let $\alpha = \theta(P_a), \beta = \theta(P_b), \gamma = \theta(P_c), \delta = \theta(P_d)$ where θ applied to a poset gives the partition of Greene's Theorem.

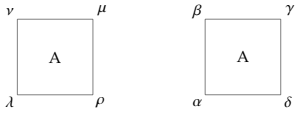

By induction $\alpha = \lambda, \beta = \nu, \gamma = \mu$. We must show that $\delta = \rho$. Note that $P_a \subsetneq P_d$ because they differ by a single maximal element, and $P_d \subsetneq P_c$ because they differ by a single minimal element. Likewise, $P_a \subsetneq P_b \subsetneq P_c$. In terms of partitions this

means $\alpha \subsetneq \delta \subsetneq \gamma$ and $\alpha \subsetneq \beta \subsetneq \gamma$. But γ is of the form above including the boxes containing a, b, while α is of the same form excluding the boxes a, b. There can be only one partition between α and γ, so $\delta = \beta = \nu = \rho$.

□

Let $k = n/2$ and consider the case when all of the X marks are contained in the $k \times k$ square inside of the fundamental region. By the previous proposition, the α partitions of the vertices can be computed by transpose Fomin rules starting from the SE vertex of the square, giving a $k \times k$ growth diagram when ignoring the β partitions. Likewise, ignoring the α partitions, the β partitions give a growth diagram starting from the NW vertex. In this case, the first k steps of the corresponding oscillating tableau are increases and the final k steps are decreases. Hence, the proof of the previous proposition gives the following corollary (see Figure 4.15).

Corollary 31. *When $\vec{\lambda} = (\omega_1, \ldots, \omega_1, \omega_1^*, \ldots, \omega_1^*)$, the $k \times k$ square inside of the fundamental region of an affine growth diagram is a Fomin growth diagram towards the northwest when ignoring the negative parts. Likewise, it is a Fomin growth diagram towards the southeast when ignoring the positive parts. A similar statement holds for the $k \times k$ square in the dual-fundamental region.*

Consider an oscillating tableau of length $n = 2k$ where the partitions increase for the first k steps and then decrease for the remaining k steps. This is the same information as a pair of same-shape standard Young tableaux (P,Q), and it is natural to ask where they map to under Proposition 27. If the pair (P, Q) maps to

$$\sigma = \begin{pmatrix} 1 & 2 & 3 & \cdots & k-1 & k \\ \sigma_1 & \sigma_2 & \sigma_3 & \cdots & \sigma_{k-1} & \sigma_k \end{pmatrix}$$

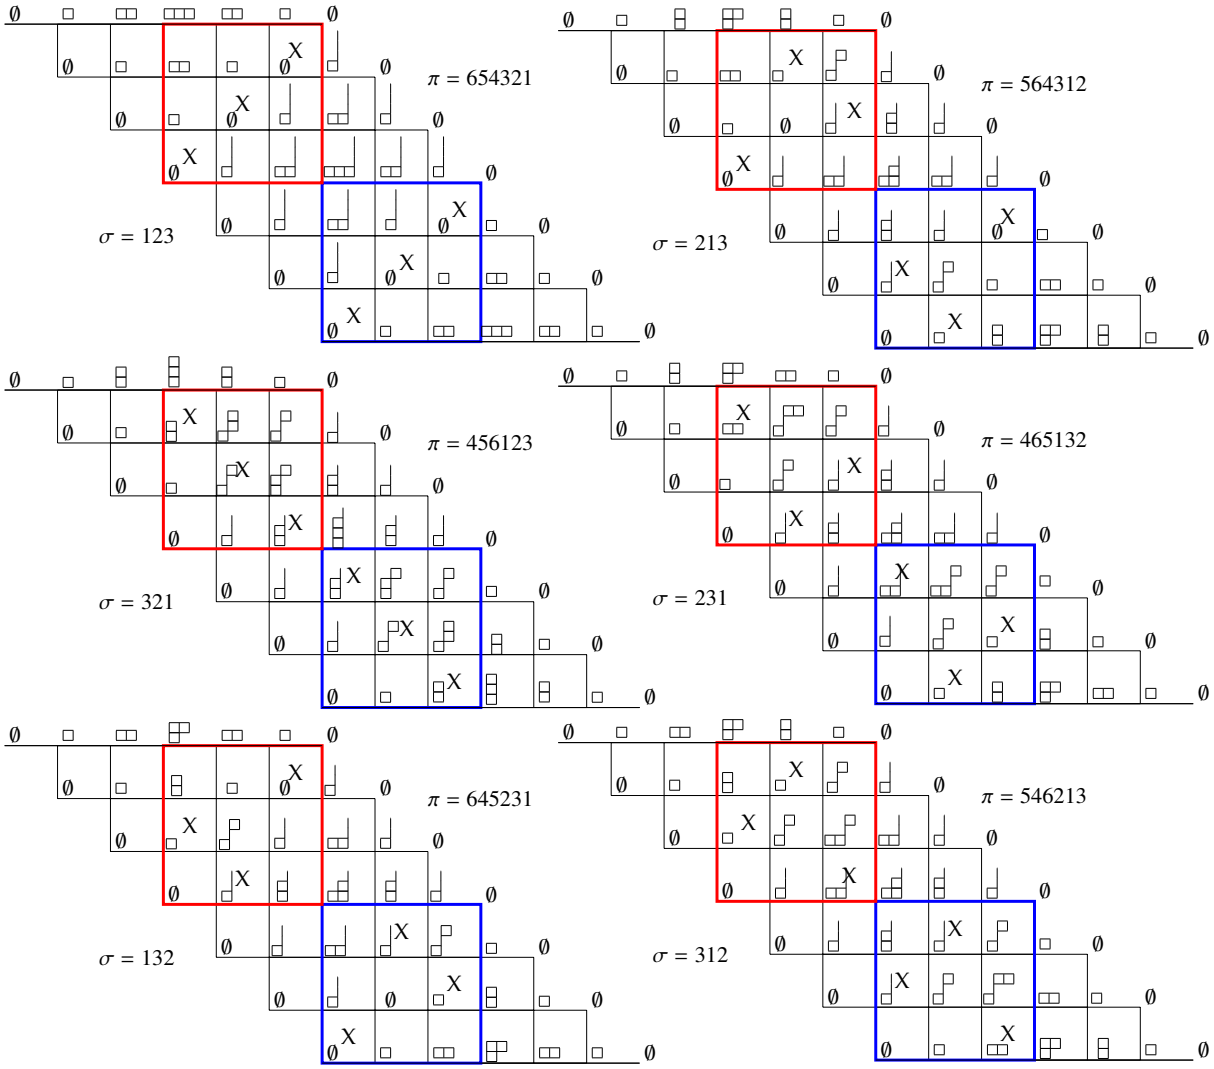

Figure 4.15: The six affine growth diagrams of the RS correspondence

under the RS correspondence, then the pair viewed as an oscillating tableau maps to the unique involution π where the reversal $\sigma_r = \left(\begin{smallmatrix} 1 & 2 & \cdots & k-1 & k \\ \sigma_k & \sigma_{k-1} & \cdots & \sigma_2 & \sigma_1 \end{smallmatrix}\right)$ gives the values of π on $k+1, \ldots, 2k$,

$$\pi = \left(\begin{smallmatrix} \sigma_1 & \sigma_2 & \cdots & \sigma_{k-1} & \sigma_k & k{+}1 & k{+}2 & \cdots & 2k{-}1 & 2k \\ 2k & 2k{-}1 & \cdots & k{+}2 & k{+}1 & \sigma_k & \sigma_{k-1} & \cdots & \sigma_2 & \sigma_1 \end{smallmatrix}\right). \tag{4.1}$$

To see this, begin by writing $[P, \emptyset], [Q, \emptyset]$ on the first line of a growth diagram to specify the two chains of partitions that meet in the middle encoded by the standard tableaux P and Q (see Figure 4.16). The $\emptyset$'s in the second slots indicate

that the corresponding dominant weights contain no negative parts. We are only interested in labeling line segments where either the first or second slot is $\emptyset$.

Consider the right triangle with one side the line segment labeled by $[P, \emptyset]$ and hypotenuse, the diagonal, labeled by $\emptyset$ at each vertex. Since P specifies a chain of partitions with only positive parts, the diagram on this triangle is a genuine growth diagram on partitions. In this case, with a square labeled as below, the sorting local rule can be rephrased in a simpler manner:

> If μ/ν is a pair of adjacent squares, then $\rho = \lambda$.
>
> Otherwise, ρ is the unique partition such that $\nu \subset \rho \subset \mu$ and $\rho \neq \lambda$.

This is precisely the rule given in [38] and [34, Chp. 7, App. 1] that computes jeu de taquin, so the other edge of the right triangle is labeled by $[ev(P), \emptyset]$. The same applies to the right triangle defined by the vertical line segment labeled by $[\emptyset, Q]$. Since this triangular portion of the diagram is labeled by negative partitions, the other side of the triangle is labeled by $[\emptyset, ev(Q)]$. By similar observations or dual symmetry, label the other segments as in Figure 4.16.

Now suppose that $(P, Q) \leftrightarrow \sigma$ under the RS correspondence. Then by (2.2), $(P^t, ev(Q)^t) \leftrightarrow \sigma \cdot w_0$ under the RS correspondence, or equivalently $(P, ev(Q)) \leftrightarrow \sigma \cdot w_0$ under the transpose RS correspondence. Consider the $k \times k$ square in the dual fundamental region, and consider just the positive parts of the weight labels, so that it is a transpose Fomin growth diagram growing towards the pair $(P, ev(Q))$ in the southeast. Thus, this square contains the permutation $\sigma \cdot w_0$, so

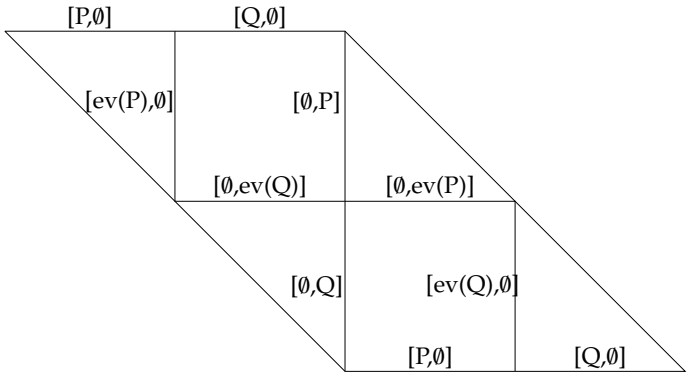

Figure 4.16: Labeling a diagram for standard tableaux (P,Q)

for example the very last row has an X in column $\sigma(1)$, the penultimate in $\sigma(2)$, and so forth. Extending to the entire affine growth diagram by dual symmetry, this gives the fixed point free involution π specified in (4.1).

Hence the RS correspondence lives inside of the bijection of Proposition 27. For $k = 3$ there are 15 fixed-point-free involutions in S_6 corresponding to 15 oscillating tableaux of length 6. Six of these oscillating tableaux are equivalent to pairs of same-shape standard tableaux, which correspond to the six permutations of S_3 embedded into the fixed-point-free involutions of S_6 as described above. See Figure 4.15. The interested reader is invited to write out all 15 growth diagrams.

4.4 Knuth version

Recall Knuth's generalization of the RS correspondence, a weight-preserving bijection between matrices of natural numbers and pairs of same-shape semistandard tableaux. Roby showed how the full RSK correspondence fits into the Fomin growth diagram framework [28, Chp. 4]. He also studied various other

bijections arising as a consequence of growth diagrams, for example a generalization of the Stanley–Sundaram bijection discussed below.

In this section, the λ^i are free to be any minuscule weights, not just ω_1 and ω_1^*. This introduces natural-number entries rather than just X's in the affine growth diagrams, and oscillating tableaux differing by vertical strips rather than just single boxes. In Theorem 36 we prove a bijection involving semistandard oscillating tableaux, within which lives the full RSK correspondence. As a corollary, we see that Fomin-Roby growth diagrams appear within affine growth diagrams for $\vec{\lambda}$ consisting of fundamental weights preceding dual-fundamental weights. We are still assuming that m is large enough, in particular $m \geq \sum|\lambda^i|$ where the sum is over the λ^i that are fundamental weights (as opposed to dual fundamental weights, which would give the same sum).

Definition 32. A **semistandard oscillating tableau** of length k is a sequence of partitions $(\mu^0 = \emptyset, \mu^1, \mu^2, \ldots, \mu^k = \emptyset)$ such that for all i, μ^i and μ^{i+1} differ by a vertical strip of boxes, i.e. no two boxes in the same row.

Note that this is slightly different from the definition in [28].

Definition 33. For an $n \times n$ symmetric matrix with natural-number entries let $\mathbf{r_i}$, respectively $\mathbf{c_i}$, denote the sum of the entries of the ith row, respectively column, up to and including the main diagonal entry.

Definition 34. A natural-number matrix is a **natural-number fixed-point-free involution** if it is a symmetric matrix such that $r_i = 0$ or $c_i = 0$ for all i.

The $r_i = 0$ or $c_i = 0$ condition prevents nonzero entries on the diagonal. These matrices are the natural-number analogues of fixed-point-free involutions, since

a fixed-point-free involution satisfies either $r_i = 0$ and $c_i = 1$, or $r_i = 1$ and $c_i = 0$ for all i. We can also view a symmetric matrix by restricting to the NE part above the diagonal, so that the $r_i = 0$ or $c_i = 0$ condition says that either the horizontal part or vertical part of each hook is empty. We need one lemma from [28], which we rephrase for our context.

Lemma 35 ([28, Thm. 4.1.4])**.** *In the transpose Fomin algorithm, let $\mu^i, \mu^{i+1}, \ldots, \mu^j$ be consecutive shapes growing along the bottom of the $n \times n$ square and let the partial permutation above them be given by $w_i, w_{i+1}, \ldots, w_j$. If $w_i < w_{i+1} < \ldots < w_j$, then μ^j and μ^i differ by a vertical strip and the boxes are added sequentially in $\mu^i, \mu^{i+1}, \ldots, \mu^j$ from top to bottom.*

Compare theorem 36 with [28, Thm. 4.4.3]. There Roby allows fixed points and situations where c_i and r_i are both nonzero for the same i. The proof will just be the relabeling scheme of Knuth's [20] original generalization of Schensted insertion and matrix-ball construction phrased in terms of growth diagrams (see [Roby, p. 63] and Figure 4.17).

Theorem 36. *There is a bijection between semistandard oscillating tableaux of length k and $k \times k$ natural-number fixed-point-free involutions. Furthermore, the row content gives the increases, and the column content gives the decreases in the oscillating tableau.*

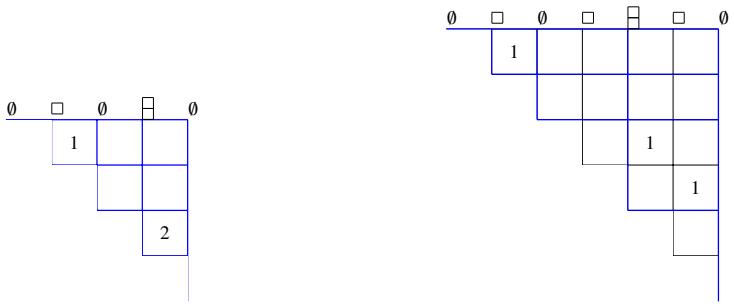

Figure 4.17: Applying Knuth's relabeling scheme

Proof. Let $\vec{\mu}$ be a length k semistandard oscillating tableau and the minuscule $\vec{\lambda}$ given by the differences $\mu^{i+1} - \mu^i = w \cdot \lambda^i$. Create a standard oscillating tableau by breaking up any step that adds a vertical strip of length j into j single box steps such that the higher boxes are added first. Likewise, break up a step that removes a vertical strip into individual steps such that lower boxes are removed first. Write this new standard oscillating tableaux as the labels of the top line of an affine growth diagram, but draw the lines that start at the labels from the original semistandard oscillating tableaux in blue. Now build the affine growth diagram as usual to obtain the X marks. Note that this requires $m \geq \sum|\lambda^i|$ where the sum is over fundamental λ^i. Finally, sum the number of X's in each blue rectangle and make each blue rectangle its own square in a $k \times k$ matrix. This is the desired matrix with natural-number entries. Note that either $c_i = 0$ or $r_i = 0$ for all i because the original semistandard oscillating tableaux either increased or decreased at each step. See Figure 4.17.

Conversely, start with a $k \times k$ symmetric matrix with natural-number entries such that $c_i = 0$ or $r_i = 0$ for all i. Convert this matrix into a $K \times K$ fixed-point-free involution matrix as follows. If $c_i \neq 0$, then divide column i and row i into c_i columns and rows by adding $c_i - 1$ extra lines, and similarly if $r_i \neq 0$. In this expanded $K \times K$ matrix, mark in blue the lines that came from the original $k \times k$ matrix. Thus the single cell (i, j) has been divided into c_j columns and r_i rows. Suppose that the NWmost nonzero natural-number entry is a_1 at cell (i_1, j_1). Fill its subdivided cell with a_1 individual X's arranged in a diagonal from the NW to SE. Continue to the next NWmost nonzero entry a_2 at (i_2, j_2). If the a_1 and a_2 entries were in the same row, i.e. $i_1 = i_2$, then the first a_1 rows are already taken, hence add a_2 X's in the northmost rows possible such that the X's are arranged in a diagonal. Likewise, if $j_1 = j_2$. Proceed in this way to consistently mark the

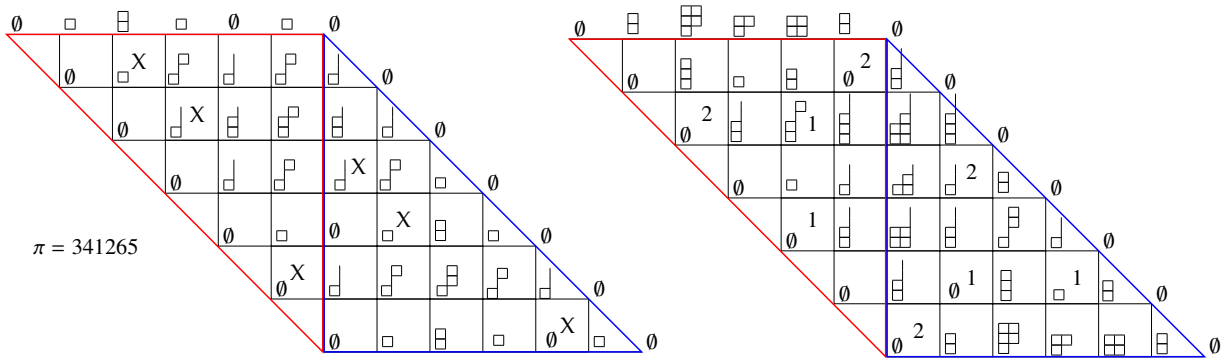

Figure 4.18: Standard and semistandard oscillating tableau examples.

refined $K \times K$ matrix to give a fixed-point-free involution.

Now fill in the weight labels of the affine growth diagram by Greene's Theorem. Note that this requires $m \geq \sum c_i = \sum r_i$. Reading off the vertex labels of the blue lines gives the semistandard oscillating tableau corresponding to the original matrix. This is a bijection by the previous lemma. □

The previous proof is a "Knuthification" of the ω_1 and ω_1^* case. However, one may fill in the affine growth diagram directly starting from a minuscule path by recording the number of places that switched from decreasing α to increasing β or vice versa (see the right diagram in Figure 4.18). These partitions are then given by the general matrix-ball construction on natural numbers.

Oscillating tableaux that increase for p consecutive steps and then decrease for q consecutive steps correspond to matrices with nonzero entries in the $p \times q$ rectangle in the fundamental region. Roby showed how to extend Fomin growth diagrams to the Knuth case and we have the following. These Fomin–Roby growth diagrams appear within affine growth diagrams (see the right diagram in Figure 4.18).

Corollary 37. *In the setting of the previous paragraph, the $p \times q$ rectangle in the fundamental region of an affine growth diagram is a Fomin-Roby growth diagram towards the northwest when ignoring the negative parts. Likewise, it is a Fomin-Roby growth diagram towards the southeast when ignoring the positive parts. A similar statement holds for the $q \times p$ rectangle in the dual-fundamental region.*

CHAPTER 5

PROOF OF THE LOCAL RULE

In this chapter we prove Theorem 12. To do so we introduce Knutson–Tao hives following [22] and [19].

5.1 Hives

Consider the triangle $\Delta_m^3 = \{(i, j, k) : i, j, k \geq 0, i + j + k = m\}$ with its $\binom{m+2}{2}$ integer points. A **size-m, 3-hive** is an equivalence class of labelings by integers, $f : \Delta_m^3 \to \mathbb{Z}$, such that for any unit rhombus, the sum across the short diagonal is at least the sum across the long diagonal. Two such functions are equivalent if they differ by a constant function. In terms of indices the rhombus inequalities are the following.

$$f_{i,j,k} + f_{i,j+1,k-1} \geq f_{i+1,j,k-1} + f_{i-1,j+1,k}$$

$$f_{i,j,k} + f_{i+1,j-1,k} \geq f_{i+1,j,k-1} + f_{i,j-1,k+1}$$

$$f_{i,j,k} + f_{i+1,j,k-1} \geq f_{i,j+1,k-1} + f_{i+1,j-1,k}$$

An equivalent condition to satisfying these inequalities is that if f is extended linearly on each unit triangle to a real-valued function, then f is convex. In particular, if the vertex labels along one side of the hive are

$$a_0 = f_{m,0,0}, \quad a_1 = f_{m-1,1,0}, \quad \ldots, \quad a_m = f_{0,m,0},$$

then $\vec{a}$ is convex, that is, $\lambda = (a_1 - a_0, a_2 - a_1, \ldots, a_m - a_{m-1})$ is a weakly decreasing list of m integers, a dominant weight of G. Note that taking the differences in

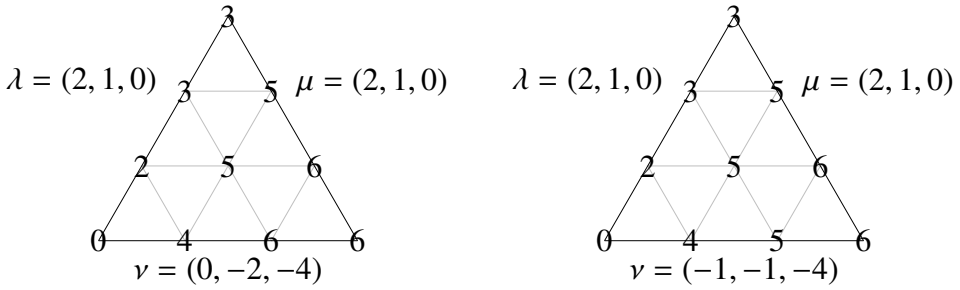

Figure 5.1: Examples of size-3, 3-hives

the reverse order gives the dual dominant weight

$$\lambda^* = (a_{m-1} - a_m, \ldots, a_0 - a_1).$$

Reading clockwise, consider the other sides of the hive and their associated weights, μ and ν. Then the triple λ, μ, ν is called the **boundary** of the 3-hive. When referring to sides of a hive, the context should make it clear whether we mean the actual hive values a_i or the dominant weights given by their successive differences, keeping in mind that switching directions gives the dual weight. See Figure 5.1 for two examples of 3-hives. Hives give a combinatorial model for Littlewood–Richardson coefficients:

Theorem 38 ([21])**.** *The number of 3-hives with boundary λ, μ, ν equals*

$$c_{\lambda,\mu,\nu} = \dim \left(V_\lambda \otimes V_\mu \otimes V_\nu \right)^G.$$

Define the size-m, $(n-1)$-dimensional simplex

$$\Delta_m^n = \left\{ (i_1, \ldots, i_n) \mid i_j \in \{0, 1, \ldots, m\}, \sum i_j = m \right\}.$$

A **size-m, n-hive** is an equivalence class of labelings, $f : \Delta_m^n \to \mathbb{Z}$, such that the restriction to every two-dimensional face is a 3-hive and the following "octahedron recurrence" is satisfied. A function f satisfies the **octahedron recurrence** if

for any $i = (i_1, i_2, \ldots, i_n)$ such that $\sum i_j = m - 2$ and for any $1 \leq a < b < c < d \leq n$,

$$f_{i+e_a+e_c} + f_{i+e_b+e_d} = \max\left(f_{i+e_a+e_b} + f_{i+e_c+e_d}, f_{i+e_a+e_d} + f_{i+e_b+e_c}\right),$$

where e_j is the vector with a 1 in the jth positions and 0s elsewhere. By abuse of notation, we will not distinguish between an n-hive and its underlying $(n-1)$-simplex, saying things like, "the 1-skeleton of the n-hive."

The **boundary of an n-hive** consists of the n dominant weights $\lambda^1, \ldots, \lambda^n$ of differences along the edges in the cycle of vertices

$$(m, 0, 0, \ldots, 0), (0, m, 0, \ldots, 0), (0, 0, m, 0, \ldots, 0), \ldots, (0, 0, \ldots, 0, m), (m, 0, 0, \ldots, 0).$$

Note that this is an n-gon inside of the 1-skeleton of an n-hive, so certainly not the boundary of the entire n-hive. Instead, it is the boundary of surfaces in the 2-skeleton that we will be interested in. Generalizing the $n = 3$ case gives the following theorem.

Theorem 39 ([22, 19, 11])**.** *The number of n-hives with boundary $\vec{\lambda}$ equals*

$$c_{\vec{\lambda}} = \dim\left(V_{\lambda^1} \otimes \cdots \otimes V_{\lambda^n}\right)^G.$$

It is possible to visualize the octahedron recurrence when $n = 4$, so that the 4-hive is a labeling of lattice points in a tetrahedron. Suppose the tetrahedron is balanced on an edge and is oriented so that the view from the top shows two faces and likewise the view from the bottom. Then, given a unit octahedron with vertices e, e', a, b, c, d where e and e' are the top and bottom vertices respectively, and a, b, c, d the vertices around the equatorial square, the octahedron recurrence says that

$$e' = \max(a + c, b + d) - e.$$

This gives a deterministic procedure to "excavate" a unit octahedron to reveal the vertex label that it was blocking from view. In particular the 4-hive is determined by its restriction to the top two faces, or the bottom two faces. For general n, certain subsets of the 2-skeleton are enough to determine the entire n-hive, as explained in Proposition 41.

5.2 Hives for minuscule weights

Consider a 3-hive with boundary λ, μ, ν, where λ is minuscule. In this case, the tensor product $V_\lambda \otimes V_\mu$ is multiplicity free, so $c_{\lambda\mu\nu}$ equals 0 or 1. For ν that do occur in the decomposition, Theorem 38 says there is exactly one hive with boundary λ, μ, ν, so specifying the boundary is enough to determine the entire hive.

For an n-hive with boundary consisting of minuscules $\vec{\lambda}$, we would like to know which edges of the 1-skeleton are sufficient to determine the entire n-hive. For example, for a 4-hive with boundary $\lambda^1, \lambda^2, \lambda^3, \lambda^4$, fixing one of the remaining two edges, say the top one, to be μ determines a unique 4-hive (if one exists for these labels): By multiplicity freeness there is exactly one 3-hive with boundary $\lambda^1, \lambda^2, \mu^*$ and exactly one with boundary $\mu, \lambda^3, \lambda^4$. This determines the top two faces of the tetrahedron, which determine the rest of the 4-hive by repeated use of the octahedron recurrence. This process generalizes to n-hives by specifying the labels μ^i along the interior edges of an extroverted triangulation.

To an extroverted triangulation we associate the corresponding triangulated disk with boundary $\vec{\lambda}$ in the 2-skeleton of the n-hive. For $n = 4$ there are two extroverted triangulations of the square, one corresponding to the disk consisting of the top two faces of a tetrahedron and the other to the disk consisting of the

bottom two faces.

Proposition 40. *For a fixed boundary $\vec{\lambda}$ of minuscule weights of an n-hive, the* 3*-hives in the disk corresponding to an extroverted triangulation are uniquely determined by the interior edge labels.*

Proof. Each triangle shares an edge with the exterior, which is labeled by a minuscule weight. Hence, the decomposition of the corresponding tensor product is multiplicity free. □

The following proposition holds for any n-hive, without minuscule assumptions.

Proposition 41. *For any extroverted triangulation of the n-gon, the n-hive values on the corresponding disk in the* 2*-skeleton are sufficient to determine the entire n-hive via the octahedron recurrence.*

Proof. Any two faces that share an edge can be viewed as the top or bottom two faces of a tetrahedron in the n-hive, and the octahedron recurrence determines every value in this tetrahedron. It is easy to see that for any extroverted triangulation this process will cover every tetrahedron in the n-hive, and then continue on to fill all of the higher-dimensional subsimplices. □

Remark 42. It is easy to see that there are subsets of the 2-skeleton, even of size $n-2$, that do not generate the entire n-hive in this way. For example, choose the four faces of a tetrahedron in a 6-hive.

Combining the two previous propositions gives the following corollary.

Corollary 43. *Fix an extroverted triangulation τ, and let $\vec{\lambda}$ be a tuple of minuscule weights. The n-hives with boundary $\vec{\lambda}$ are uniquely determined by their values on the interior edges of τ.*

5.3 Bijection between components and hives

By Theorem 39 and the geometric Satake correspondence (Theorem 4), the Littlewood–Richardson coefficient, $c_{\vec{\lambda}}$, is both the number of n-hives with boundary $\vec{\lambda}$ and the number of components of $\mathrm{Poly}(\vec{\lambda})$. To see this bijection directly Kamnitzer, with a suggestion from Speyer, defined in [19] a constructible function,

$$H : \mathrm{Poly}(\vec{\lambda}) \to \mathbb{Z}^{\Delta_m^n},$$

as follows. For each tuple $\vec{i} = (i_1, \ldots, i_n) \in \Delta_m^n$, let $W_{\vec{i}} = V_{\Lambda^{i_1}} \otimes V_{\Lambda^{i_2}} \otimes \cdots \otimes V_{\Lambda^{i_n}}$ where $V_{\Lambda^k} = \Lambda^k \mathbb{C}^m$ is the kth fundamental representation of $G = GL_m$. As a G-representation each $W_{\vec{i}}$ contains a single copy of the one-dimensional determinant representation. Fix a basis vector $\xi_{\vec{i}} \in W_{\vec{i}}$ for this subrepresentation, but now think of $W_{\vec{i}}$ as a GL_m^n representation. To eventually coincide exactly with the definition of hives that we have given from [22] and [19], let $\mathcal{O}' = \mathbb{C}[[t^{-1}]]$ and $\mathcal{K}' = \mathbb{C}((t^{-1}))$. This choice will coincide with taking a maximum in the octahedron recurrence rather than a minimum, as well as $(a_1 - a_0, \ldots, a_m - a_{m-1})$ being weakly decreasing rather than $(a_0 - a_1, \ldots, a_{m-1} - a_m)$ being weakly decreasing. Since G^n acts on $W_{\vec{i}}$, $G^n(\mathcal{K}')$ acts on $W_{\vec{i}} \otimes \mathcal{K}'$, which has a filtration,

$$\cdots \subset W_{\vec{i}} \otimes t^{-2}\mathcal{O}' \subset W_{\vec{i}} \otimes t^{-1}\mathcal{O}' \subset W_{\vec{i}} \otimes \mathcal{O}' \subset W_{\vec{i}} \otimes t\mathcal{O}' \subset W_{\vec{i}} \otimes t^2\mathcal{O}' \subset \cdots .$$

Define the valuation function val: $W_{\vec{i}} \otimes \mathcal{K}' \to \mathbb{Z}$, such that for a vector $v \in W_{\vec{i}} \otimes \mathcal{K}'$, $\mathrm{val}(v) = k$ if $v \in W_{\vec{i}} \otimes t^k\mathcal{O}'$, but $v \notin W_{\vec{i}} \otimes t^{k-1}\mathcal{O}'$. Note that $G^n(\mathcal{O}')$ preserves the

filtration, hence for any $h \in G^n(\mathcal{O}')$ and $v \in W_{\vec{i}} \otimes \mathcal{K}'$, $\mathrm{val}(h \cdot v) = \mathrm{val}(v)$. Thus, for any $([g_0], \ldots, [g_{n-1}]) \in Gr^n = G^n(\mathcal{O}')\backslash G^n(\mathcal{K}')$ we have the action $(g_0, \ldots, g_{n-1}) \cdot \xi_{\vec{i}}$, and a well-defined function $H_{\vec{i}} : Gr^n \to \mathbb{Z}$ given by

$$H_{\vec{i}}(g_0, \ldots, g_{n-1}) = \mathrm{val}((g_0, \ldots, g_{n-1}) \cdot \xi_{\vec{i}}).$$

The function H associates to each component of $\mathrm{Poly}(\vec{\lambda})$ an n-hive. The following theorem of Kamnitzer holds for general $\vec{\lambda}$, not just minuscule.

Theorem 44 ([19])**.** *The generic value of H on each component of $\mathrm{Poly}(\vec{\lambda})$ is an n-hive with boundary $\vec{\lambda}$, and this gives a bijection.*

The following proposition is an easy combination of [8] and [19]. First, observe that in the definition of $H_{\vec{i}}$, $\xi_{\vec{i}}$ is an eigenvector for the diagonal $G(\mathcal{O}')^\Delta$ in $G^n(\mathcal{K}')$ with eigenvalue $\det : G(\mathcal{O}') \to (\mathcal{O}')^\times$. Then for $h \in G(\mathcal{O}')^\Delta$,

$$H_{\vec{i}}([gh]) = \mathrm{val}(gh \cdot \xi_{\vec{i}}) = \mathrm{val}(g \cdot \det(h)\xi_{\vec{i}}) = \mathrm{val}(\det(h)g \cdot \xi_{\vec{i}}) = \mathrm{val}(g \cdot \xi_{\vec{i}}) = H_{\vec{i}}([g]).$$

Hence, H is $G(\mathcal{O}')^\Delta$ invariant.

Proposition 45. *Fix a component X of $\mathrm{Poly}(\vec{\lambda})$ and its corresponding n-hive F given by the function H. The 1-skeleton of F is labeled by the generic distances of X.*

Proof. Let $([g_0], [g_1], \ldots, [g_{n-1}])$ be generic in X, and consider $[g_j]$ and $[g_k]$. The j vertex of F is $(0, \ldots, 0, m, 0, \ldots, 0)$ where the m is in the j position, and likewise the k vertex of F has the m in the k position. The values of H along the edge connecting the j and k vertices depend only on $[g_j]$ and $[g_k]$ because the representations in this case are $W_a = 1 \otimes \cdots \otimes \Lambda^{m-a} \otimes 1 \otimes \cdots \otimes 1 \otimes \Lambda^a \otimes \cdots \otimes 1$ with Λ^{m-a} in the j position and Λ^a in the k position with the trivial representation in all other positions. Hence, for $0 \le a \le m$, let $H_a([g_j], [g_k])$ denote the value of H on the ath lattice point on the edge connecting vertex j and k.

There exists an $h \in G(\mathcal{O}')^{\Delta}$, such that $[g_j h] = [t^{\alpha}]$ and $[g_k h] = [g'_k]$. Note that $d([g_j], [g_k]) = d([t^{\alpha}], [g'_k])$ holds. Using the $G(\mathcal{O}')^{\Delta}$ invariance, $H_a([g_j], [g_k]) = H_a([t^{\alpha}], [g'_k])$. The set of points $([t^{\alpha}], [g'_k])$ is in bijection with the set of points $([t^0], [g'_k t^{-\alpha}]) = ([t^0], [g''_k])$, and note that $d([t^{\alpha}], [g'_k]) = d([t^0], [g'_k t^{-\alpha}])$. Also, $H_a([t^{\alpha}], [g'_k]) = |\alpha| + H_a([t^0], [g''_k])$. Suppose $[g''_k] \in Gr(\lambda)$, so $d(t^0, [g''_k]) = \lambda$. To calculate $H_a([t^0], [g''_k])$, use the $G(\mathcal{O}')^{\Delta}$ invariance again, so that $H_a([t^0], [g''_k]) = H_a([t^0], [t^{\lambda}])$.

Let $\xi_a \in W_a$ be the vector defined above. Finally, compute that $(t^0, t^{\lambda}) \cdot \xi_a \in W_a \otimes t^{\lambda_1 + \cdots + \lambda_a} \mathcal{O}'$, hence $H_a([t^0], [t^{\lambda}]) = \lambda_1 + \cdots + \lambda_a$. Hence,

$$\begin{aligned} H_a\left([t^{\alpha}], [g'_k]\right) - H_{a-1}\left([t^{\alpha}], [g'_k]\right) &= |\alpha| + H_a\left([t^0], [g''_k]\right) - |\alpha| - H_{a-1}\left(([t^0], [g''_k]\right) \\ &= \sum_{i=1}^{a} \lambda_i - \sum_{i=1}^{a-1} \lambda_i \\ &= \lambda_a, \end{aligned}$$

as desired. □

Thus, the generic distances of a component can be understood by studying the combinatorics of the 1-skeleton of the associated n-hive. Recall that the local rule of Theorem 12 implements a quadrilateral flip as in Figure 3.4. The quadrilateral has two external edges (labeled by minuscule weights) and two internal edges of the n-gon. These four edges, together with the two diagonals of the quadrilateral form the edges of a tetrahedron, i.e. a 4-hive, inside of the n-hive. See Figure 5.2 for the top view of a 4-hive balanced on an edge with edges labeled by corresponding weights such that the visible weights are read west to east, and ρ is read north to south. Hence, to prove Theorem 12, we need to understand the excavation of a single 4-hive when two of the external edges are minuscule. In other words, given λ, μ, ν, j, k as in the setup of Figure 5.2,

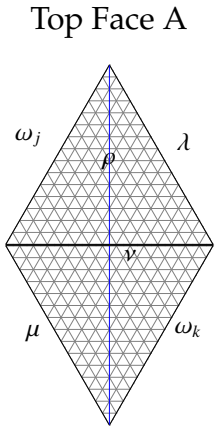

Figure 5.2: A transparent top view of a 4-hive

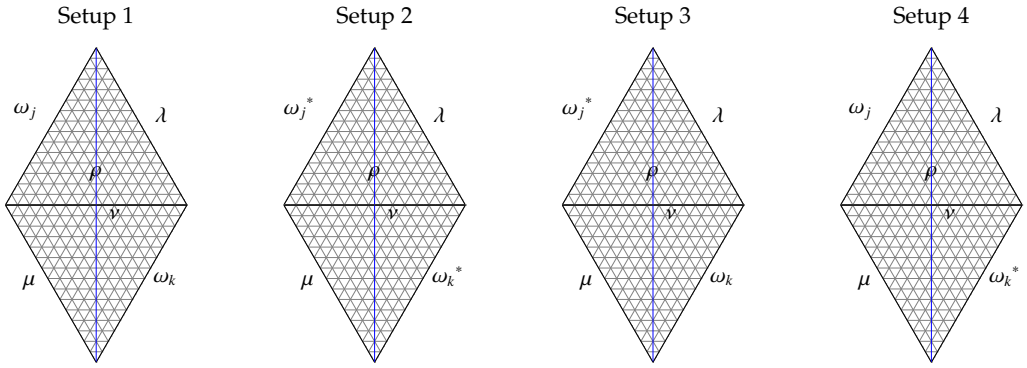

Figure 5.3: The four possible orientations

determine ρ. Let N, E, S, W denote the compass directions. We will keep this orientation of the 4-hive fixed, so that the closer faces are called the top Face A and the top Face B, while the farther faces are called the left bottom face and the right bottom face. The four possible setups for the two minuscule weights when read west to east are listed in figure 5.3, however we will focus on setup 1.

5.4 Proof

In this section, we prove Theorem 12. We will focus on the hive setup in figure 5.2, with the others following similarly.

Recall that a hive is labeled at the lattice points by integers, and the differences of the lattice points along the boundary edges give the dominant weights. In our case, the NW edge has differences ω_j when read from SW to NE, so the integer labels along this edge increase by one j times and are constant thereafter. Restrict to the Face A 3-hive. A lemma in [22] gives an easy description of the differences along the strip one step in from this NW external edge. It is restated here in terms of our Figure 5.2.

Lemma 46 ([22, Lem. 2])**.** *If the NW external edge of a 3-hive has differences ω_j, then the strip one step in from the NW edge has differences ω_j or ω_{j-1}. Furthermore, the first case occurs if and only if $\nu_1 = \lambda_1$ and the second case if and only if $\nu_1 = \lambda_1 + 1$.*

Let the second strip be ω_{j_2} where j_2 must be j or $j-1$. Since this strip can be viewed as the NW external edge of a smaller 3-hive contained in Face A, the third strip must have differences ω_{j_2} or ω_{j_2-1}. This pattern continues so that there are j decreases from the first strip to the $n+1$st strip, which is just a single lattice point (the east corner). A decrement occurs between strip i and $i+1$ if and only if $\nu_i = \lambda_i + 1$, i.e. a box is added to λ in the ith row. The left of Figure 5.4 shows the top of a 4-hive with Face A such that $n = 5$, $j = 3$, $\lambda = (3, 3, 1, 1, 1)$, $\nu = (4, 3, 2, 2, 1)$. The blue path indicates the lattice points after which the strips are constant. We call this the **break path**. By the previous lemma, the break path determines the skew shape ν/λ.

Likewise, the top Face B of the 4-hive has ω_k along its SE edge with the ω_k

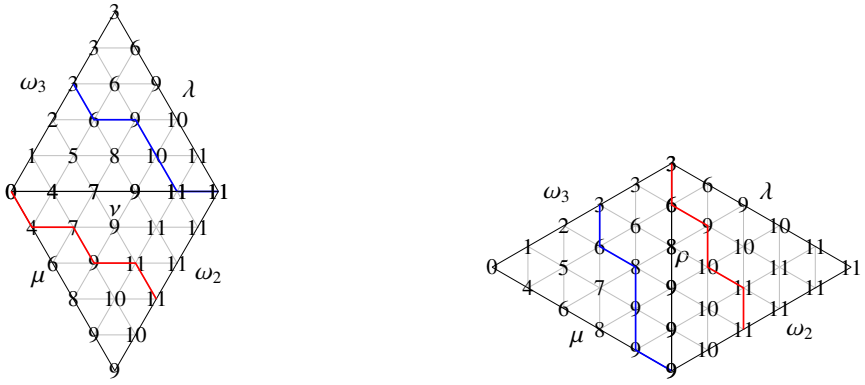

Figure 5.4: The top and bottom of a 4-hive with break paths indicated

being read from the SW to the NE. The strip one step in from this boundary edge is either ω_k or ω_{k-1}, and so forth with k decrements between the first strip and the $n+1$st strip (the west corner). The break path of Face B connects the SE external edge to the W lattice point. In order for the numbering of the strips to coincide with top Face A, we can view the break path of Face B as starting at the W lattice point with the first strip being ω_0, so that there are k increases before reaching ω_k on the SE edge. As with Face A, there is an increase between strip i and $i+1$ if and only if $\nu_i = \mu_i + 1$, and no increase if and only if $\nu_i = \mu_i$. Thus, ν/μ is determined by the B break path. The left of Figure 5.4 shows the top view of a 4-hive with Face B such that $k = 2$, $\mu = (4, 2, 2, 1, 0)$. Note that a horizontal step in the break path of Face B corresponds to an increase, while a slanted step in the break path of A corresponds to a decrease.

The right of Figure 5.4 shows the bottom two faces of a 4-hive with top two faces given on the left. As with the top faces, μ together with the decreases of the bottom left face's break path determine ρ. More specifically, $\rho_i = \mu_i$ if there is no decrease between the i and $i+1$ strip and $\rho_i = \mu_i + 1$ if there is a decrease. We will use the octahedron recurrence to excavate individual lattice points to reveal these decreases, and hence determine ρ.

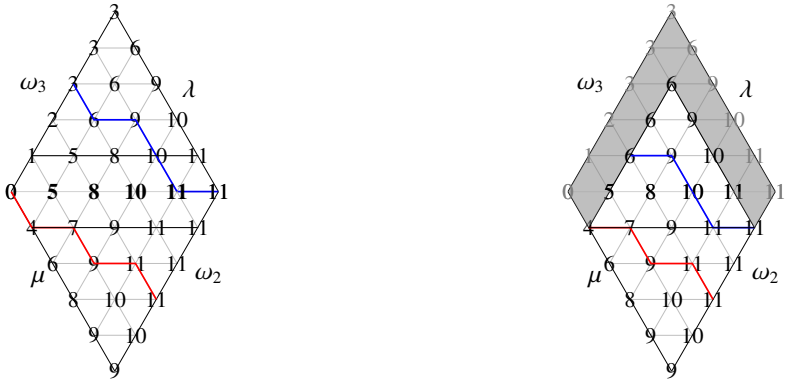

Figure 5.5: Excavating a 4-hive

The bottom left 3-hive has the same setup as top Face A with ρ and μ playing the roles of λ and ν respectively. Hence the decreases in the break path are the rows where a box is added to ρ to get μ. But since the goal is to determine ρ, we view the decreases in the break path as corresponding to the positions of μ where a box was subtracted to get ρ. Note that the decreases are the vertical steps in the break path. As stated previously, since we know μ, if we know the ith strip then all we need to know is the $i+1$ strip to determine ρ_i.

Notice that the second strip of the bottom left face is found by excavating the second strip of Face A just once. Thus to reveal the second strip of the bottom face, and equivalently determine ρ_1, it suffices to shave off Face A by a depth of 1. This leaves a smaller tetrahedron, whose main horizontal edge visible from the top is one row down on Face B from the original main horizontal edge. The right of Figure 5.5 demonstrates this first shave step for our previous example with the new tetrahedron indicated. At this point the first two strips of the bottom left face are known, so in particular $\rho_1 = 3$ can be seen either by directly reading off the difference $6 - 3$, or by noting that there is a decrease from ω_3 to ω_2 between the the first and second strips and that $\mu_1 = 4$.

The newly exposed, smaller tetrahedron has a new top Face B, whose values were the previously exposed values of the original top Face B, and a new top Face A, whose values were revealed. To get ρ_2, or equivalently the third strip of the bottom left face, the top Face A of this new tetrahedron is shaved off and so forth. We call such a step in our process a **shave step** to differentiate from a single application of the octahedron recurrence at a specific lattice point. We already observed that every strip of the original top Face A is an ω_i for some i. The same reasoning applies to the top Face A of the new tetrahedron, and induction gives the following lemma.

Lemma 47. *Every strip of a newly exposed top Face A has differences ω_i for some i.*

Proof. The NW external edge of each new top Face A is a strip of the bottom left face, which we know by Lemma 46 has differences ω_i for some i. Thus, by Lemma 46 again each strip in this face has differences ω_j for some $j \leq i$. □

To get each ρ_i we must understand which ω_i appear after a single shave step, or equivalently, how the break path changes after a shave step. Since the new top Face B is obtained simply by ignoring the previous main horizontal, it is easy to see that the break path contracts along the slanted step from the main horizontal to the horizontal one row down. If there is no such slanted step, then the break path must itself be the main horizontal in which case the new break path is the new main horizontal. Analyzing the change in the break path of top Face A requires looking at a number of cases. The goal is the following proposition.

Proposition 48. *After a single shave step, the strip underneath a strip labeled by ω_i is labeled by either ω_i or ω_{i+1}. The ω_{i+1} case occurs only if $\mu_k = \mu_{k+1}$ where k is the strip in*

consideration. Furthermore, there can be at most one interval $[a, b]$ *of strip indices in which the strip label* ω_i *changes to* ω_{i+1}.

Consider shaving the original 4-hive. A single shave step is built out of inductive steps, each consisting of excavating a horizontal row (first the main horizontal, then the one above, and so forth). The base case of this induction is the excavation of the main horizontal of the original 4-hive, as in the left of Figure 5.5.

By the octahedron recurrence, excavating a single position depends on the value of the position itself and the four values that are directly NE, SE, SW, NW of the position in question. However, in analyzing the different cases arising from different positions on the break path, it will be easier to keep track of two extra points. The cases will depend on values of the following relative positions, where we are excavating the middle position of the middle row. The values of these relative positions will be called the **local data**.

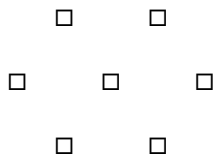

5.4.1 Excavating below the break path

Consider first excavating points strictly below the break path of Face A. Even more specifically, begin with a point on the main horizontal. There are three subcases depending on the break path of Face B, as shown in Figure 5.6. If this

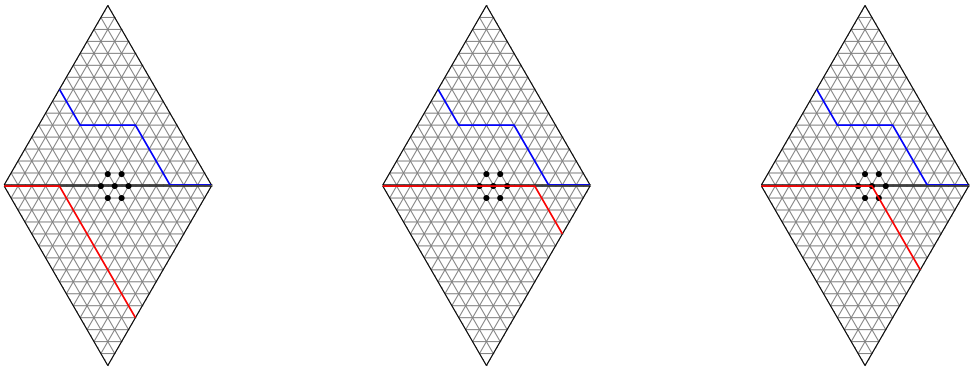

Figure 5.6: Case A (main horizontal)

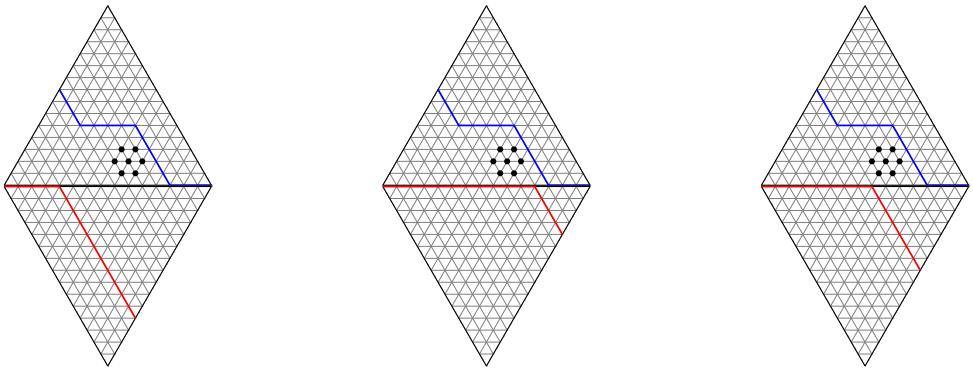

Figure 5.7: Case A

is the $k+1$ strip, then the local data looks like the following.

$$\begin{array}{ccc} \sum_1^{k-1} \nu_i + 1 & & \sum_1^{k} \nu_i + 1 \\ \sum_1^{k-1} \nu_i & \sum_1^{k} \nu_i & \sum_1^{k+1} \nu_i \\ \sum_1^{k} \nu_i - x & & \sum_1^{k+1} \nu_i - y \end{array}$$

Since we are excavating the main horizontal, the bottom row positions are on Face B, and the x, y variables are 0 or 1 depending on the break path of Face B. Only the values of x and y will affect excavation at this point, so the possible cases are $x = 0$ and $y = 0$, or $x = 1$ and $y = 1$, or $x = 1$ and $y = 0$, as shown in Figure 5.6.

Let z be the value of the lattice point that the current excavation will reveal.

The value of z must be either $\sum_1^k \nu_i - x$ or $\sum_1^k \nu_i - x + 1$ because every strip of the new face will be an ω_i for some i, and the value of the first lattice point of this strip is $\sum_1^k \nu_i - x$. In each case we show that z is actually $\sum_1^k \nu_i - x + 1$, hence there is a difference of 1 between the first and second positions of the revealed strip.

(i) $(x = 0, y = 0)$

$z = \max\left(\sum_1^k \nu_i + 1, \sum_1^{k-1} \nu_i + 1 + \nu_{k+1}\right) = \sum_1^k \nu_i + 1$ because $\nu_{k+1} \leq \nu_k$.

(ii) $(x = 1, y = 1)$

$z = \max\left(\sum_1^k \nu_i, \sum_1^{k-1} \nu_i + \nu_{k+1}\right) = \sum_1^k \nu_i$ because $\nu_{k+1} \leq \nu_k$.

(iii) $(x = 1, y = 0)$

$z = \max\left(\sum_1^k \nu_i, \sum_1^{k-1} \nu_i + \nu_{k+1} + 1\right) = \sum_1^k \nu_i$.

In this case, the second term in the maximum will be $z = \sum_1^k \nu_i + 1$ if $\nu_{k+1} = \nu_k$, but this is impossible because the newly excavated line has to be an ω_i for some i and these values give a difference of 2. Or we can see directly that $x = 1, y = 0$ and $\nu_k = \nu_{k+1}$ imply that $\mu_k + 1 = \nu_k = \nu_{k+1} = \mu_{k+1}$ which contradicts μ being dominant.

Excavating the main horizontal leaves a trench, where the newly visible lattice points now play the role of the bottom row of the 7 local positions. We call these the **trench values**. In the case just analyzed of excavating the portion of the main horizontal that is below Face A's break path, the trench values were the values on Face B on the horizontal one down from the main horizontal. In all three cases above, the **trench difference** x between the trench and the excavation point is preserved, an observation we will need for the inductive step.

Now consider excavating a point that is on the $k + 1$ strip and on the cth line above the main horizontal, but still strictly below the break path of Face

A. Assume that the two relevant lattice points below this location have already been excavated. The local data looks like the following.

$$\begin{array}{ccccc} & \sum_1^{k-1} \nu_i + c + 1 & & \sum_1^{k} \nu_i + c + 1 & \\ \sum_1^{k-1} \nu_i + c & & \sum_1^{k} \nu_i + c & & \sum_1^{k+1} \nu_i + c \\ & \sum_1^{k} \nu_i + c - x & & \sum_1^{k+1} \nu_i + c - y & \end{array}$$

By induction the exposed portion of the new $k+1$ strip increases by one at every step thus far, and the trench differences of x and y in the local data are the same as the trench differences given by the break path of Face B. In particular, the same three subcases may arise, as indicated in Figure 5.7. Since the constant c does not affect the octahedron recurrence, the calculation of z is as before, so we may conclude in this case that both the strip difference and the trench difference propagate. In particular, note that the trench differences given by the break path of Face B propagate all the way up to excavating the points directly on the A break path.

5.4.2 Excavating on the break path

Now consider excavating a point on the break path of A. Consider case B of Figure 5.8. Since the two positions of the bottom row of the local data were excavated by instances of case A, the trench differences x, y take on the same values originally given by the break path of Face B. In particular, the same three possible subcases can arise. The local data is as follows where the constant c is omitted.

$$\begin{array}{ccccc} & \sum_1^{k-1} \nu_i & & \sum_1^{k} \nu_i & \\ \sum_1^{k-1} \nu_i & & \sum_1^{k} \nu_i & & \sum_1^{k+1} \nu_i \\ & \sum_1^{k} \nu_i - x & & \sum_1^{k+1} \nu_i - y & \end{array}$$

(i) $(x = 0, y = 0)$

$z = \max\left(\sum_1^k \nu_i, \sum_1^{k-1} \nu_i + \nu_{k+1}\right) = \sum_1^k \nu_i$ because $\nu_{k+1} \leq \nu_k$.

(ii) $(x = 1, y = 1)$

$z = \max\left(\sum_1^k \nu_i - 1, \sum_1^{k-1} \nu_i + \nu_{k+1} - 1\right) = \sum_1^k \nu_i - 1$ because $\nu_{k+1} \leq \nu_k$.

(iii) $(x = 1, y = 0)$

$z = \max\left(\sum_1^k \nu_i - 1, \sum_1^{k-1} \nu_i + \nu_{k+1}\right) = \sum_1^k \nu_i - 1$

In this case the second term of the maximum is $z = \sum_1^k \nu_i$ if $\nu_{k+1} = \nu_k$, but as before this is impossible: The values $x = 1$ and $y = 0$ are given by the break path of B, hence imply that $\nu_k = \mu_k + 1$ and $\nu_{k+1} = \mu_{k+1}$. Together with $\nu_{k+1} = \nu_k$ this would imply, $\mu_{k+1} = \nu_{k+1} = \nu_k = \mu_k + 1$, contradicting μ dominant. In every subcase, the strip equality and the trench difference propagate to the strip below.

At this point of excavating a strip that falls under case B, the revealed strip has c increasing steps (which we knew from case A), followed by a constant step.

Now consider case C, the interesting case where the excavation point is at an **out elbow** in the break path of A. In this case, the break path of A says that to obtain ν from λ a box was added to λ_k but not to λ_{k+1}. Then $\nu_k \neq \nu_{k+1}$ because otherwise this would imply $\lambda_k < \lambda_{k+1}$ contradicting that λ is dominant. Hence, for case C, $\nu_{k+1} \leq \nu_k - 1$. The local data is as follows.

$$\begin{array}{ccc} \sum_1^{k-1} \nu_i + 1 & & \sum_1^k \nu_i \\ \sum_1^{k-1} \nu_i & \sum_1^k \nu_i & \sum_1^{k+1} \nu_i \\ \sum_1^k \nu_i - x & & \sum_1^{k+1} \nu_i - y \end{array}$$

As before, the trench differences x and y are determined by instances of case A, hence were propagated from the break path of B. The three subcases are as

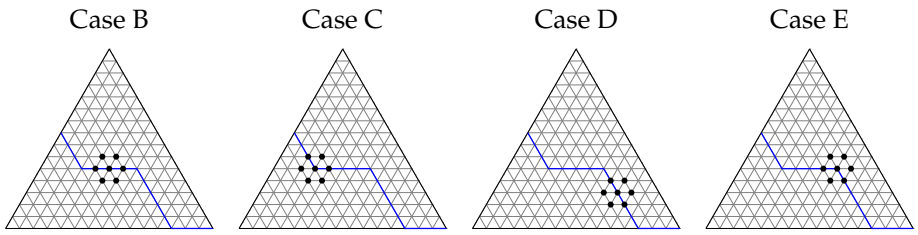

Figure 5.8: All of the "on break path" cases

follows.

(i) $(x = 0, y = 0)$

$z = \max\left(\sum_1^k \nu_i, \sum_1^{k-1} \nu_i + \nu_{k+1} + 1\right) = \sum_1^k \nu_i$ where we have used the fact that $\nu_{k+1} \leq \nu_k - 1$.

(ii) $(x = 1, y = 1)$

$z = \max\left(\sum_1^k \nu_i - 1, \sum_1^{k-1} \nu_i + \nu_{k+1}\right) = \sum_1^k \nu_i - 1$ where we have used the fact that $\nu_{k+1} \leq \nu_k - 1$.

(iii) $(x = 1, y = 0)$

$z = \max\left(\sum_1^k \nu_i - 1, \sum_1^{k-1} \nu_i + \nu_{k+1}\right)$

In this case, $z = \sum_1^k \nu_i - 1$ if $\nu_{k+1} < \nu_k - 1$ and $z = \sum_1^k \nu_i$ if $\nu_{k+1} = \nu_k - 1$, which can actually occur. The second case corresponds to a difference of 1 between the current two positions of the revealed strip, even though there was a difference of 0 between the two positions in the strip above. Also in this special case the trench difference $x = 1$ turns into $x = 0$.

Furthermore, if the Face B break path gives $x = 1$, $y = 0$, then $\nu_{k+1} = \nu_k - 1$ if and only if $\mu_k = \mu_{k+1}$. This is because $x = 1$ implies $\nu_k = \mu_k + 1$ and $y = 0$ implies $\nu_{k+1} = \mu_{k+1}$.

Hence, excavation of a case C strip has thus far revealed c increasing steps,

followed possibly by another increasing step if and only if the B break path has an out elbow at the same strip and $\mu_k = \mu_{k+1}$.

Now consider case D, the other interesting case, with local data is as follows.

$$\begin{array}{ccc} & \sum_1^{k-1} \nu_i + 2 & \sum_1^{k} \nu_i + 1 \\ \sum_1^{k-1} \nu_i + 1 & \sum_1^{k} \nu_i + 1 & \sum_1^{k+1} \nu_i \\ & \sum_1^{k} \nu_i + 1 - x & \sum_1^{k+1} \nu_i - y \end{array}$$

As before, the position containing x was excavated by an instance of case A, so this trench value propagated from the original trench value given by break path B. The position containing y however could have been excavated by an instance of C, or by another instance of D. If it was excavated by C, then it is possible that the special case of C changed a trench difference of 1 to a 0. Either way, we have the same three subcases for x, y values.

(i) $(x = 0, y = 0)$

$z = \max\left(\sum_1^k \nu_i + 1, \sum_1^{k-1} \nu_i + \nu_{k+1} + 1\right) = \sum_1^k \nu_i + 1$

(ii) $(x = 1, y = 1)$

$z = \max\left(\sum_1^k \nu_i, \sum_1^{k-1} \nu_i + \nu_{k+1}\right) = \sum_1^k \nu_i$

(iii) $(x = 1, y = 0)$

$z = \max\left(\sum_1^k \nu_i, \sum_1^{k-1} \nu_i + \nu_{k+1} + 1\right)$

In case (iii), $z = \sum_1^k \nu_i$ if $\nu_{k+1} < \nu_k$ and $z = \sum_1^k \nu_i + 1$ if $\nu_{k+1} = \nu_k$, which can actually occur. The second case corresponds to an increase in the current position of the revealed strip, even though there was no such increase in the original strip. Also in this second case the trench difference $x = 1$ turns into $x = 0$. Let's consider when this special case can actually happen. One way to end up with $x = 1$ and $y = 0$ is if the break path of the top Face B starts by giving $x = 1$ and $y = 0$ at

the very first step. Then $\nu_{k+1} = \nu_k$ forces $\mu_k = \nu_k - 1 < \nu_k = \nu_{k+1} = \mu_{k+1}$ which is a contradiction. The only other way to end up with $x = 1$ and $y = 0$ is if the break path of top Face B starts by giving $x = 1$ and $y = 1$ on these strips, and the excavation of the position directly to the right changes $y = 1$ into $y = 0$ as in the special case of C(iii). In this case, $\nu_k = \nu_{k+1}$ if and only if $\mu_k = \mu_{k+1}$.

If this special case of D(iii) occurs then the strip difference of 0 becomes a difference of 1, and the trench difference of $x = 1$ turns into a trench difference of $x = 0$. By induction, an instance of case D can only have the above three possible values for x and y. In particular, the special case of D(iii) can trigger another instance of the special case of D(iii) directly to the NW. Such a string of special case D(iii)'s can only occur if the special case C(iii) occurs at the out elbow of the same NW-SE strip.

Finally, for case E the position involving y could have been excavated by an instance of case C or D, in which case we know that $y = 0$ or $y = 1$. We leave the routine analysis to the reader.

This concludes the cases of excavating a point on the break path. At this point we have shown that strip difference and trench difference always propagate when below the break path, and also when on the break path, except for the special cases of C(iii) and D(iii). If the special cases of C(iii) and D(iii) occur then this implies $\mu_k = \mu_{k+1}$. In both special cases, the strip difference of 0 flips to a 1, while the trench difference of 1 flips to a 0.

There are four remaining cases of excavating a point above the break path. We claim that there are no new interesting cases, i.e. the strip difference and trench difference propagate as they did below the break path, and we leave this

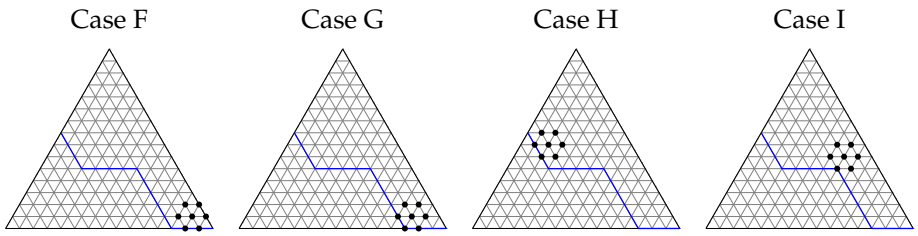

Figure 5.9: All of the "above break path" cases

analysis to the reader. Note that case F and G must be analyzed before case H, likewise H before I, because the point containing y must have been excavated by one of these instances.

This completes the proof of the proposition. In fact, we have shown more.

Proposition 49 (Proposition 48 restated)**.** *If the break path B doesn't contain a proper portion of the main horizontal, then each strip is labeled by the same fundamental weight as the one beneath it. Otherwise, if the break path of B is horizontal exactly up to the $k+1$st strip, then this could trigger the special case of C(iii) and a sequence of special cases of D(iii). In each case $\mu_j = \mu_{j+1}$ is necessary. For this interval of strips $[i, k+1]$, each fundamental weight is incremented by one upon excavation. All other strips remain the same upon excavation.*

The only assumptions that we have made on the original 4-hive is that the NW and SE edges are fundamental weights and that λ, μ, ν are dominant weights. After one shave step we see that these assumptions still hold, so the same applies to the smaller 4-hive. Note though that the break paths change. Each shave step can increment the fundamental weights of at most one nonempty interval of strips.

5.4.3 Completing the proof

Recall that a box is added to λ_k to get ν_k if and only if the kth strip of Face A is labeled by ω_i and the $k+1$st by ω_{i-1}. Likewise, a box is removed from μ_k to get ρ_k if and only if the kth strip of the bottom left face is labeled by ω_i and the $k+1$st by ω_{i-1}. This information is revealed on the kth shaving of the hive.

For example, let $k = 1$. In order to discover whether a box needs to be removed from μ_1 to get ρ_1 we need to know whether or not there is a decrement between the first and second strips of the bottom face. But the second strip of the bottom face is directly underneath the second strip of the top face, so it is exposed after a single shaving. See Figures 5.4 and 5.5 for a reminder. Therefore, if a special case does not arise along the second strip, then a nondecrease between the first two strips of the top face gets copied to a nondecrease of the first two strips of the bottom face, which corresponds to no box being added to λ_1 and no box being removed from μ_1. Also if a special case does not arise along the second strip, then a decrease between the first two strips of the top face gets copied to a decrease of the first two strips of the bottom face, and this corresponds to a box being added to λ_1 and a box being removed from μ_1. The next shaving reveals the third line of the bottom face and so forth.

Thus, if $\mu_k > \mu_{k+1}$ for all k, then by the proposition the special cases can never occur, so each shave simply copies the ω_i from the top face to the next. This means that a box is added to λ_k to get ν_k if and only if a box is removed from μ_k to get ρ_k. On the other hand, if $\mu_k = \mu_{k+1}$ and a box is added to λ_k to get ν_k, then removing a box from μ_k would result in $\mu_k < \mu_{k+1}$, i.e. not a dominant weight. Instead, sorting has to occur in order to remove a box from the first μ_j where this is possible. The special cases of C(iii) and D(iii) facilitate the sorting, as we

will see in the proof of the following main theorem.

Theorem 50. *Consider the setup of the 4-hive in Figure 5.2 and the j boxes added to λ in distinct rows to get ν. Then subtracting a box from the same rows of μ and sorting gives the dominant weight ρ.*

Proof. For a dominant weight μ let $i_k < j_k$ be the indices denoting the stretches of equalities of the parts of μ,

$$\mu_1 > \mu_2 > \cdots > \mu_{i_1} = \mu_{i_1+1} = \cdots = \mu_{j_1} > \mu_{j_1+1} > \cdots > \mu_{i_2} = \mu_{i_2+1} = \cdots = \mu_{j_2} > \cdots > \mu_n.$$

Note it is possible that $i_{k+1} = j_k + 1$. Recall that the strips of the 4-hive are indexed $1, \ldots, n+1$ where the last strip is always a single lattice point, while the parts of the dominant weights are indexed $1, \ldots, n$ and are given by the differences of the labels of the $1, \ldots, n+1$ lattice points.

Let $\alpha \notin [i_k, j_k]$ for all k, which is equivalent to $\mu_\alpha \neq \mu_{\alpha+1}$ and $\mu_\alpha \neq \mu_{\alpha-1}$. Suppose a box is not added to the α part of λ, so that the break path step between the α and $\alpha+1$ strips is horizontal. The only way the step could end up being slanted after shaving is if one of the special cases occurs at strip α, but this is impossible because $\mu_{\alpha-1} \neq \mu_\alpha$. Thus a box is not removed from the α part of μ to give the α part of ρ.

Suppose a box is added to the α part of λ, so that the break path step between the α and $\alpha+1$ strips is slanted. The only way the step could end up being horizontal after shaving is if one of the special cases occurs at strip $\alpha+1$, but again this is impossible because $\mu_\alpha \neq \mu_{\alpha+1}$. Thus a box is removed from the α part of μ to give the α part of ρ.

Now consider the entire interval $[i_k, j_k]$ for some k, and restrict ν and μ to this interval. Then to get ν from μ the only possibility is that boxes were added

to some amount of consecutive parts $[i_k, h_k]$ of μ, where $h_k \leq j_k$, followed by no additions thereafter. Then the break path of Face B on the interval $[i_k, j_k]$ is horizontal for the steps between strips i_k to h_k and slanted between strips h_k to j_k. See Figure 5.10 where the strips i_k to j_k are bold.

If for every $\alpha \in [i_k, j_k]$ a box is not added to the α part of λ, then the break path of Face A across these gaps is horizontal. Then the special cases cannot occur at any of the strips with index in $[i_k, j_k]$. Thus the strips remain the same after every shave, so that no box is removed from μ to give ρ.

Now suppose that there is at least one index in $[i_k, j_k]$ such that the corresponding part of λ has a box added to it. We would like to understand the possibilities for the break path of Face A. On the interval $[i_k, j_k]$, the general μ, ν, and λ look like the following where we have omitted a fixed number of initial columns on the left.

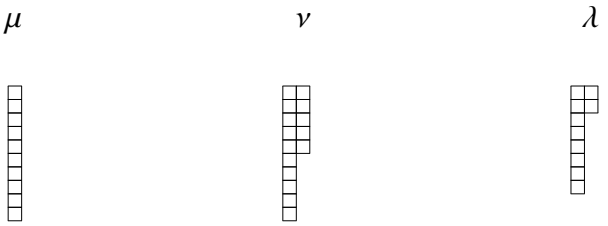

Note that if μ_i is negative then we get the following picture where we omitted a number of initial "negative" columns on the right.

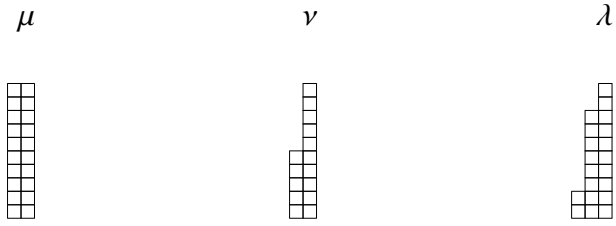

This corresponds to a portion of Face A's break path that starts with consecutive horizontal steps, then a number of consecutive slanted steps, followed by

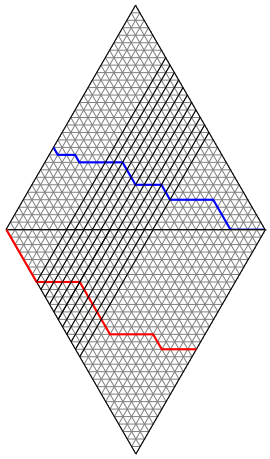

Figure 5.10: Break paths

horizontal steps, and finally slanted steps. It is possible that any one of these four consecutive segments is empty, but in any case there is at most one spot where a slanted step is followed by a horizontal step, i.e. there is at most one out elbow. Furthermore, if the break path of Face A does have an out elbow at say strip k, then the break path of B must also have an out elbow at strip k because of the values of μ, ν, λ on this interval. Thus $[i_k, j_k]$ contains two subintervals, $[\alpha_1, \beta_1]$ and $[\alpha_2, \beta_2]$, that consist of indices for which the corresponding part of λ has a box added to it. See Figure 5.10.

Let the ith shave step excavate Face A_i to reveal the smaller Face A_{i+1}. This eliminates the ith strip, so Face A_{i+1} consists of the strips indexed by $i+1$ to $n+1$. In particular, the ith shave reveals the ith step of the break path of the bottom left face and hence determines ρ_i from μ_i. If there is a slant between the ith and $i+1$st strips of Face A_i and the $i+1$st strip is not increased during the shaving then a box is removed from μ_i to get ρ_i. If this happened for an $i \in [\alpha_1, \beta_1]$, then μ would yield a nondominant ρ. We will show that this doesn't happen because the special cases of the shave steps perform the desired sorting.

Note that at the i_kth shave step, the break path of B_{i_k} is horizontal until strip β_1 and slanted at least until j_k. For each $i \in [i_k, \alpha_1]$, the ith step in A's break path is horizontal, so the ith shave reveals that $\rho_i = \mu_i$. Consider the α_1 shave step. Since there is an out elbow at β_1 in the break paths of both A_i and B_i and $\mu_i = \mu_{i+1}$, the special case of C(iii) occurs and triggers a string of special cases of D(iii) in each of the strips α_i to $\beta_i - 1$. Thus, although there was a slant between the α_i and $\alpha_i + 1$ strips, the α_i shave reveals that $\rho_i = \mu_i$. The break path of A_{α_1+1} has the same number of slanted steps, from $\alpha_1 + 1$ to $\beta_1 + 1$ and α_2 to β_2, but one fewer horizontal step, from $\beta_1 + 1$ to α_2. The break paths of A_{α_1+1} and B_{α_1+1} now have out elbows at strip $\beta_1 + 1$, so the same process occurs at the $\alpha_1 + 1$ shave step giving $\rho_i = \mu_i$. This continues for $\alpha_2 - \beta_1$ steps with the slanted intervals of the A break path moving closer together each time, until they finally coalesce. At this point the break path of A consists of $\beta_1 - \alpha_1 + \beta_2 - \alpha_2$ slanted steps between the strip indices $[\alpha_1 + \alpha_2 - \beta_1, \beta_2]$.

At this point the special case of C(iii) cannot occur because the break path of B cannot have an out elbow at the β_2 strip. Thus, each remaining shave step up until β_2 reveals $\rho_i = \mu_i - 1$. The end result is that on the interval $[i_k, j_k]$, μ is decremented in the final $\beta_2 - \alpha_2 + \beta_1 - \alpha_1$ positions to yield a dominant ρ. □

The bottom of the 4-hive looks like setup 2 of Figure 5.3 where the edge labels are oriented from W to E as before. Since this is the same setup as in our proofs, $\nu^* = \text{sort}(\lambda^* + \mu^* - \rho^*)$, which is equivalent to $\nu = \text{sort}(\lambda + \mu - \rho)$.

There are however, two more orientations of the external edge labels that can arise, setup 3 and setup 4 in Figure 5.3. They are the top and bottom of the same hive. It is not hard to see that if $\rho = \text{sort}(\lambda + \mu - \nu)$, then going back is given by the same formula $\nu = \text{sort}(\lambda + \mu - \rho)$. Thus, we need only consider setup 3.

In this case, the NW strip of Face A is constant for some time followed by j decreases. Lemma 46 applies in an analogous way where a horizontal step in the break path of A corresponds to a decrease from ω_i^* to ω_{i-1}^*. The cases that arise can be analyzed similarly with a special case analogous to C(iii) at an out elbow that can trigger a sequence of special cases analogous to D(iii). These occur only if $\mu_k = \mu_{k+1}$, which allows for sorting as in the proof of the main theorem. Before, k boxes are added to μ to get ν, whereas in this case, j boxes are removed from λ to get ν. The relevant picture on an interval $[i_k, j_k]$ is now the following.

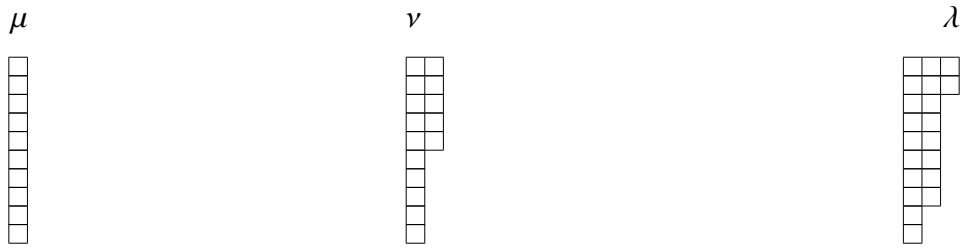

The final steps of the proof proceed similarly, and the details are left to the reader.

BIBLIOGRAPHY

[1] Jared Anderson and Mikhail Kogan. Mirković-Vilonen cycles and polytopes in type A. *International Mathematics Research Notices*, 2004(12):561–591, 2004.

[2] Alexander A. Beilinson and Vladimir G. Drinfeld. Quantization of Hitchin's fibration and Langlands' program. In *Algebraic and geometric methods in mathematical physics (Kaciveli, 1993)*, volume 19 of *Math. Phys. Stud.*, pages 3–7. Kluwer Acad. Publ., Dordrecht, 1996.

[3] Thomas Britz and Sergey Fomin. Finite posets and Ferrers shapes. *Advances in Mathematics*, 158(1):86 – 127, 2001.

[4] Michael Chmutov, Max Glick, and Pavlo Pylyavskyy. The Berenstein–Kirillov group and cactus groups. arXiv:1609.02046, 2017.

[5] Sergey V. Fomin. Finite, partially ordered sets and Young diagrams. *Dokl. Akad. Nauk SSSR*, 243(5):1144–1147, 1978.

[6] Sergey V. Fomin. The generalized Robinson-Schensted-Knuth correspondence. *Zap. Nauchn. Sem. Leningrad. Otdel. Mat. Inst. Steklov. (LOMI)*, 155(Differentsial′naya Geometriya, Gruppy Li i Mekh. VIII):156–175, 195, 1986.

[7] Bruce Fontaine and Joel Kamnitzer. Cyclic sieving, rotation, and geometric representation theory. *Selecta Math. (N.S.)*, 20(2):609–625, 2014.

[8] Bruce Fontaine, Joel Kamnitzer, and Greg Kuperberg. Buildings, spiders, and geometric Satake. *Compos. Math.*, 149(11):1871–1912, 2013.

[9] William Fulton. *Young tableaux*, volume 35 of *London Mathematical Society Student Texts*. Cambridge University Press, Cambridge, 1997.

[10] Victor Ginzburg. Perverse sheaves on a loop group and Langlands duality. alg-geom/9511007, 1995.

[11] Alexander Goncharov and Linhui Shen. Geometry of canonical bases and mirror symmetry. *Invent. Math.*, 202(2):487–633, 2015.

[12] Curtis Greene. An extension of Schensted's theorem. *Advances in Math.*, 14:254–265, 1974.

[13] Curtis Greene. Some partitions associated with a partially ordered set. *Journal of Combinatorial Theory, Series A*, 20(1):69 – 79, 1976.

[14] Thomas J. Haines. Structure constants for Hecke and representation rings. *Int. Math. Res. Not.*, (39):2103–2119, 2003.

[15] André Henriques. A periodicity theorem for the octahedron recurrence. *J. Algebraic Combin.*, 26(1):1–26, 2007.

[16] André Henriques and Joel Kamnitzer. The octahedron recurrence and $\mathfrak{gl}_n$ crystals. *Adv. Math.*, 206(1):211–249, 2006.

[17] Wim H. Hesselink. A classification of the nilpotent triangular matrices. *Compositio Math.*, 55(1):89–133, 1985.

[18] Jiuzu Hong and Linhui Shen. Tensor invariants, saturation problems, and Dynkin automorphisms. *Advances in Mathematics*, 285:629–657, 11 2015.

[19] Joel Kamnitzer. Hives and the fibres of the convolution morphism. *Selecta Math. (N.S.)*, 13(3):483–496, 2007.

[20] Donald E. Knuth. Permutations, matrices, and generalized Young tableaux. *Pacific J. Math.*, 34(3):709–727, 1970.

[21] Allen Knutson and Terence Tao. The honeycomb model of $\mathrm{GL}_n(\mathbb{C})$ tensor products I: Proof of the saturation conjecture. *J. Amer. Math. Soc.*, 12(4):1055–1090, 1999.

[22] Allen Knutson, Terence Tao, and Christopher Woodward. A positive proof of the Littlewood-Richardson rule using the octahedron recurrence. *Electron. J. Combin.*, 11(1):Research Paper 61, 18, 2004.

[23] Peter Littelmann. Characters of representations and paths in $\mathfrak{h}^*_{\mathbb{R}}$. In *Representation theory and automorphic forms (Edinburgh, 1996)*, volume 61 of *Proc. Sympos. Pure Math.*, pages 29–49. Amer. Math. Soc., Providence, RI, 1997.

[24] D. E. Littlewood and A. R. Richardson. Group characters and algebra. *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character*, 233:99–141, 1934.

[25] George Lusztig. Singularities, character formulas, and a q-analog of weight multiplicities. In *Analysis and topology on singular spaces, II, III (Luminy,*

1981), volume 101 of *Astérisque*, pages 208–229. Soc. Math. France, Paris, 1983.

[26] Ivan Mirković and Kari Vilonen. Geometric Langlands duality and representations of algebraic groups over commutative rings. *Ann. of Math. (2)*, 166(1):95–143, 2007.

[27] Stephan Pfannerer, Martin Rubey, and Bruce W. Westbury. Promotion, evacuation and cactus groups. *Sém. Lothar. Combin.*, 78B:Art. 71, 12, 2017.

[28] Thomas Walton Roby, V. *Applications and extensions of Fomin's generalization of the Robinson-Schensted correspondence to differential posets*. 1991. Thesis (Ph.D.)–Massachusetts Institute of Technology.

[29] Daniele Rosso. Classic and mirabolic Robinson-Schensted-Knuth correspondence for partial flags. *Canadian Journal of Mathematics*, 64, 05 2010.

[30] Bruce E. Sagan. *The symmetric group*, volume 203 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, second edition, 2001. Representations, combinatorial algorithms, and symmetric functions.

[31] Craige Schensted. Longest increasing and decreasing subsequences. *Canad. J. Math.*, 13:179–191, 1961.

[32] Nicolas Spaltenstein. The fixed point set of a unipotent transformation on the flag manifold. *Nederl. Akad. Wetensch. Proc. Ser. A* **79**=*Indag. Math.*, 38(5):452–456, 1976.

[33] David E Speyer. Schubert problems with respect to osculating flags of stable rational curves. *Algebr. Geom.*, 1(1):14–45, 2014.

[34] Richard P. Stanley. *Enumerative combinatorics. Vol. 2*, volume 62 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 1999. With a foreword by Gian-Carlo Rota and appendix 1 by Sergey Fomin.

[35] Robert Steinberg. On the desingularization of the unipotent variety. *Invent. Math.*, 36:209–224, 1976.

[36] Robert Steinberg. An occurrence of the Robinson-Schensted correspondence. *Journal of Algebra*, 113(2):523–528, 1988.

[37] Sheila Sundaram. *On the combinatorics of representations of $Sp(2n, \mathbb{C})$*. 1986. Thesis (Ph.D.)–Massachusetts Institute of Technology.

[38] Marc A. A. van Leeuwen. The Robinson-Schensted and Schützenberger algorithms, an elementary approach. *Electron. J. Combin.*, 3(2):Research Paper 15, approx. 32 pp. 1996. The Foata Festschrift.

[39] Marc A. A. van Leeuwen. An analogue of jeu de taquin for Littelmann's crystal paths. *Sém. Lothar. Combin.*, 41:Art. B41b, 23 pp. 1998.

[40] Marc A. A. van Leeuwen. Flag varieties and interpretations of Young tableau algorithms. *J. Algebra*, 224(2):397–426, 2000.

[41] Xavier Viennot. Une forme geometrique de la correspondance de Robinson-Schensted. In Dominique Foata, editor, *Combinatoire et Représentation du Groupe Symétrique: Actes de la Table Ronde du C.N.R.S. tenue à l'Université Louis-Pasteur de Strasbourg, 26 au 30 avril 1976*, pages 29–58. Springer Berlin Heidelberg, Berlin, Heidelberg, 1977.

[42] Bruce W. Westbury. Coboundary categories and local rules. arXiv:1705.07141, 2018.