

Dual equivalence graphs, ribbon tableaux and Macdonald polynomials

by

Sami Hayes Assaf

B.A. (University of Notre Dame) 2001

A dissertation submitted in partial satisfaction of the
requirements for the degree of
Doctor of Philosophy

in

Mathematics

in the

GRADUATE DIVISION

of the

UNIVERSITY OF CALIFORNIA, BERKELEY

Committee in charge:
Professor Mark Haiman, Chair
Professor Nicolai Reshetikhin
Professor Mary Gaillard

Spring 2007

Dual equivalence graphs, ribbon tableaux and Macdonald polynomials

Copyright 2007

by

Sami Hayes Assaf

Abstract

Dual equivalence graphs, ribbon tableaux and Macdonald polynomials

by

Sami Hayes Assaf

Doctor of Philosophy in Mathematics

University of California, Berkeley

Professor Mark Haiman, Chair

We make a systematic study of a new combinatorial construction called a dual equivalence graph. We axiomatize such constructions and prove that the generating functions of these graphs are Schur positive. We construct a graph on k -ribbon tableaux which we conjecture to be a dual equivalence graph, and we prove the conjecture for $k \leq 3$. This implies the Schur positivity of the k -ribbon tableaux generating functions introduced by Lascoux, Leclerc and Thibon. From Haglund's formula for the transformed Macdonald polynomials, this has the further consequence of a combinatorial expansion of the Macdonald-Kostka polynomials indexed by a partition with at most 3 columns.

Professor Mark Haiman
Dissertation Committee Chair

To Gloria, for those precious eighteen years.
To Sean, for all the rest.

Contents

1	Introduction	1
2	Background	5
2.1	Partitions and tableaux	5
2.2	Symmetric functions	8
2.3	Ribbon tableaux	10
2.4	LLT polynomials	14
3	Dual equivalence graphs	16
3.1	The standard dual equivalence graph	16
3.2	Axiomatization of dual equivalence	17
3.3	Proof of Theorem 3.11	20
3.4	Constructing edges by induction	27
4	LLT polynomials	32
4.1	The vertices and signature function	32
4.2	Constructing the edges	35
4.3	The graph on 3-ribbon tableaux is a dual equivalence graph	39
5	Macdonald polynomials	48
5.1	Haglund's formula for Macdonald polynomials	49
5.2	A combinatorial formula for Kostka-Macdonald polynomials	51
	Bibliography	53

Acknowledgments

This dissertation was only possible with guidance and encouragement from my academic advisor, Mark Haiman. His breadth and depth of knowledge is truly an inspiration, and I am eternally grateful for the countless hours of conversations during which he helped me to develop the ideas and results which became this thesis.

Without encouragement from Sean Borman, I would never have considered majoring in mathematics. I was fortunate to have many inspiring professors as an undergraduate, most notably Professors Peter Cholak and Frank Connolly. It was the support and encouragement from these individuals which helped me to develop a strong foundation in mathematics and an intense enthusiasm for the subject, and for that I am forever grateful.

I am grateful, as well, to the tremendous faculty at the Berkeley Mathematics Department for providing me with a stimulating environment in which to study, and also to the tireless staff for helping me to maneuver through the obstacle course of bureaucracy so that I could focus on mathematics. I am indebted to Mark Haiman for introducing me to combinatorics which has become my passion as well as my profession, and for finding a research problem which has kept me enthralled for so many years. I must also thank Monica Vazirani for being a constant source of information and encouragement throughout the arduous process of graduating, and Stephanie Somersille for going through the ups and downs of graduate school (and the fire trail) with me.

I owe thanks to the National Science Foundation and to the Department of Defense for funding my tenure in graduate school (and beyond) with generous fellowship support.

My most sincere thanks go to Monica Vazirani, Ian Grojnowski, Jim Haglund and especially Mark Haiman for many stimulating conversations and insightful suggestions which helped me to develop a more elegant formulation of these results which I hope will ultimately lead to a complete generalization. Finally, my thanks to Mark Haiman, Nicolai Reshetikhin and Mary Gaillard for serving on my dissertation committee.

Chapter 1

Introduction

The primary subject of this dissertation is symmetric function theory, which plays an important role in many areas of mathematics including algebraic combinatorics, representation theory, Lie groups and Lie algebras, algebraic geometry and the theory of special functions. Multiplicities of irreducible components, dimensions of algebraic varieties, and various other algebraic constructions that require the computation of certain integers may be translated to the computation of the coefficients in the expansion of certain generalizations of the Schur basis. Often the coefficients can be identified as generating functions of tableau-like structures, providing a useful and often insightful *combinatorial* formula.

Since their introduction in 1988, Macdonald polynomials have been intensely studied and have been found to have applications in such areas as representation theory, algebraic geometry, group theory, statistics, and quantum mechanics. Unfortunately, given the indirect definition of these polynomials as the unique functions satisfying certain conditions, most results require difficult technical machinery. Recent work by Haglund, Haiman and Loehr has connected the study of Macdonald polynomials to that of LLT polynomials. Though both Macdonald polynomials and LLT polynomials have been shown to be Schur positive using geometric methods, finding a *combinatorial* proof of positivity remains an important open problem in this area.

The immediate purpose of this thesis is to establish a combinatorial formula for the Schur expansion of the k -ribbon tableaux generating functions known as LLT polynomials when $k \leq 3$. As a corollary, this yields a combinatorial formula for the Kostka-Macdonald polynomials for partitions with at most 3 columns. Furthermore, we conjecture that the construction used generalizes to arbitrary k . Our real purpose, however, is not only to obtain the above results, but also to introduce a new combinatorial construction, called a *dual equivalence graph*, by which one can establish the Schur positivity of functions which are expressed in terms of monomials.

The LLT polynomials $\tilde{G}_\mu^{(k)}(x; q)$, defined by Lascoux, Leclerc and Thibon in 1997 [LLT97], are the q -generating functions of k -ribbon tableaux of shape μ weighted by a statistic called *cospin*.

By the Stanton-White correspondence [SW85], LLT polynomials may be realized as q -analogs of products of Schur functions.

Using Fock space representations of quantum affine Lie algebras constructed by Kashiwara, Miwa and Stern [KMS95], Lascoux, Leclerc and Thibon [LLT97] proved that $\tilde{G}_\mu^{(k)}(x; q)$ is a symmetric function. Thus we may define the Schur coefficients, $\tilde{K}_{\lambda, \mu}^{(k)}(q)$, by

$$\tilde{G}_\mu^{(k)}(x; q) = \sum_{\lambda} \tilde{K}_{\lambda, \mu}^{(k)}(q) s_{\lambda}(x).$$

In 2000, using Kazhdan-Lusztig theory, Leclerc and Thibon [LT00] proved that $\tilde{K}_{\lambda, \mu}^{(k)}(q) \in \mathbb{N}[q]$ for straight shapes μ . This has recently been extended by Grojnowski and Haiman [GH07] to skew shapes. The question remains to find a combinatorial interpretation for $\tilde{K}_{\lambda, \mu}^{(k)}(q)$.

Macdonald polynomials were first defined by Macdonald in 1988 [Mac88] as the unique functions satisfying certain triangularity and orthogonality conditions. From their definition, the transformed Macdonald polynomials, $\tilde{H}_\mu(x; q, t)$, are known to be symmetric functions. Therefore we may define the Kostka-Macdonald polynomials, $\tilde{K}_{\lambda, \mu}(q, t)$, which give the Schur expansion of Macdonald polynomials, by

$$\tilde{H}_\mu(x; q, t) = \sum_{\lambda} \tilde{K}_{\lambda, \mu}(q, t) s_{\lambda}(x).$$

The Macdonald positivity conjecture (now theorem) states that $\tilde{K}_{\lambda, \mu}(q, t) \in \mathbb{N}[q, t]$. Following an idea outlined by Procesi, Garsia and Haiman [GH93] conjectured that the transformed Macdonald polynomials $\tilde{H}_\mu(x; q, t)$ could be realized as the bigraded characters of certain modules for the diagonal action of S_n on two sets of variables. By analyzing the algebraic geometry of the Hilbert scheme of n points in the plane, Haiman [Hai01] proved this conjecture and consequently establish that $\tilde{K}_{\lambda, \mu}(q, t) \in \mathbb{N}[q, t]$. However, finding a combinatorial proof of Macdonald positivity, or better still a combinatorial description of $\tilde{K}_{\lambda, \mu}(q, t)$, remains an important open question.

In 2004, Haglund [Hag04] conjectured a combinatorial formula for the monomial expansion of the transformed Macdonald polynomials. This formula, proven by Haglund, Haiman and Loehr [HHL05a], does not give a combinatorial proof of $\tilde{K}_{\lambda, \mu}(q, t) \in \mathbb{N}[q, t]$ since monomials, and even the monomial symmetric functions m_{λ} , are not Schur positive. However, combining Theorem 2.3, Proposition 3.4 and equation (23) from [HHL05a], to give a combinatorial description of $\tilde{K}_{\lambda, \mu}(q, t)$ it suffices to give a description of $\tilde{K}_{\lambda, \nu}^{(\mu_1)}(q)$ for certain skew shapes ν which depend upon μ .

Some progress has been made towards finding combinatorial formulas for $\tilde{K}_{\lambda, \mu}^{(k)}(q)$ and $\tilde{K}_{\lambda, \mu}(q, t)$. In 1995, Carré and Leclerc [CL95] presented a combinatorial interpretation of $\tilde{K}_{\lambda, \mu}^{(2)}(q)$ in their study of 2-ribbon tableaux. Though their result is correct, the proof which they present is incomplete. In 2005, van Leeuwen [vL05] gave the first complete proof of their result using the theory of crystal graphs. In 2004, Haglund [Hag04] gave a combinatorial formula for $\tilde{K}_{\lambda, \mu}(q, t)$ when μ is a partition with 2 columns. In both cases, finding extensions for these formulae has proven elusive.

In 2005, Haiman suggested looking at the dual equivalence relation on standard Young tableaux defined in [Hai92]. From this relation, he suggested defining an edge-colored graph on tableaux and investigating how this graph may be related to the crystal graph on tableaux. The result of this idea is a new combinatorial method for establishing the Schur positivity of a function which is expressed in terms of monomials. In particular, this method has successfully been applied to LLT polynomials to obtain a combinatorial formula for $\tilde{K}_{\lambda,\mu}^{(k)}(q)$ when $k \leq 3$. From Haglund's formula for Macdonald polynomials, this result has been extended to give a combinatorial formula for $\tilde{K}_{\lambda,\mu}(q, t)$ when μ is a partition with at most 3 columns. The constructions used to obtain these formulae may be extended beyond 3, and it is conjectured that doing so will give a combinatorial formula for $\tilde{K}_{\lambda,\mu}^{(k)}(q)$ and $\tilde{K}_{\lambda,\mu}(q, t)$.

The thesis is organized as follows. In Chapter 2, we introduce familiar objects from the theory of symmetric functions, for the most part following the notation of Macdonald [Mac95]. The main reference for the presentation of ribbon tableau is the classic text on representation theory by James and Kerber [JK81]. We also review the original definition of LLT polynomials given in [LLT97], though we will use an alternate definition in Chapter 4.

Chapter 3 is devoted to the development of the theory of dual equivalence graphs. We review the original definition of dual equivalence given in [Hai92], and in Section 3.1 show how from this we can obtain a graph whose vertices are indexed by standard Young tableaux and whose connected components are indexed by partitions. We define the generating function of a vertex-signed, edge-colored graph using quasi-symmetric functions, and show that the generating function of the *standard dual equivalence graph associated to λ* is the Schur function s_λ . In Section 3.2, a *dual equivalence graph* is defined to be a vertex-signed, edge-colored graph satisfying certain conditions. We then present our first main result, Theorem 3.11, which states that every connected component of a dual equivalence graph is isomorphic to one of the standard dual equivalence graphs constructed in Section 3.1. In terms of symmetric functions, this gives Corollary 3.12, which states that the generating function of a dual equivalence graph is Schur positive. The proof of Theorem 3.11 is carried out in Section 3.3, where we investigate the structure of dual equivalence graphs. Section 3.4 explores how the edges of a dual equivalence graph may be constructed by induction, which lays the foundation for Chapter 4.

With these constructions in place, Chapter 4 contains the first application of dual equivalence graphs from which we obtain a combinatorial interpretation of the Schur expansion of LLT polynomials when $k \leq 3$. We begin, in Section 4.1, by reformulating the notion of a ribbon tableaux so as to encode the vertices of the proposed graph in a more convenient form. In Section 4.2, we define the basic involutions from which we construct the edges of the graph inductively. As a special case, the construction when $k = 2$ is given directly from these involutions, and as a consequence we obtain a simple combinatorial proof that $\tilde{K}_{\lambda,\mu}^{(k)}(q) \in \mathbb{N}[q]$ in contrast to the much more complicated crystal-theoretic proof by van Leeuwen. Section 4.3 is devoted to proving that the graph constructed on 3-ribbon tableaux is also a dual equivalence graph. Using the theory developed in Chapter 3, we

obtain Corollary 4.22 which gives a combinatorial interpretation of $\tilde{K}_{\lambda,\mu}^{(k)}(q)$ when $k \leq 3$.

The application to LLT polynomials is extended in Chapter 5 using the combinatorial expansion of Macdonald polynomials in terms of LLT polynomials from [Hag04, HHL05a]. We present this formula in Section 5.1, and in Section 5.2 mimic the constructions in Chapter 4 for the case of Macdonald polynomials. The section concludes with a combinatorial description of $\tilde{K}_{\lambda,\mu}(q, t)$ when $\mu_1 \leq 3$.

Chapter 2

Background

2.1 Partitions and tableaux

We represent an integer *partition* λ by the decreasing sequence of its (nonzero) parts

$$\lambda = (\lambda_1, \lambda_2, \dots, \lambda_l), \quad \lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_l > 0.$$

When it is convenient to highlight the multiplicities of each integer in λ , we may also write

$$\lambda = (1^{m_1}, 2^{m_2}, \dots)$$


where m_i , called the *multiplicity of i in λ* , is the number of times i occurs as a part of λ . We denote the size of λ by $|\lambda| = \sum_i \lambda_i$ and the length of λ by $l(\lambda) = \max\{i : \lambda_i > 0\}$. If $|\lambda| = n$, we say that λ is a *partition of n* .

A *composition* π is a sequence of nonnegative integers

$$\pi = (\pi_1, \pi_2, \dots, \pi_m), \quad \pi_i \geq 0.$$

As with partitions, the size of a composition π is denoted by $|\pi| = \sum_i \pi_i$. If $|\pi| = n$, we say that π is a *composition of n* . The *rectification* of a composition π is the unique partition obtained by rearranging the parts of π into a decreasing sequence of nonzero integers. For example, the rectification of $\pi = (4, 0, 4, 1, 0, 5, 0)$ is $\lambda_\pi = (5, 4, 4, 1)$. Note that if λ_π is the rectification of π , then $|\lambda_\pi| = |\pi|$.


The *Young diagram* of a partition λ is the set of points (i, j) in the lattice quadrant $\mathbb{Z}_{>0} \times \mathbb{Z}_{>0}$ plane such that $1 \leq i \leq \lambda_j$. We draw the diagram so that each point (i, j) is represented by the unit cell southwest of the point. Abusing notation, we will write λ for both the partition and its diagram. With this notation, the cell of λ in position (i, j) lies in the i th column and j th row of λ .

Figure 2.1: The Young diagram for $(5, 4, 4, 1)$.


The *conjugate partition* λ' of λ is defined by

$$\lambda'_j = \sum_{i \geq j} m_i.$$

The Young diagram of λ' may be obtained by reflecting the diagram of λ across the main diagonal. For example, compare the Young diagram for $(5, 4, 4, 1)$ in Figure 2.1 with the Young diagram of its conjugate $(4, 3, 3, 3, 1)$ in Figure 2.2.

Figure 2.2: The Young diagram for $(5, 4, 4, 1)' = (4, 3, 3, 3, 1)$.

For partitions λ, μ , we write $\mu \subset \lambda$ whenever the diagram of μ is contained within the diagram of λ , equivalently $\mu_i \leq \lambda_i$ for all i . In this case, we define the *skew diagram* λ/μ to be the set theoretic difference $\lambda - \mu$. The conjugate of a skew diagram λ/μ is λ'/μ' .

Figure 2.3: The skew Young diagram for $(5, 4, 4, 1)/(4, 3, 2)$.

A *connected skew diagram* is one for which exactly one cell has no cell immediately north or west of it, and exactly one cell has no cell immediately south or east of it. For example, the skew diagram depicted in Figure 2.3 is the disjoint union of three connected skew diagrams.

A skew diagram λ/μ forms a *horizontal strip* if no cell of λ/μ lies immediately north of another. Similarly, a skew diagram forms a *vertical strip* if no cell of the diagram lies immediately east of another.

Let \geq denote the *dominance partial ordering* on partitions of n , defined by

$$\lambda \geq \mu \Leftrightarrow \lambda_1 + \lambda_2 + \cdots + \lambda_i \geq \mu_1 + \mu_2 + \cdots + \mu_i \quad \forall i. \quad (2.1)$$

It is well known [Mac95] that conjugation inverts this ordering, i.e. $\lambda \geq \mu$ if and only if $\lambda' \leq \mu'$.

For a cell x in the diagram of λ , define the *arm of x* to be the set of cells east of x , and the *leg of x* to be the set of cells north x . The *hook of x* is x together its arm and leg. Denote by

$$a(x), \quad l(x), \quad h(x) = 1 + a(x) + l(x)$$

the sizes of the arm, leg and hook of x . For the example in Figure 2.4, we have $a(x) = 2$, $l(x) = 1$ and $h(x) = 4$.




Figure 2.4: The arm and leg of the cell x .

For a partition diagram, the *content* of a cell indexes the diagonal on which it occurs, i.e. $c(x) = i - j$ when the cell x lies in position $(i, j) \in \mathbb{Z}_{>0} \times \mathbb{Z}_{>0}$ as shown in Figure 2.5. We will often be concerned with the cells of a given row, column or content of λ .




Figure 2.5: The contents of the cells of the partition $(5, 4, 4, 1)$.

A *filling* of a (skew) diagram λ is a map $S : \lambda \rightarrow \mathbb{N}$. A *semi-standard Young tableau* is a filling which is weakly increasing along each row of λ and strictly increasing along each column. Equivalently, a filling is semi-standard if and only if the cells labeled i form a horizontal strip for each i , and the union of the cells with labels $< i$ form a (skew) diagram for all i . A semi-standard Young tableau is *standard* if it is a bijection from λ to $[n]$, where $[n] = \{1, 2, \dots, n\}$. If T is a standard Young tableau, then the conjugate tableau T' , obtained by reflecting T along the main diagonal, is as well, however this is not necessarily the case for semi-standard Young tableaux.

For λ a diagram of size n and π a composition of n , we define

$$\begin{aligned} \text{SSYT}(\lambda) &= \{\text{semi-standard tableaux } T : \lambda \rightarrow \mathbb{N}\} \\ \text{SSYT}(\lambda, \pi) &= \{\text{semi-standard tableaux } T : \lambda \rightarrow \mathbb{N} \text{ with entries } 1^{\pi_1}, 2^{\pi_2}, \dots\} \\ \text{SYT}(\lambda) &= \{\text{standard tableaux } T : \lambda \twoheadrightarrow [n]\} = \text{SSYT}(\lambda, (1^n)). \end{aligned}$$

For $T \in \text{SSYT}(\lambda, \pi)$, we say that T has *shape* λ and *weight* π . Note that if $T \in \text{SSYT}(\lambda, \mu)$ for partitions λ and μ , then $\lambda \geq \mu$.

4				
3	3	4	5	
2	2	3	4	
1	1	1	1	2

Figure 2.6: A semi-standard Young tableau of shape $(5, 4, 4, 1)$ and weight $(4, 3, 3, 3, 1)$.

It is often useful to represent a tableaux by a sequence of integers. The *row reading word* of T , denoted $r(T)$, is the word obtained by reading the entries of T from west to east, north to south. Similarly, the *column reading word*, denoted $c(T)$, is obtained by reading the columns of T from north to south, west to east. From the example in Figure 2.6,

$$r(T) = 43345223411112 \quad \text{and} \quad c(T) = 43213214315412.$$

Notice that $c(T)$ is obtained by reversing $r(T')$, where T' is the conjugate of T . More importantly, if the shape of T is a partition, then T can be recovered from either of these reading words: read the word from left to right, beginning a new row whenever $r(T)_i > r(T)_{i+1}$ or a new column whenever $c(T)_i \leq c(T)_{i+1}$.

The *content reading word* of T , $w(T)$, is obtained by reading the entries of T along diagonals starting with the smallest content (northwesternmost cell). Again from the same example,

$$w(T) = 43231241351412.$$

Each diagonal may be identified by looking for descents unless $\lambda_1 > \lambda_2$, in which case there is a seemingly long diagonal at the end. By comparing the first (smallest) letter from each diagonal, there is a unique way to realign the diagonals which results in a partition shape. When aligning the diagonals, the “false” diagonal at the end of $w(T)$ will be apparent by the fact that the result is not a partition shape.

2.2 Symmetric functions

We have the familiar bases for Λ , the ring of symmetric functions, from [Mac95]: the monomial symmetric functions m_λ , the elementary symmetric functions e_λ , the complete homogeneous symmetric functions h_λ , the power sum symmetric functions p_λ and, most importantly, the Schur functions s_λ .

The notation $f(x)$ will stand for a function in countably many variables $f(x_1, x_2, \dots)$. The *monomial symmetric function*, m_λ , is defined by

$$m_\lambda(x) = \sum_{\pi} x^\pi \tag{2.2}$$

where the sum is over all compositions π which rectify to λ .

For each integer $r \geq 0$, the r th *elementary symmetric function*, e_r , is the sum of all products of r distinct variables. For a partition $\lambda = (\lambda_1, \lambda_2, \dots)$, we have

$$e_\lambda = e_{\lambda_1} e_{\lambda_2} \cdots \quad \text{where} \quad e_r(x) = \sum_{i_1 < i_2 < \cdots < i_r} x_{i_1} x_{i_2} \cdots x_{i_r} = m_{(1^r)}(x). \quad (2.3)$$

For each integer $r \geq 0$, the r th *complete homogeneous symmetric function*, h_r , is the sum of all monomials of total degree r . For a partition $\lambda = (\lambda_1, \lambda_2, \dots)$, we have

$$h_\lambda = h_{\lambda_1} h_{\lambda_2} \cdots \quad \text{where} \quad h_r(x) = \sum_{|\lambda|=r} m_\lambda(x). \quad (2.4)$$

For $r \geq 1$, the r th *power sum symmetric function*, p_r , is the sum of the r th powers of the variables x_1, x_2, \dots . For $\lambda = (\lambda_1, \lambda_2, \dots)$, we have

$$p_\lambda = p_{\lambda_1} p_{\lambda_2} \cdots \quad \text{where} \quad p_r(x) = \sum x_i^r = m_{(r)}(x). \quad (2.5)$$

Whereas the m_λ , e_λ and h_λ form \mathbb{Z} -bases of $\Lambda_{\mathbb{Z}}$, the p_λ only form a \mathbb{Q} -basis for $\Lambda_{\mathbb{Q}}$. For example, $h_2 = \frac{1}{2}(p_1^2 + p_2)$ does not have integral coefficients when expressed in terms of the p_λ .

Finally we come to the *Schur functions* s_λ , which may be defined in several ways. For the purposes of this paper, we take the tableau approach and define

$$s_\lambda(x) = \sum_{T \in \text{SSYT}(\lambda)} x^T, \quad (2.6)$$

where x^T is the monomial $x_1^{\pi_1} x_2^{\pi_2} \cdots$ when T has weight π . This formula also defines the *skew Schur functions*, $s_{\lambda/\mu}$ by taking the sum over semi-standard Young tableaux of shape λ/μ . The Schur functions form a \mathbb{Z} -basis of Λ , and in many ways form the “best” such basis.

Recall the Hall inner product $\langle -, - \rangle$ on symmetric functions defined by either of the equations

$$\langle s_\lambda, s_\mu \rangle = \delta_{\lambda, \mu} = \langle m_\lambda, h_\mu \rangle.$$

The *Kostka numbers*, $K_{\lambda, \mu}$, give the change of basis from the complete homogeneous symmetric functions to the Schur functions,

$$\langle h_\mu, s_\lambda \rangle = K_{\lambda, \mu} = \langle s_\lambda, m_\mu \rangle,$$

or equivalently

$$h_\mu = \sum_{\lambda} K_{\lambda, \mu} s_\lambda; \quad s_\lambda = \sum_{\mu} K_{\lambda, \mu} m_\mu.$$

In particular, $K_{\lambda, \mu}$ is the number of semi-standard Young tableaux of shape λ and weight μ . Throughout this paper, we are interested in certain one- and two-parameter generalizations of the Kostka numbers.

2.3 Ribbon tableaux

A *ribbon*, also called a *rim hook*, is a connected skew diagram which contains no 2×2 block. A k -*ribbon* is a ribbon comprised of exactly k cells.

Given a partition λ , we form a k -ribbon tiling as follows. Begin by tiling the diagram of λ with k -ribbons in such a way that successively removing the ribbons gives a partition diagram at every stage. Let $\lambda_{(k)}$ denote the partition which remains when no further k -ribbons can be removed. The result is a k -*ribbon tiling* of $\lambda/\lambda_{(k)}$. See Figure 2.7 for two possible 3-ribbon tilings of $(5, 4, 4, 1)/(2)$.




Figure 2.7: Two 3-ribbon tilings of $(5, 4, 4, 1)/(2)$.

The partition $\lambda_{(k)}$ is called the k -*core* of λ and does not depend upon the choice of which k -ribbons are removed from λ [JK81]. That is to say, the k -*core* of λ , denoted $\lambda_{(k)}$, is the unique partition which remains after successively removing all k -ribbons from λ . We will use the abacus model presented in [JK81] to see this and to define a k -tuple of partitions associated to λ called the k -*quotient* of λ .

Consider the diagram of λ lying in the $\mathbb{N} \times \mathbb{N}$ plane with infinite positive axes. Walk in unit steps along the boundary of λ placing a bead (\bullet) on each vertical step and a spacer (\circ) on each horizontal step. Then straighten the boundary to get a doubly infinite rod, which is a binary string uniquely representing λ . For example, see Figure 2.8.




Figure 2.8: Bead and spacer model for $(5, 4, 4, 1)$.

Define the content of a bead to be the content of the cell immediately west and the content of a spacer to be the content of the cell immediately north. Notice that there is a unique bead or spacer at each content. Moreover, a binary string of this form is the string associated to a partition if and only if the number of beads with content ≥ 0 is finite and is equal to the number of spacers with content < 0 .

Along any k -ribbon which may be removed from the diagram, there is a total of $k + 1$ beads and spacers. Removing a k -ribbon from a Young diagram changes a horizontal step to a vertical step at the beginning of the ribbon, and a vertical step to a horizontal step at the end. On the rod, the bead with content $i + k$ becomes a spacer, and the spacer with content i becomes a bead. Therefore being able to remove a k -ribbon from the Young diagram is equivalent to having a spacer k positions left of a bead on the rod.




Figure 2.9: Two 3-ribbons which may be removed from $(5, 4, 4, 1)$.

Since removing k -ribbons always changes beads and spacers which are exactly k positions apart, it makes sense to break the rod into an abacus consisting of k rods in such a way that all beads and spacers with the same content modulo k appear in order on the same rod. Place the rods in rows beginning with beads and spacers with content congruent to 0 modulo k . Then removing a k -ribbon amounts to moving a single bead one space to the left on a given rod, e.g. see Figure 2.10.




Figure 2.10: 3 rod abacus model for $(5, 4, 4, 1)$.

On the abacus model, the k -core of λ is obtained by sliding all of the beads on each rod to the left as far as possible. Since the result is independent of the order in which the beads are moved, the k -core of λ depends only on the shape of λ and not on the choice of a k -ribbon tiling. From Figure 2.11, the 3-core of $(5, 4, 4, 1)$ is (2) .




Figure 2.11: The 3-core of $(5, 4, 4, 1)$.

Returning to the original k -tuple of rods which corresponds to λ , by considering each rod as representing a partition we obtain a k -tuple of partitions associated with λ called the k -quotient of λ , denoted $(\lambda^{(0)}, \dots, \lambda^{(k-1)})$. Again, it is clear from the abacus model that the k -quotient of λ depends only on λ and not on the choice of a k -ribbon tiling of $\lambda/\lambda_{(k)}$.




Figure 2.12: The 3-quotient of $(5, 4, 4, 1)$.

Keeping in mind that each rod is doubly infinite, in order to make the correspondence between λ and $(\lambda^{(0)}, \dots, \lambda^{(k-1)})$ a bijection we must include additional information with the k -quotient which indicates how the k rods should be aligned with respect to one another. One way to do this is to include the k -core, hence λ may be uniquely recovered from its k -core and k -quotient [JK81]. There is another way to achieve this bijection which will be useful for later discussions.

By remembering the content of each bead or spacer from λ , we obtain the *shifted content* for the beads and spacers of the k -quotient and consequently for the cells of the k -quotient. For each $\lambda^{(i)}$, let s_i be the shifted content of the main diagonal of $\lambda^{(i)}$. From earlier remarks all cells of $\lambda^{(i)}$ will have shifted contents congruent to i modulo k , in particular, $s_i \equiv i \pmod k$. Moreover, given the s_i , we can recover the *shifted content*, \tilde{c} , of the cells of the quotient by shifting the usual content for a cell x of $\lambda^{(i)}$ by

$$\tilde{c}(x) = kc(x) + s_i . \tag{2.7}$$

Therefore λ corresponds uniquely to the k -tuple of partitions $(\lambda^{(0)}, \dots, \lambda^{(k-1)})$ together with the k -tuple of integers (s_0, \dots, s_{k-1}) such that $s_i \equiv i \pmod k$. Back to our running example, Figure 2.5 shows the shifted contents for each piece of the 3-quotient from which we see that the partition $(5, 4, 4, 1)$ corresponds to 3-quotient $(\emptyset, (1, 1), (2))$ together with the 3-tuple $(-3, 4, -1)$.

Just as with ordinary tableaux, we may form *skew k -ribbon tableaux*, though in this case the constraints are slightly more subtle. In order for a k -ribbon tiling of λ to restrict to a k -ribbon tiling of μ , it is not enough to stipulate containment, $\mu \subset \lambda$. In addition, we must ensure that λ




Figure 2.13: Shifted contents for the 3-quotient of $(5, 4, 4, 1)$.

and μ have the same k -core, $\mu_{(k)} = \lambda_{(k)}$, and there must be containment of k -quotients meaning that $\mu^{(i)} \subset \lambda^{(i)}$ for $i = 0, \dots, k - 1$. When these conditions hold, the k -core of λ/μ is $\lambda_{(k)}$, and the k -quotient is $(\lambda^{(0)}/\mu^{(0)}, \dots, \lambda^{(k-1)}/\mu^{(k-1)})$.

A skew diagram λ/μ forms a *horizontal k -ribbon strip* if it has a k -ribbon tiling such that the ribbons may be removed in increasing order of content. In this case, there is a unique such tiling with this property. If λ/μ is a horizontal k -ribbon strip, then $\lambda^{(i)}/\mu^{(i)}$ is a horizontal strip for each $i = 0, \dots, k - 1$.

We may label the ribbons of a k -ribbon tiling of $\lambda/\lambda_{(k)}$ with positive integers to form a *k -ribbon filling* of $\lambda/\lambda_{(k)}$. Such a filling gives a *semi-standard k -ribbon tableau* if the ribbons labeled i form a horizontal k -ribbon strip for each i , and the union of the ribbons with labels $< i$ form a skew k -ribbon tiling for all i . A semi-standard k -ribbon tableau is called *standard* if the ribbons are labeled from 1 to n without repetition. Amending prior notation, define

$$\begin{aligned} \text{SSYT}_k(\lambda) &= \{\text{semi-standard } k\text{-ribbon tableaux of shape } \lambda/\lambda_{(k)}\}, \\ \text{SYT}_k(\lambda) &= \{\text{standard } k\text{-ribbon tableaux of shape } \lambda/\lambda_{(k)}\}. \end{aligned}$$

Analogous to tableaux, the *weight* of a k -ribbon tableau is the composition π , where π_i is the number of k -ribbons labeled i . Define the *content* of a ribbon to be the content of its southeastern-most cell. By convention, we also label each ribbon in this distinguished cell.




Figure 2.14: Two elements of $\text{SSYT}_3(5, 4, 4, 1)$.

Using the bijection between partitions with a given k -core and their k -quotients, we obtain a bijection between k -ribbon fillings of shape $\lambda/\lambda_{(k)}$, and fillings of the k -tuple of shapes

$(\lambda^{(0)}, \dots, \lambda^{(k-1)})$ with shifted contents (s_0, \dots, s_{k-1}) . In particular, this gives the Stanton-White correspondence [SW85] between semi-standard k -ribbon tableau of shape $\lambda/\lambda_{(k)}$ and k -tuples of semi-standard Young tableau of shapes $(\lambda^{(0)}, \dots, \lambda^{(k-1)})$. For an example, see Figures 2.14 and 2.15. Note that under this correspondence, the content of a ribbon corresponds to the shifted content of a cell in the quotient.




Figure 2.15: The corresponding 3-quotients for Figure 2.14.

As with tableaux, we can form the row, column and content reading words for a k -ribbon tableaux, though care must be taken as to when a particular ribbon is read. In this paper, we are interested primarily in the content reading word for a k -ribbon tableaux which is obtained by reading the k -ribbons by content. In the case of Figure 2.14, the content words are

$$w(T_A) = 3231 \quad \text{and} \quad w(T_B) = 1321.$$

The content reading word may also be obtained from the k -quotient by interspersing the usual content reading words of the given tableaux.

2.4 LLT polynomials

Following the analogy with semi-standard Young tableaux and Schur functions, we may form the generating function for k -ribbon tableaux

$$G_\lambda^{(k)}(x) = \sum_{T \in \text{SSYT}_k(\lambda)} x^T, \quad (2.8)$$

where again x^T is the monomial $x_1^{\pi_1} x_2^{\pi_2} \dots$ when T has weight π . Using the Stanton-White correspondence between semi-standard k -ribbon tableaux on $\lambda/\lambda_{(k)}$, and k -tuples of semi-standard tableaux on $(\lambda^{(0)}, \dots, \lambda^{(k-1)})$, equation (2.8) reduces to a product of Schur functions.

$$G_\lambda^{(k)}(x) = \sum_{T \in \text{SSYT}_k(\lambda)} x^T = \prod_{i=0}^{k-1} \sum_{T^{(i)} \in \text{SSYT}(\lambda^{(i)})} x^{T^{(i)}} = \prod_{i=0}^{k-1} s_{\lambda^{(i)}}(x) \quad (2.9)$$

In 1997, Lascoux, Leclerc and Thibon [LLT97] introduced a new family of symmetric functions which are q -analogs of products of Schur functions. In order to define LLT polynomials, we weight equation (2.8) with an additional statistic, called *cospin*.

Define the *spin* of a ribbon R to be

$$s(R) = \frac{\text{ht}(R) - 1}{2}, \quad (2.10)$$

where $\text{ht}(R)$ denotes the height of the ribbon. Extending this to a ribbon tableau T , define $s(T)$ to be the sum of the spins of the ribbons of T . Note $s(T) \in \frac{1}{2}\mathbb{N}$.

Given a shape λ , define

$$s_k^*(\lambda) = \max\{s(T) \mid T \in \text{SSYT}_k(\lambda)\}. \quad (2.11)$$

The *cospin*, denoted $\tilde{s}(T)$, of a k -ribbon tableau T of shape λ is given by

$$\tilde{s}(T) = s_k^*(\lambda) - s(T). \quad (2.12)$$

It is shown in [HHL⁺05b] that $s(T_1) - s(T_2) \in \mathbb{Z}$ whenever T_1 and T_2 are two k -ribbon tableaux of the same shape. Consequently, $\tilde{s}(T) \in \mathbb{N}$ for any k -ribbon tableaux T .

The LLT polynomial $\tilde{G}_\mu^{(k)}(x; q)$ is defined by

$$\tilde{G}_\mu^{(k)}(x; q) = \sum_{T \in \text{SSYT}_k(\mu)} q^{\tilde{s}(T)} x^T. \quad (2.13)$$

Using Fock space representations of quantum affine Lie algebras constructed by Kashiwara, Miwa and Stern [KMS95], Lascoux, Leclerc and Thibon proved that $\tilde{G}_\mu^{(k)}(x; q)$ is a symmetric function. Thus we may define the Schur coefficients, $\tilde{K}_{\lambda, \mu}^{(k)}(q)$, by

$$\tilde{G}_\mu^{(k)}(x; q) = \sum_{\lambda} \tilde{K}_{\lambda, \mu}^{(k)}(q) s_\lambda(x). \quad (2.14)$$

Using Kazhdan-Lusztig theory, Leclerc and Thibon [LT00] proved $\tilde{K}_{\lambda, \mu}^{(k)}(q) \in \mathbb{N}[q]$ for straight shapes μ . This has recently been extended by Grojnowski and Haiman [GH07] to skew shapes.

An incomplete combinatorial proof of $\tilde{K}_{\lambda, \mu}^{(2)}(q) \in \mathbb{N}[q]$ was given by Carré and Leclerc in [CL95]. This proof was later completed by van Leeuwen [vL05] using the theory of crystals. The proof is quite involved and relies on special properties of $k = 2$ which fail for $k \geq 3$.

One of the main purposes of this paper is to find a combinatorial description of the Schur coefficients $K_{\lambda, \mu}^{(k)}(q)$ for $k \leq 3$, thereby giving a combinatorial proof that $K_{\lambda, \mu}^{(3)}(q) \in \mathbb{N}[q]$. Furthermore, this formula is conjectured to hold for all k .

Chapter 3

Dual equivalence graphs

Dual equivalence was first defined by Haiman [Hai92] as a relation on tableaux which is “dual” to *jeu de taquin* equivalence under the Schensted correspondence. In Section 3.1, we use this relation to construct a graph whose vertices are given by standard Young tableaux and whose connected components are indexed by partitions. Using quasi-symmetric functions, we define the generating function on the vertices of these graphs which provides the connection with symmetric functions. In Section 3.2, we present an axiomatization for when a signed, colored graph is a *dual equivalence graph*. The main justification of this definition is given in Theorem 3.11, where we prove that the connected components of a graph satisfying these axioms are isomorphic to the connected components of the standard dual equivalence graph on tableaux. On the symmetric function level, Corollary 3.12 shows that the generating function for a dual equivalence graph is Schur positive. The proof of Theorem 3.11 is carried out in Section 3.3, where we analyze the structure of dual equivalence graphs. Section 3.4 shows how a dual equivalence graph may be extended, thereby laying the foundation for the dual equivalence graphs we will construct in Chapter 4.

3.1 The standard dual equivalence graph

We begin by recalling the definition of dual equivalence on permutations regarded as words on $[n]$, which can easily be extended to standard Young tableaux via the content reading word.

Definition 3.1. An *elementary dual equivalence* on three consecutive letters $i-1, i, i+1$ of a permutation is given by switching the outer two letters whenever the middle letter is not i .

$$\cdots i \cdots i \pm 1 \cdots i \mp 1 \cdots \equiv^* \cdots i \mp 1 \cdots i \pm 1 \cdots i \cdots$$

Two permutations are *dual equivalent* if they differ by some sequence of elementary dual equivalences. Two standard tableaux of the same shape are *dual equivalent* if their content reading words are.

Construct an edge-colored graph on standard tableaux of partition shape from the dual equivalence relation in the following way. Whenever two standard tableaux T, U have content reading words which differ by an elementary dual equivalence for $i-1, i, i+1$, connect T and U with an edge colored by i . The connected components of the graph so constructed are the dual equivalence classes of standard tableaux. Let \mathcal{G}_λ denote the subgraph on tableaux of shape λ . The following proposition tells us that the \mathcal{G}_λ exactly give the connected components of the graph.

Proposition 3.2 ([Hai92]). *Two standard tableaux on partition shapes are dual equivalent if and only if they have the same shape.*

For any subset $D \subset [n-1]$, Gessel [Ges84] defined the *quasi-symmetric function*

$$Q_{n,D}(x) = \sum_{\substack{i_1 \leq \dots \leq i_n \\ i_j = i_{j+1} \Rightarrow j \notin D}} x_{i_1} \cdots x_{i_n}. \quad (3.1)$$

We can use quasi-symmetric functions to define a generating function on the vertices of a graph. First, we add a signature for each vertex, which may be regarded as an indicator function for a subset of $\{1, 2, \dots, n-1\}$, setting $i \in D$ if and only if $\sigma_i = -1$.

Let T be a standard tableau on $[n]$ with content reading word $w(T)$. Define the *descent signature* $\sigma(T) \in \{\pm 1\}^{n-1}$ by

$$\sigma(T)_i = \begin{cases} +1 & \text{if } i \text{ appears to the left of } i+1 \text{ in } w(T) \\ -1 & \text{if } i+1 \text{ appears to the left of } i \text{ in } w(T) \end{cases}. \quad (3.2)$$

It is worth noting that replacing the content reading word with the row or column reading word in 3.2 will not change the signature.

The generating function associated to \mathcal{G}_λ , $|\lambda| = n$, is defined by

$$g_\lambda(x) = \sum_{v \in V(\mathcal{G}_\lambda)} Q_{n, \sigma(v)}(x). \quad (3.3)$$

But this is nothing more than Gessel's quasi-symmetric function expansion for a Schur function.

Proposition 3.3 ([Ges84]). *The Schur function $s_\lambda(x)$, $|\lambda| = n$, can be expressed in terms of quasi-symmetric functions by*

$$s_\lambda(x) = \sum_{T \in \text{SYT}(\lambda)} Q_{n, \sigma(T)}(x).$$

The fact that $g_\lambda = s_\lambda$ shows that the generating function of any vertex-signed graph whose connected components are isomorphic to graphs \mathcal{G}_λ is automatically Schur positive.

3.2 Axiomatization of dual equivalence

The purpose of this section is to characterize when a given edge-colored graph \mathcal{G} with signed vertices has connected components which look like the graphs of Section 3.1. We begin by specifying the data needed to form a signed, colored graph of type (n, N) .

Definition 3.4. A signed, colored graph of type (n, N) consists of the following data:

- a vertex set V ;
- a signature function $\sigma : V \rightarrow \{\pm 1\}^{N-1}$; and
- for each $1 < i < n$, a collection E_i of pairs of distinct vertices of V .

We denote such a graph by $\mathcal{G} = (V, \sigma, E)$,

Definition 3.5. For $m \leq n$ and $M \leq N$, the (m, M) -restriction of a signed, colored graph \mathcal{G} of type (n, N) consists of the vertex set V , signature function $\sigma : V \rightarrow \{\pm 1\}^{M-1}$ obtained by truncating σ at $M - 1$, and edge set $E = E_2 \cup \dots \cup E_{m-1}$.

Definition 3.6 presents a list of 5 axioms which a signed, colored graph must satisfy to be considered a dual equivalence graph. The axiomatization comes from analyzing the local properties of \mathcal{G}_λ solely in terms of signatures and edge colors.

Definition 3.6. A signed, colored graph $\mathcal{G} = (V, \sigma, E)$ of type (n, N) is a *dual equivalence graph of type (n, N)* if $n \leq N$ and the following hold:

(ax1) For $w \in V$ and $1 < i < n$, $\sigma(w)_{i-1} = -\sigma(w)_i$ if and only if there exists $x \in V$ such that $\{w, x\} \in E_i$. Moreover, x is unique when it exists.

(ax2) Whenever $\{w, x\} \in E_i$,


$$\begin{aligned} \sigma(w)_j &= -\sigma(x)_j & \text{for } j = i-1, i; \\ \sigma(w)_h &= \sigma(x)_h & \text{for } h < i-2 \text{ and } i+1 < h. \end{aligned}$$

(ax3) Whenever $\{w, x\} \in E_i$,


(a) if $\sigma(w)_{i-2} = -\sigma(x)_{i-2}$, then $\{w, x\} \in E_{i-1}$;

(b) if $\sigma(w)_{i+1} = -\sigma(x)_{i+1}$, then $\{w, x\} \in E_{i+1}$ if $i+1 < n$, and $\sigma(w)_i = -\sigma(x)_i$ if $i+1 = n$.

(ax4) For $3 < i < n$, every non-trivial connected component of the subgraph $(V, \sigma, E_{i-2} \cup E_{i-1} \cup E_i)$ is one of the following:



(ax5) Whenever $|i - j| \geq 3$, $\{w, x\} \in E_i$ and $\{x, y\} \in E_j$, there exists $v \in V$ such that $\{w, v\} \in E_j$ and $\{v, y\} \in E_i$.



Observe that axiom 1 implies that, in a dual equivalence graph, each vertex has at most one i -neighbor. Furthermore, axioms 1 and 2 together imply that there are at most two edges of different colors between a given pair of vertices, and that when this happens the colors of the edges are consecutive.

Remark 3.7. Every connected component of a dual equivalence graph of type (n, N) is again a dual equivalence graph of type (n, N) . Also, for $m \leq n, M \leq N$, the (m, M) -restriction of a dual equivalence graph of type (n, N) is a dual equivalence graph of type (m, M) .

The following proposition is the first step in justifying Definition 3.6. It also allows us to refer to \mathcal{G}_λ as the *standard dual equivalence graph corresponding to λ* .

Proposition 3.8. *For λ a partition of n , \mathcal{G}_λ is a dual equivalence graph of type (n, n) .*

Proof. For $T \in \text{SYT}(\lambda)$, $\sigma(T)_{i-1} = -\sigma(T)_i$ if and only if i does not lie between $i-1$ and $i+1$ in the content reading word of T . In this case, there exists $U \in \text{SYT}(\lambda)$ such that T and U differ by an elementary dual equivalence for $i-1, i, i+1$. Therefore U is obtained from T by swapping i with $i-1$ or $i+1$, whichever lies further away, with the result that $\sigma(T)_j = -\sigma(U)_j$ for $j = i-1, i$ and also $\sigma(T)_h = \sigma(U)_h$ for $h < i-2$ and $i+1 < h$. This verifies axioms 1 and 2.

For axiom 3, if $\sigma(T)_{i-2} = -\sigma(U)_{i-2}$ as well, then i and $i-1$ must have interchanged positions with $i-2$ lying between, so that T and U also differ by an elementary dual equivalence for $i-2, i-1, i$. A similar argument demonstrates the analogous situation for $i+1$. Given this characterization of double edges, axiom 4 becomes a straightforward check.

Finally, if $|i - j| \geq 3$, then $\{i-1, i, i+1\} \cap \{j-1, j, j+1\} = \emptyset$, so the elementary dual equivalences for $i-1, i, i+1$ and for $j-1, j, j+1$ commute with one another, thereby demonstrating axiom 5. \square

Remark 3.9. For partitions $\lambda \subset \mu$, with $|\lambda| = n$ and $|\mu| = N$, choose a tableau A of shape μ/λ with entries $n+1, \dots, N$. Define the set of standard Young tableaux of shape λ augmented by A , denoted $\text{ASYT}(\lambda, A)$, to be those $T \in \text{SYT}(\mu)$ such that T restricted to μ/λ is A . Let $\mathcal{G}_{\lambda, A}$ be the signed, colored graph of type (n, N) constructed on $\text{ASYT}(\lambda, A)$ with i -edges given by elementary dual equivalences for $i-1, i, i+1$ with $i < n$. Then $\mathcal{G}_{\lambda, A}$ is a dual equivalence graph of type (n, N) , and the (n, n) -restriction of $\mathcal{G}_{\lambda, A}$ is \mathcal{G}_λ .

The final justification of our axiomatization comes in the form of Theorem 3.11 below, which is, in a sense, the converse of Proposition 3.8. Before stating the theorem, we first define an

isomorphism of signed, colored graphs which, simply put, is a bijection on vertex sets which respects signatures and color-adjacency.

Definition 3.10. Two signed, colored graphs of type (n, N) , say $\mathcal{G} = (V, \sigma, E)$ and $\mathcal{H} = (W, \tau, F)$, are *isomorphic* if there exists a bijection $\phi : V \rightarrow W$ such that

- $\sigma(v)_i = \tau(\phi(v))_i$ for every $v \in V$, $1 \leq i < N$, and
- $\{u, v\} \in E_i$ if and only if $\{\phi(u), \phi(v)\} \in F_i$ for every $u, v \in V$, $1 < i < n$.

While the definition of an isomorphism between two signed, colored graphs assumes that they have the same type, we will often restrict one or both of the graphs we wish to compare in order to obtain the desired isomorphism.

Theorem 3.11. *Every connected component of a dual equivalence graph of type (n, n) is isomorphic to \mathcal{G}_λ for a unique partition λ of n .*

Theorem 3.11 will be extremely useful in the sections to follow, because it provides a concrete way to understand the structure of a dual equivalence graph by appealing to the known structure of the standard dual equivalence graph on tableaux.

We conclude this section by interpreting Theorem 3.11 in terms of symmetric functions. For a given signed, colored graph \mathcal{G} of type (n, N) for which every vertex is assigned some additional statistic α , we define the generating function $G(x; q)$ by

$$G(x; q) = \sum_{v \in V(\mathcal{G})} q^{\alpha(v)} Q_{N, \sigma(v)}(x). \quad (3.4)$$

We can, of course, include multivariate statistics, but as our immediate purpose is to apply this theory to LLT polynomials, a single parameter will suffice.

Theorem 3.11 and Proposition 3.3 together give a criterion implying that $G(x; q)$ is symmetric and Schur positive, and establish a combinatorial interpretation of the Schur expansion.

Corollary 3.12. *Let \mathcal{G} be a dual equivalence graph of type (n, n) with a vertex statistic α which is constant on connected components of \mathcal{G} . Let $C(\lambda)$ denote the set of connected components of \mathcal{G} which are isomorphic to \mathcal{G}_λ . Then*

$$G(x; q) = \sum_{\lambda} \sum_{c \in C(\lambda)} q^{\alpha(c)} s_{\lambda}(x).$$

3.3 Proof of Theorem 3.11

In order to avoid cumbersome notation, as we investigate the connection between arbitrary dual equivalence graphs and the standard one, we will often abuse notation by simultaneously referring to σ, E as the signatures, edge sets for both graphs.

The following result shows that a slightly stronger version of Theorem 3.11 holds. Recall the notion of augmenting a partition λ by a skew tableau A and the resulting dual equivalence graph $\mathcal{G}_{\lambda,A}$ from Remark 3.9. Then once Theorem 3.11 has been proved, it may be restated as: Let \mathcal{G} be a connected dual equivalence graph of type (n, N) , then \mathcal{G} is isomorphic to $\mathcal{G}_{\lambda,A}$ for a unique partition λ and some skew tableau A of shape μ/λ , $|\mu| = N$, with entries $n+1, \dots, N$.

Proposition 3.13. *Let $\mathcal{G} = (V, \sigma, E)$ be a connected dual equivalence graph of type (n, N) , and let ϕ be an isomorphism from the (n, n) -restriction of \mathcal{G} to \mathcal{G}_λ for some partition λ of n . Then there exists a semi-standard tableau A of shape μ/λ , $|\mu| = N$, with entries $n+1, \dots, N$ such that ϕ gives an isomorphism from \mathcal{G} to $\mathcal{G}_{\lambda,A}$. Moreover, the position of the cell of A containing $n+1$ is unique.*

Proof. By axiom 2 and the hypothesis that \mathcal{G} is connected, σ_n is constant on \mathcal{G} for $h \geq n+1$. Therefore once a suitable cell for $n+1$ has been chosen, the cells for $n+2, \dots, N$ may be chosen in any way which gives the correct signature. One solution is to place j north of the first column if $\sigma_{j-1} = -1$ or east of the first row if $\sigma_{j-1} = +1$ for $j = n+2, \dots, N$. Therefore it suffices to find the unique position for the cell containing $n+1$.

First consider the case when σ_n is constant on \mathcal{G} . Then we require σ_n to be constant with the same value on $\mathcal{G}_{\lambda,A}$ as well. Since n may occupy any northeastern corner of λ for some vertex of \mathcal{G}_λ by Proposition 3.2, the only way for σ_n to remain constant is for $n+1$ to be placed north of the first column or east of the first row of λ . The former ensures that $\sigma_n \equiv -1$, and the latter ensures $\sigma_n \equiv +1$. Therefore there is a unique position for $n+1$ in A which gives the correct signature.

Now assume that σ_n is not constant on all of \mathcal{G} . By dual equivalence axiom 2, σ_n is constant on connected components of the $(n-1, N)$ -restriction of \mathcal{G} . By Proposition 3.2, a connected component of the $(n-1, n-1)$ -restriction of \mathcal{G}_λ consists of all standard Young tableaux where n lies in a particular northeastern cell of λ , that is, at the eastern end of a row i where $\lambda_i > \lambda_{i+1}$. Let \mathcal{C}_i be the connected component of the $(n-1, N)$ -restriction of \mathcal{G} which corresponds via ϕ^{-1} to those tableaux with n at the eastern end of row i .

We claim that if $\sigma(\mathcal{C}_i)_n = -1$, then $\sigma(\mathcal{C}_j)_n = -1$ for all $j < i$, and that if $\sigma(\mathcal{C}_i)_n = +1$, then $\sigma(\mathcal{C}_k)_n = +1$ for all $k > i$. Given this, there exists a unique row i such that $\sigma(\mathcal{C}_j)_n = -1$ and $\sigma(\mathcal{C}_k)_n = +1$ for $j \leq i < k$. In this case, the cell containing $n+1$ must be placed at the eastern end of row $i+1$, and doing so extends ϕ to an isomorphism between $(n, n+1)$ graphs.




Figure 3.1: T and $U = E_{n-1}(T)$.

To prove the claim, let i, j be indices such that $\lambda_j > \lambda_{j+1} = \dots = \lambda_i > \lambda_{i+1}$, and suppose,

to the contrary, that $\sigma(\mathcal{C}_i)_n = -1$ and $\sigma(\mathcal{C}_j)_n = +1$. Let $T \in \text{SYT}(\lambda)$ be some tableau with n in row i and $n-1$ in row j . Then $n-2$ may lie immediately south of n whenever $j+1 < i$ or $\lambda_j - \lambda_i = 1$, and $n-2$ may lie immediately west of $n-1$ whenever $j+1 = i$ or $\lambda_j - \lambda_i > 1$. Since at least one option is always possible, we may choose T so that $n-2$ lies between $n-1$ and n in the content reading word, as in Figure 3.1.

Let U be the result of an elementary dual equivalence on T for $n-2, n-1, n$, which will necessarily interchange $n-1$ and n . Let $w = \phi^{-1}(T)$ and $x = \phi^{-1}(U)$, which are vertices of \mathcal{C}_i and \mathcal{C}_j , respectively. Then $\sigma(w)_{n-2, n-1, n} = + - -$ and $\sigma(x)_{n-2, n-1, n} = - + +$. However, since ϕ is an isomorphism of signed, colored graphs of type (n, n) , $\{w, x\} \in E_{n-1}$, which contradicts axiom 3(b). \square

The following lemma shows that any two connected components of the $(n-1, N)$ -restriction of a connected dual equivalence graph \mathcal{G} of type (n, N) are connected in \mathcal{G} by an E_{n-1} edge whenever the (n, n) -restriction of \mathcal{G} is isomorphic to a standard dual equivalence graph.

Lemma 3.14. *For any tableaux $T, U \in \text{SYT}(\lambda)$, $|\lambda| = n$, there exists a sequence of edges in \mathcal{G}_λ connecting T to U which contains at most one E_{n-1} edge.*

Proof. Let \mathcal{C}_T (resp. \mathcal{C}_U) denote the connected component of the $(n-1, n-1)$ -restriction of \mathcal{G}_λ containing T (resp. U). Let μ (resp. ν) be the shape of T (resp. U) with the cell containing n removed. Then $\mathcal{C}_T \cong \mathcal{G}_\mu$ and $\mathcal{C}_U \cong \mathcal{G}_\nu$.

If $\mu = \nu$, then, by Proposition 3.2, $\mathcal{C}_T = \mathcal{C}_U$ and the result follows. Assume, then, that $\mu \neq \nu$. Since $\mu, \nu \subset \lambda$ and $|\mu| = |\nu| = |\lambda| - 1$, both cells λ/μ and λ/ν must be northeastern corners of λ . Therefore, by the same argument as in the proof of Proposition 3.13, there exists $T' \in \text{SYT}(\lambda)$ with $n-1$ in position λ/ν , n in position λ/μ , and $n-2$ between $n-1$ and n in the content reading word of T' . Let U' be the result of swapping $n-1$ and n in T' , in particular, $\{T', U'\} \in E_{n-1}$.

By Proposition 3.2, T' is in \mathcal{C}_T and U' is in \mathcal{C}_U , hence there exists a sequence of edges from T to T' and a sequence from U' to U each consisting only of edges E_h , $h < n-1$. Since $\{T', U'\} \in E_{n-1}$, the result follows. \square

Next we will show that the \mathcal{G}_λ are mutually non-isomorphic and have no nontrivial automorphisms, the former of which justifies the uniqueness claim of Theorem 3.11. Both results stem from the same observation that \mathcal{G}_λ contains a unique vertex for which the composition formed by the lengths of the runs of $+1$'s in the signature gives the largest possible partition in dominance order.

Proposition 3.15. *Let λ, μ be partitions. Then $\text{Aut}(\mathcal{G}_\lambda) = \{\text{id}\}$, and if $\mathcal{G}_\lambda \cong \mathcal{G}_\mu$, then $\lambda = \mu$.*

Proof. Let $T_0 \in \text{SYT}(\lambda)$ be the tableau obtained by filling the numbers 1 through n into the rows of λ from left to right, bottom to top, in which case $\sigma(T_0) = +^{\lambda_1-1}, -, +^{\lambda_2-1}, -, \dots$. See Figure 3.2 for an example. For any standard tableau U such that $\sigma(U) = \sigma(T_0)$, the numbers 1 through λ_1 , and

also $\lambda_1 + 1$ through $\lambda_1 + \lambda_2$, and so on, must form horizontal strips. In particular, if $\sigma(U) = \sigma(T_0)$ for some U of shape μ , then $\lambda \leq \mu$ with equality if and only if $U = T_0$.




Figure 3.2: T_0 for $\lambda = (5, 4, 4, 1)$

For $\phi \in \text{Aut}(\mathcal{G}_\lambda)$ and T_0 as defined above, $\phi(T_0) = T_0$. For $U \in \text{SYT}(\lambda)$ such that $\{T_0, U\} \in E_i$, we have $\{T_0, \phi(U)\} \in E_i$, so $\phi(U) = U$ by dual equivalence axiom 1. Extending this, every tableau connected to a fixed point by some sequence of edges is also a fixed point for ϕ , hence $\phi = \text{id}$ on each \mathcal{G}_λ by Proposition 3.2.

Suppose $\phi : \mathcal{G}_\lambda \rightarrow \mathcal{G}_\mu$ is an isomorphism. Let T_0 be as above for λ , and let U_0 be the corresponding tableau for μ . Then since $\sigma(\phi(T_0)) = \sigma(T_0)$, $\lambda \leq \mu$. Conversely, since $\sigma(\phi^{-1}(U_0)) = \sigma(U_0)$, $\mu \leq \lambda$. Therefore $\lambda = \mu$. □

Stepping back from \mathcal{G}_λ for a moment, we next explore the structure of dual equivalence graphs of type (n, N) when $n < N$. To do this, we must first introduce some terminology.

Definition 3.16. Let $\mathcal{G} = (V, \sigma, E)$ be a dual equivalence graph of type (n, N) . For $1 < i < N$, we say that a vertex $w \in V$ admits an i -neighbor if $\sigma(w)_{i-1} = -\sigma(w)_i$.

The motivation for Definition 3.16 is dual equivalence axiom 1. For $1 < i < n$, if $\sigma(w)_{i-1} = -\sigma(w)_i$ for some $w \in V$, then axiom 1 implies the existence of $x \in V$ such that $\{w, x\} \in E_i$. That is, if w admits an i -neighbor for some $1 < i < n$, then w has an i -neighbor in \mathcal{G} .

For $n \leq i < N$, though i -edges do not exist in \mathcal{G} , if \mathcal{G} were the restriction of a dual equivalence graph of type $(i+1, N)$, then the condition $\sigma(w)_{i-1} = -\sigma(w)_i$ would imply the existence of a vertex x for which $\{w, x\} \in E_i$ in the type $(i+1, N)$ graph.

Lemma 3.17. For $\mathcal{G} = (V, \sigma, E)$ a dual equivalence graph of type (n, N) , the connected component of $(V, \sigma, E_{n-2} \cup E_{n-1})$ of a vertex which admits an n -neighbor is one of the cases depicted in Figure 3.3.

Remark 3.18. If $\{w, x\} \in E_{n-1}$ and both w and x admit an n -neighbor, then if \mathcal{G} is the (n, N) -restriction of a dual equivalence graph of type $(n+1, N)$, axiom 3(b) forces $\{w, x\} \in E_n$ in the type $(n+1, N)$ graph. This motivates the dashed line connecting the two leftmost vertices of B in Figure 3.3.




Figure 3.3: Possible connected components of $(V, \sigma, E_{n-2} \cup E_{n-1})$ where dashed lines indicate those vertices admitting an n -neighbor.

For the remaining vertices which admit an n -neighbor, for each of A, B, and C, there is exactly one vertex which also admits an $n-1$ -neighbor, and exactly one which does not. Two identical components of C are shown in Figure 3.3-C in order to parallel axiom 4.

Proof. Let $w \in V$ admit an n -neighbor. If w admits neither an $n-2$ -neighbor nor an $n-1$ -neighbor, then w is the right-hand side of A.

If w admits an $n-1$ -neighbor, say $\{w, v\} \in E_{n-1}$, but does not admit an $n-2$ -neighbor, then by axiom 2, v must admit an $n-2$ -neighbor, say $\{v, u\} \in E_{n-2}$. If v admits an n -neighbor then u must also admit an n -neighbor by axiom 2, and w, v, u form the left-hand side of B. If v does not admit an n -neighbor, then by axiom 2 neither does u , so u, v, w form the left-hand side of A.

If w admits an $n-2$ -neighbor, say $\{w, v\} \in E_{n-2}$, but does not admit an $n-1$ -neighbor, then v must admit an $n-1$ -neighbor by axiom 2, say $\{v, u\} \in E_{n-1}$. Again, by axiom 2, since w admits an n -neighbor, so too must v . If u admits an n -neighbor as well, then u, v, w form the left-hand side of B. If u does not admit an n -neighbor, then u, v, w are as in C.

Finally, consider the case when w admits both an $n-1$ -neighbor and an $n-2$ -neighbor, say $\{w, v\} \in E_{n-1}$ and $\{w, u\} \in E_{n-2}$. If $u = v$, then v does not admit an n -neighbor, and w, v are as the right-hand side of B. Otherwise, u must also admit an n -neighbor by axiom 2 and may not admit an $n-1$ -neighbor by axiom 3. Similarly, v may not admit an $n-2$ -neighbor by axiom 3. If v admits an n -neighbor, then v, w, u are as the left-hand side of B, and if v does not admit an n -neighbor, then v, w, u are as in C. \square

Lemma 3.17 justifies defining the n -type of a vertex admitting an n -neighbor as A, B or C according to the characterization of Figure 3.3. We will often treat separately the two leftmost vertices of B for reasons discussed in Section 3.4. If v admits an n -neighbor but is not one of these two cases, then the connected component of $(V, \sigma, E_{n-2} \cup E_{n-1})$ of v is completely determined by the n -type of v and whether or not v admits an $n-1$ -edge.

Remark 3.19. Let \mathcal{G} be a dual equivalence graph of type (n, N) . Let w and x be vertices such that

$\{w, x\} \in E_i$ and $\{w, x\} \notin E_{i-1}$ for some $i < n$. We may define the i -type of w and x by considering the (i, N) -restriction of \mathcal{G} . Note that the condition $\{w, x\} \notin E_{i-1}$ rules out the possibility that w and x are the two leftmost vertices of B in Figure 3.3. By dual equivalence axioms 1, 2 and 3 and the preceding remarks, $\sigma(w)_h$ and the i -type of w are completely determined by $\sigma(x)_h$ and the i -type of x for $h \leq i$.

The following result is, in essence, the inductive step for the proof of Theorem 3.11. However, since the result will be useful in later sections, we present it independently.

Proposition 3.20. *Let $\mathcal{G} = (V, \sigma, E)$ be a connected dual equivalence graph of type (n, n) . Let \mathcal{C} be a connected component of the $(n-1, n)$ -restriction of \mathcal{G} , which, by Remark 3.7, is a dual equivalence graph of type $(n-1, n)$. Assume there exists a partition μ of $n-1$ such that the $(n-1, n-1)$ -restriction of \mathcal{C} is isomorphic to \mathcal{G}_μ . Then there exists a unique partition λ of n with $\mu \subset \lambda$ such that $\mathcal{G} \cong \mathcal{G}_\lambda$.*

Proof. Let ϕ be the isomorphism from the $(n-1, n-1)$ -restriction of \mathcal{C} to \mathcal{G}_μ . By Proposition 3.13, there exists a unique partition λ of n and a unique tableau A of shape λ/μ with entry n such that ϕ gives an isomorphism from \mathcal{C} to $\mathcal{G}_{\mu, A}$. We will show that $\mathcal{G} \cong \mathcal{G}_\lambda$ for this particular λ .

If no vertex of \mathcal{C} admits an $n-1$ -neighbor, then no vertex of $\mathcal{G}_{\mu, A}$ admits an $n-1$ -neighbor. Therefore $n-1$ must always lie between $n-2$ and n in the content reading word of any tableaux in $\text{ASYT}(\mu, A)$. Therefore λ is a single row or column, and so $\mathcal{G}_{\mu, A} = \mathcal{G}_\lambda$. Since \mathcal{G} is assumed to be connected and \mathcal{C} is a single vertex which admits no neighbors, $\mathcal{C} = \mathcal{G}$. Therefore ϕ gives an isomorphism from \mathcal{G} to \mathcal{G}_λ . In particular, the result holds for $n = 2$, so we may proceed by induction on n .

Choose a vertex w of \mathcal{C} which admits an $n-1$ -neighbor, say $\{w, x\} \in E_{n-1}$ for some vertex x of \mathcal{G} . Let \mathcal{D} be the connected component of the $(n-1, n)$ -restriction of \mathcal{G} which contains x . Let $T = \phi(w)$, $T \in \text{ASYT}(\mu, A)$. Since ϕ is an isomorphism of $(n-1, n)$ graphs, T admits an $n-1$ -neighbor, so let U be such that $\{T, U\} \in E_{n-1}$ in \mathcal{G}_λ . Let ν be the shape of U with the cell containing n removed, and let B be the tableau of shape λ/ν with entry n . Then $U \in \text{ASYT}(\nu, B)$.

We will construct an isomorphism from \mathcal{D} to $\mathcal{G}_{\nu, B}$ which maps x to U . Given this, by Lemma 3.14, each connected component of the $(n-1, n)$ -restriction of \mathcal{G} will be isomorphic to a unique connected component of the $(n-1, n)$ -restriction of \mathcal{G}_λ , and conversely. By Proposition 3.15 \mathcal{G}_λ has no nontrivial automorphisms, so these isomorphisms must combine to give an isomorphism from \mathcal{G} to \mathcal{G}_λ .

We may assume $\mathcal{D} \neq \mathcal{C}$, the contrary case being trivial. By Proposition 3.15, this implies that $\nu \neq \mu$. Let \mathcal{C}_j (resp. \mathcal{D}_j , $(\mathcal{G}_{\mu, A})_j$, and $(\mathcal{G}_{\nu, B})_j$) denote the connected component of the (j, j) -restriction of \mathcal{C} (resp. \mathcal{D} , $\mathcal{G}_{\mu, A}$, and $\mathcal{G}_{\nu, B}$) containing w (resp. x , T , and U).

By dual equivalence axiom 2, σ_{n-2} and σ_{n-1} are constant on \mathcal{C}_{n-4} , \mathcal{D}_{n-4} , $(\mathcal{G}_\mu)_{n-4}$ and $(\mathcal{G}_\nu)_{n-4}$, therefore every vertex on these components must admit an $n-1$ -neighbor. By axioms 2 and 5, E_{n-1} may be regarded as an isomorphism from \mathcal{C}_{n-4} to \mathcal{D}_{n-4} which maps w to x , and similarly from $(\mathcal{G}_\mu)_{n-4}$ to $(\mathcal{G}_\nu)_{n-4}$ mapping T to U . Let ψ be the conjugate of ϕ by the isomorphism given

in this way by E_{n-1} , see Figure 3.4. Then ψ is an isomorphism from \mathcal{D}_{n-4} to $(\mathcal{G}_\nu)_{n-4}$. We will use induction and Remark 3.19 to extend ψ to an isomorphism from \mathcal{D} to $\mathcal{G}_{\nu,B}$.




Figure 3.4: The construction of $\psi : \mathcal{D}_{n-4} \rightarrow (\mathcal{G}_\nu)_{n-4}$.

Let y be a vertex of \mathcal{D}_{n-4} and let v be the vertex of \mathcal{C}_{n-4} such that $\{v, y\} \in E_{n-1}$ in \mathcal{G} . Then, by the definition of ψ , $\{\phi(v), \psi(y)\} \in E_{n-1}$ in \mathcal{G}_λ .

By the assumption that $\mathcal{C} \neq \mathcal{D}$, $\{v, y\}, \{\phi(v), \psi(y)\} \notin E_{n-2}$, so Remark 3.19 may be applied to y and also to $\psi(y)$. Since ϕ is an isomorphism of type $(n-1, n)$ graphs, $\sigma(v)_h = \sigma(\phi(v))_h$ for $h \leq n-1$ and the $n-1$ -type of v is the $n-1$ -type of $\phi(v)$. Combining these observations, $\sigma(y)_h = \sigma(\psi(y))_h$ for $h \leq n-1$ and the $n-1$ -type of y is the $n-1$ -type of $\psi(y)$. In particular, $\sigma(y)_{n-3} = \sigma(\psi(y))_{n-3}$. Therefore, by induction and the uniqueness clause of Proposition 3.13, ψ extends to an isomorphism from \mathcal{D}_{n-3} to $(\mathcal{G}_\nu)_{n-3}$.

In order to extend ψ to an isomorphism from \mathcal{D}_{n-2} to $(\mathcal{G}_\nu)_{n-2}$, we need only show $\sigma(z)_{n-2} = \sigma(\psi(z))_{n-2}$ for all vertices z of \mathcal{D}_{n-3} . This holds for all vertices of \mathcal{D}_{n-4} by the preceding argument. By Lemma 3.14, we may assume $\{y, z\} \in E_{n-3}$ for some vertex y of \mathcal{D}_{n-4} . Since ψ is an isomorphism of type $(n-3, n-3)$ graphs, $\{\psi(y), \psi(z)\} \in E_{n-3}$. Then $\sigma(z)_h = \sigma(y)_h$ (resp. $\sigma(\psi(z))_h = \sigma(\psi(y))_h$) for $h = n-2, n-1$ if and only if $\{y, z\} \notin E_{n-2}$ (resp. $\{\psi(y), \psi(z)\} \notin E_{n-2}$), which is determined by the $n-1$ -type of y (resp. $\psi(y)$). By the previous paragraph, $\sigma(y)_h = \sigma(\psi(y))_h$ for $h = n-2, n-1$ and the $n-1$ -type of y is the $n-1$ -type of $\psi(y)$. In particular, $\sigma(z)_{n-2} = \sigma(\psi(z))_{n-2}$, and ψ extends to an isomorphism from \mathcal{D}_{n-2} to $(\mathcal{G}_\nu)_{n-2}$.

Finally, to extend ψ to an isomorphism from \mathcal{D} to $\mathcal{G}_{\nu,B}$, we must show $\sigma(z)_{n-1} = \sigma(\psi(z))_{n-1}$ for all vertices z of \mathcal{D} . Since this holds for vertices of \mathcal{D}_{n-3} , we can mimic the previous argument to obtain the result. \square

Proof of Theorem 3.11. Let \mathcal{G} be a dual equivalence graph of type (n, n) . We will proceed by induction on n . When $n = 2$, the result follows immediately from Proposition 3.13 since \mathcal{G} has no edges. Now let us assume the result for dual equivalence graphs of type $(n-1, n-1)$. Let \mathcal{C} be a connected component of the $(n-1, n-1)$ -restriction of \mathcal{G} . By induction, there is a unique partition ν of $n-1$ such that $\mathcal{C} \cong \mathcal{G}_\nu$. By Proposition 3.20, there exists a unique partition λ of n with $\nu \subset \lambda$ such that $\mathcal{G} \cong \mathcal{G}_\lambda$. \square

3.4 Constructing edges by induction

The purpose of this section is to investigate how to extend a dual equivalence graph of type (i, N) to a dual equivalence graph of type $(i+1, N)$ by defining a suitable collection E_i of i -colored edges. We will use this idea of extending dual equivalence graphs in Chapter 4, where we will build a dual equivalence graph by induction.

Throughout this section, let $\mathcal{G} = (V, \sigma, E)$ be a dual equivalence graph of type (i, N) with $i < N$. For $h \leq i$, let $C_{h-2} \subset E_{h-2}$ denote the collection of those E_{h-2} edges which connect vertices of h -type C. The first task is to group together certain vertices which admit an i -neighbor.

Definition 3.21. Let w be a vertex of \mathcal{G} which admits an i -neighbor. Define the i -package of w , denoted \mathcal{P}_w , to be the connected component of $(V, \sigma, E_2 \cup \dots \cup E_{i-3} \cup C_{i-2})$ containing w , and define the *reduced i -package* of w , denoted \mathcal{P}'_w , to be the connected component of the $(i-2, N)$ -restriction of \mathcal{P}_w .

Note that \mathcal{P}_w is a signed, colored graph of type $(i-1, N)$, but is not, in general, a dual equivalence graph. However, since \mathcal{P}'_w is the connected component of the $(i-2, N)$ -restriction of \mathcal{G} containing w , \mathcal{P}'_w is a dual equivalence graph of type $(i-2, N)$.

Remark 3.22. By Proposition 3.15 and Theorem 3.11, there are no nontrivial automorphisms of a reduced i -package. Therefore a nontrivial automorphism of an i -package \mathcal{P}_w may not map any connected component of the $(i-2, N)$ -restriction of \mathcal{P}_w to itself. However, for λ a partition of $i-1$, the connected components of the $(i-2, i-1)$ -restriction of \mathcal{G}_λ are pairwise non-isomorphic by Proposition 3.15 and the uniqueness claim of Proposition 3.13. By Theorem 3.11, there exists a unique partition λ of $i-1$ such that the connected component of the $(i-1, i-1)$ -restriction of \mathcal{G} containing \mathcal{P}_w is isomorphic to \mathcal{G}_λ . Therefore the connected components of the $(i-2, N)$ -restriction of \mathcal{P}_w are pairwise non-isomorphic. Hence i -packages have no nontrivial automorphisms.

By dual equivalence axioms 2, 3 and 4 and the definition of i -type C, every vertex of \mathcal{P}_w will admit an i -neighbor. Furthermore, if there exists a collection E_i of pairs of vertices of V such that $(V, \sigma, E \cup E_i)$ is a dual equivalence graph of type $(i+1, N)$, then by dual equivalence axioms 4 and 5, knowing $\{v, y\} \in E_i$ for some $v \in \mathcal{P}_w$ and some $y \in \mathcal{P}_x$ determines E_i on all of \mathcal{P}_w and \mathcal{P}_x . With this in mind, constructing a suitable collection E_i amounts to choosing an i -neighbor for one vertex on each i -package and then extending this definition to the entire i -package. More precisely, we have the following definition.

Definition 3.23. Let w, x be vertices of \mathcal{G} which admit an i -neighbor, and let ϕ be an isomorphism from the $(i-1, i-2)$ -restriction of \mathcal{P}_w to the $(i-1, i-2)$ -restriction of \mathcal{P}_x which maps w to x . Let F be a collection of pairs of vertices of \mathcal{G} such that $\{w, x\} \in F$. Then F may be *extended along \mathcal{P}_w and \mathcal{P}_x at w and x* by adding to F the pair $\{v, \phi(v)\}$ for each vertex v of \mathcal{P}_w .

If F is a collection of pairs of vertices of \mathcal{G} which admit an i -neighbor such that for each pair $\{w, x\} \in F$ there exists an isomorphism $\phi_{w,x}$ from the $(i-1, i-2)$ -restriction of \mathcal{P}_w to the $(i-1, i-2)$ -

restriction of \mathcal{P}_x which maps w to x , then F may be *extended along i -packages* by extending F along \mathcal{P}_w and \mathcal{P}_x at w and x for every pair $\{w, x\} \in F$.

If ϕ and ψ are two isomorphisms from the $(i-1, i-2)$ -restriction of \mathcal{P}_w to the $(i-1, i-2)$ -restriction of \mathcal{P}_x which map w to x , then $\phi \circ \psi^{-1}$ gives an automorphism of \mathcal{P}_w . Therefore by Remark 3.22, $\phi \circ \psi^{-1}$ must be the identity. By symmetry, $\phi = \psi$. In particular, an extension of F along an i -package is independent of the isomorphism used.

The condition that there exists an isomorphism from the $(i-1, i-2)$ -restriction of \mathcal{P}_w to the $(i-1, i-2)$ -restriction of \mathcal{P}_x which maps w to x is necessary for E_i to be extended along i -packages, but it is not sufficient to ensure that $(V, \sigma, E \cup E_i)$ will be a dual equivalence graph of type $(i+1, N)$. The following definition characterizes when two i -packages may be paired by E_i in a way that respects the dual equivalence axioms.

Definition 3.24. Two i -packages \mathcal{P}_w and \mathcal{P}_x are *i -compatible at w and x* if there exists an isomorphism ϕ from the $(i-1, i-2)$ -restriction of \mathcal{P}_w to the $(i-1, i-2)$ -restriction of \mathcal{P}_x mapping w to x such that for every vertex v of \mathcal{P}_w :

- (c2) $\sigma(v)_j = -\sigma(\phi(v))_j$ for $j = i-1, i$, and
 $\sigma(v)_h = \sigma(\phi(v))_h$ for $i+1 < h$;
- (c3) if $\sigma(v)_{i-2} = -\sigma(\phi(v))_{i-2}$, then $\{v, \phi(v)\} \in E_{i-1}$, and
if $\sigma(v)_{i+1} = -\sigma(\phi(v))_{i+1}$, then v and $\phi(v)$ admit an $i+1$ -neighbor;
- (c4) if v has i -type C then $\phi(v)$ has i -type C.

Similarly, two reduced i -packages \mathcal{P}'_w and \mathcal{P}'_x are *i -compatible at w and x* if there exists an isomorphism ϕ from the $(i-2, i-2)$ -restriction of \mathcal{P}'_w to the $(i-2, i-2)$ -restriction of \mathcal{P}'_x mapping w to x which satisfies the above conditions.

If \mathcal{G} is a dual equivalence graph of type $(i+1, N)$ and $\{w, x\} \in E_i$, then we can define a map $\phi : \mathcal{P}_w \rightarrow \mathcal{P}_x$ by $\phi(v) = y$ when $\{v, y\} \in E_i$. Then ϕ is an i -compatibility maps from \mathcal{P}_w to \mathcal{P}_x at w and x . In Proposition 3.29, we will prove a converse of this statement. Before doing so, we have a few more results to establish.

There are two kinds of i -packages which it is important to distinguish.

Definition 3.25. An i -package \mathcal{P} is called *internal* if there exists some $w \in \mathcal{P}$ such that $\{w, x\} \in E_{i-1}$ and $\sigma(w)_i = -\sigma(x)_i$ for some vertex x . Otherwise \mathcal{P} is called *external*.

The motivation for this terminology is that if \mathcal{P}_w is internal, say with $\{w, x\} \in E_{i-1}$ such that $\sigma(w)_i = -\sigma(x)_i$, then any collection E_i of i -colored edges which hopes to satisfy the dual equivalence axioms is forced to have $\{w, x\} \in E_i$ by axiom 3(b). As remarked before, this will force the i -neighbors for all of \mathcal{P}_w , and so E_i will pair vertices from the connected component of \mathcal{G} containing w . Furthermore, if \mathcal{P}_w and \mathcal{P}_v are i -compatible at w and v for some vertex v of \mathcal{G} , then

by condition c3 of Definition 3.24, we must have $v = x$. Given this, it is reasonable to ask if \mathcal{P}_w and \mathcal{P}_x are i -compatible at w and x . Happily, this is the case.

Lemma 3.26. *If $\{w, x\} \in E_{i-1}$ is such that $\sigma(w)_i = -\sigma(x)_i$, then \mathcal{P}_w and \mathcal{P}_x are i -compatible at w and x .*

Proof. By Theorem 3.11, there exists an isomorphism, say ψ , from the connected component of \mathcal{G} containing w (and x) to $\mathcal{G}_{\lambda, A}$ for some partition λ of i and some tableau A of shape μ/λ , $|\mu| = i+1$, with entry $i+1$. Dual equivalence axiom 3(b) ensures that $\{\psi(w), \psi(x)\} \in E_{i-1} \cap E_i$ in \mathcal{G}_μ . As mentioned previously, this implies that $\mathcal{P}_{\psi(w)}$ and $\mathcal{P}_{\psi(x)}$ are i -compatible at $\psi(w)$ and $\psi(x)$, say via an isomorphism ϕ . Therefore $\psi^{-1}\phi\psi$ establishes the i -compatibility of \mathcal{P}_w and \mathcal{P}_x at w and x . \square

Remark 3.27. We can use the isomorphism between the connected component of \mathcal{G} containing w and $\mathcal{G}_{\lambda, A}$ to say a bit more about internal i -packages. Let w be such that $\{w, x\} \in E_{i-1}$ and $\sigma(w)_i = -\sigma(x)_i$ for some $x \in V$. Let $T \in \text{ASYT}(\lambda, A)$ correspond to w . Then $i+1$ must lie between i and $i-1$ in the content reading word of T , and the positions of i and $i+1$ are fixed on \mathcal{P}_T . Moreover, it's an easy check that under these circumstances, a C_{i-2} edge on \mathcal{P}_T may only act by interchanging $i-3$ and $i-2$. Therefore the position of $i-1$ is also fixed on \mathcal{P}_T . By Proposition 3.2, this implies that every vertex of the i -package \mathcal{P}_w actually lies on the reduced i -package \mathcal{P}'_w .

Proposition 3.28. *Let w and x be vertices of \mathcal{G} such that \mathcal{P}'_w and \mathcal{P}'_x are i -compatible at w and x . Then \mathcal{P}_w and \mathcal{P}_x are i -compatible at w and x .*

Proof. Let ϕ be the isomorphism from the $(i-2, i-2)$ -restriction of \mathcal{P}'_w to the $(i-2, i-2)$ -restriction of \mathcal{P}'_x . By Remark 3.27, if \mathcal{P}_w is an internal i -package, then every vertex of \mathcal{P}_w is a vertex of \mathcal{P}'_w . Therefore by the definition of an internal i -package, there exists some vertex v of \mathcal{P}'_w such that $\{v, y\} \in E_{i-1}$ and $\sigma(v)_i = -\sigma(y)_i$ for some vertex y of \mathcal{G} . By condition c3 of Definition 3.24, we must have $\phi(v) = y$, in which case \mathcal{P}_x is also internal. By Lemma 3.26, we may conclude that \mathcal{P}_w and \mathcal{P}_x are i -compatible at w and x .

Assume, then, that \mathcal{P}_w and \mathcal{P}_x are both external i -packages. In this case, condition c3 of Definition 3.24 ensures that $\sigma(v)_{i-2} = \sigma(\phi(v))_{i-2}$ for every vertex v of \mathcal{P}'_w . Therefore by Proposition 3.20, ϕ extends to an isomorphism from the connected component of the $(i-1, i-1)$ -restriction of \mathcal{G} containing w to the one containing x . We will show that ϕ restricted to \mathcal{P}_w establishes the i -compatibility of \mathcal{P}_w and \mathcal{P}_x at w and x .

By Lemma 3.14, every vertex of the connected component of the $(i-1, i-1)$ -restriction of \mathcal{G} containing w (resp. x) which does not lie on \mathcal{P}'_w (resp. \mathcal{P}'_x) lies on a connected component of the $(i-2, i-2)$ -restriction of \mathcal{G} containing some vertex v such that $\{v, u\} \in E_{i-2}$ for some vertex u of \mathcal{P}'_w (resp. \mathcal{P}'_x). Since i -packages are connected, this implies that every vertex of \mathcal{P}_w (resp. \mathcal{P}_x) which is not a vertex of \mathcal{P}'_w (resp. \mathcal{P}'_x) lies on the reduced i -package of a vertex v such that $\{v, u\} \in C_{i-2}$ for some vertex u of \mathcal{P}'_w (resp. \mathcal{P}'_x). Since u has i -type C if and only if $\phi(u)$ has i -type C for every

vertex u of \mathcal{P}'_w , $\phi(v)$ is a vertex of \mathcal{P}_x for every vertex v of \mathcal{P}_w . Therefore ϕ restricts to a map from \mathcal{P}_w to \mathcal{P}_x .

Since i -packages include only those E_{i-2} edges which connect vertices with i -type C, dual equivalence axioms 2, 3 and 4 and the definition of i -type C ensure that σ_h is constant on external i -packages for all $h \geq i-1$. Using the assumption that \mathcal{P}'_w and \mathcal{P}'_x are i -compatible at w and x , for every vertex v of \mathcal{P}_w we have $\sigma(v)_j = -\sigma(\phi(v))_j$ for $j = i-1, i$ and $\sigma(v)_h = \sigma(\phi(v))_h$ for $i+1 < h$, thereby establishing condition c2 of Definition 3.24. Furthermore, if $\sigma(v)_{i+1} = -\sigma(\phi(v))_{i+1}$, then v and $\phi(v)$ admit an $i+1$ -neighbor. Also, since ϕ is an isomorphism of $(i-1, i-1)$ graphs, $\sigma(v)_{i-2} = \sigma(\phi(v))_{i-2}$. Therefore condition c3 is also established.

Let v be a vertex of \mathcal{P}_w , and let $y = \phi(v)$ be the corresponding vertex of \mathcal{P}_x . Since \mathcal{P}_w and \mathcal{P}_x are assumed to be external i -packages, neither v nor y can be one of the two the leftmost vertices of i -type B as depicted in Figure 3.3. If v has i -type C, then v admits an $i-2$ -neighbor. Since $\sigma(v)_j = \sigma(y)_j$ for $j = i-3, i-2$, y also admits an $i-2$ -neighbor. Therefore if y does not have i -type C, then it must be one of the two rightmost vertices as depicted in Figure 3.3. Let u, z be vertices of \mathcal{G} such that $\{v, u\}, \{y, z\} \in E_{i-2}$. Since v has i -type C, u is a vertex of \mathcal{P}_w , and so $z = \phi(u)$ is a vertex of \mathcal{P}_x . Inspecting Figure 3.3, if y does not admit an $i-1$ -neighbor, then z lies on an internal i -package which contradicts that assumption that \mathcal{P}_x is external. If y admits an $i-1$ -neighbor, then z does not admit an i -neighbor which contradicts the fact that every vertex of an i -package admits an i -neighbor. Therefore y must also have i -type C, proving the final condition. \square

The main result of this section is the following proposition which presents a method for defining a suitable collection E_i of i -colored edges which extends \mathcal{G} to a dual equivalence graph of type $(i+1, N)$.

Proposition 3.29. *Let W be a subset of vertices of \mathcal{G} which admit an i -neighbor such that every i -package of \mathcal{G} contains at least one vertex of W . Let $\varphi : W \rightarrow V$ be a map satisfying the following conditions for all $v, w \in W$:*

1. φ is an involution on W ;
2. v and w lie on the same i -package if and only if $\varphi(v)$ and $\varphi(w)$ lie on the same i -package;
3. \mathcal{P}'_w and $\mathcal{P}'_{\varphi(w)}$ are i -compatible at w and $\varphi(w)$.

Define E_i to be the set of pairs $\{w, \varphi(w)\}$ for $w \in W$. Then E_i may be extended along i -packages, and doing so makes $(V, \sigma, E \cup E_i)$ into a dual equivalence graph of type $(i+1, N)$.

Proof. By condition 3 and Proposition 3.28, for each $w \in W$, \mathcal{P}_w and $\mathcal{P}_{\varphi(w)}$ are i -compatible at w and $\varphi(w)$, say via the map ϕ_w . Therefore by Definition 3.23, E_i may be extended along i -packages by including $\{v, \phi_w(v)\}$ in E_i for each vertex w of W and each vertex v of \mathcal{P}_w . It remains to show that the result is a dual equivalence graph of type $(i+1, N)$.

By assumption, each i -package contains at least one vertex of W . Since E_i is extended along i -packages, for each $u \in V$, u admits an i -edge if and only if there exists $x \in V$ such that $\{u, x\} \in E_i$. Therefore to prove dual equivalence axiom 1, we must show that whenever x exists, it is unique.

For any vertex y of $\mathcal{P}_{\varphi(w)}$, $y \neq \varphi(w)$, \mathcal{P}_w and $\mathcal{P}_{\varphi(w)}$ are not i -compatible at w and y since this would imply the existence of a nontrivial automorphism of $\mathcal{P}_{\varphi(w)}$, contradicting Remark 3.22. Let $v, w \in W$ be such that v and w lie on the same i -package, i.e. $\mathcal{P}_v = \mathcal{P}_w$. By condition 2, $\varphi(v)$ and $\varphi(w)$ also lie on the same i -package, i.e. $\mathcal{P}_{\varphi(v)} = \mathcal{P}_{\varphi(w)}$. By condition 3 applied to v , \mathcal{P}_v and $\mathcal{P}_{\varphi(v)}$ are i -compatible at v and $\mathcal{P}_{\varphi(v)}$, and by condition 3 applied to w , \mathcal{P}_v and $\mathcal{P}_{\varphi(v)}$ are i -compatible at v and $\phi_w(v)$. Therefore $\varphi(v) = \phi_w(v)$, and so extending E_i along \mathcal{P}_w and $\mathcal{P}_{\varphi(w)}$ at w and $\varphi(w)$ gives the same result as extending E_i along \mathcal{P}_w and $\mathcal{P}_{\varphi(w)}$ at v and $\varphi(v)$. Since, by condition 1, $\varphi(\varphi(w)) = w$, this proves the desired uniqueness.

Dual equivalence axioms 2 and 3 follow immediately from conditions c2 and c3, respectively. Axiom 4 follows from axioms 2 and 3 and condition c4. Finally, axiom 5 is a consequence of the isomorphisms ϕ_w and the definition of extending E_i along i -packages. \square

Chapter 4

LLT polynomials

In order to demonstrate one of the uses of dual equivalence graphs, we present the following application to ribbon tableaux generating functions known as LLT polynomials. For $k \leq 3$, we will establish the existence of a dual equivalence graph structure on standard k -ribbon tableaux of a given (skew) shape for which the cospin statistic is constant on connected components. By Corollary 3.12, this will prove that the coefficient of s_λ in $\tilde{G}_\mu^{(k)}$ q -counts the number of connected components of the graph which are isomorphic to \mathcal{G}_λ .

4.1 The vertices and signature function

The vertices for the graph we wish to create are standard k -tuples of Young tableau, which arise as the k -quotients of standard k -ribbon tableaux as discussed in Section 2.4. For this to be a viable approach, we need to translate the cospin statistic on k -ribbon tableaux, defined in equation (2.12), into a statistic on k -tuples of tableaux. This was done in [SSW03, HHL⁺05b]. We will use the statistics presented in [HHL⁺05b].

Recall that the *content* of a cell x , denoted $c(x)$, indexes the diagonal on which it lies. To each piece of the k -quotient of λ is assigned a distinct integer modulo k , say (s_0, \dots, s_{k-1}) , with $s_i \equiv i \pmod{k}$, which gives the shifted content of a cell x of $\lambda^{(i)}$ by

$$\tilde{c}(x) = kc(x) + s_i .$$

Relating back to ribbon tableaux, the vector (s_0, \dots, s_{k-1}) is chosen so that the contents of the k -ribbons correspond precisely with the shifted contents of the cells of the k -quotient.

Let $T \in \text{SSYT}_k$. Let $T(x)$ denote the entry of the cell x in T . Define the *set of k -inversions of T* , denoted $\text{Inv}_k(T)$, by

$$\text{Inv}_k(T) = \{(x, y) \mid k > \tilde{c}(y) - \tilde{c}(x) > 0 \text{ and } T(x) > T(y)\}, \quad (4.1)$$

Then the k -inversion number of T , denoted $\text{inv}_k(T)$, is

$$\text{inv}_k(T) = |\text{Inv}_k(T)|. \quad (4.2)$$

The following proposition relates the spin of T , defined in equation (2.10), and the k -inversion number of T . In particular, Proposition 4.1 shows that an operation on k -tuples of tableaux preserves cospin if and only if it preserves the k -inversion number.

Proposition 4.1 ([HHL⁺05b]). *Given a (skew) shape μ , there exists a constant a such that for every $T \in \text{SYT}_k(\mu)$, we have $s(T) = a - \text{inv}_k(T)$.*

In the classical setting of tableaux, an elementary dual equivalence on a standard Young tableau is defined in terms of an elementary dual equivalence on a suitable reading word. Mimicking this idea for k -tuples of tableau, we need a reading word for which the k -inversion number is transparent.

Construct the *content reading word*, denoted $w(T)$, of a k -tuple of tableaux by defining the j th letter of $w(T)$ to be the set of entries of T with shifted content j . That is,

$$w(T)_j = \{T(x) \mid \tilde{c}(x) = j\}. \quad (4.3)$$

Since cells with the same shifted content come from the same diagonal of the same tableau, the letters of the content reading word of a semi-standard k -tuple of tableau are ordinary subsets of $[n]$ (i.e. not multi-sets).

We picture the word w in equation (4.3) as a sequence of columns, with the j th column displaying the members of $w(T)_j$ in increasing order from bottom to top, as in Example 4.2 below. If x is a position in one of the columns of w depicted in this way, let $w(x)$ denote the entry in position x , and let $\tilde{c}(x)$ denote the index of the column in which x lies. Analogous to equation (4.1), define the *set of k -inversions* of w by

$$\text{Inv}_k(w) = \{(x, y) \mid k > \tilde{c}(y) - \tilde{c}(x) > 0 \text{ and } w(x) > w(y)\}, \quad (4.4)$$

and define the k -inversion number of w by

$$\text{inv}_k(w) = |\text{Inv}_k(w)|. \quad (4.5)$$

As remarked previously, $T^{(i)}$ contributes all entries x with $\tilde{c}(x) \equiv i \pmod{k}$. Therefore since $w(T)$ records the content of each entry, the shape of $T^{(i)}$ may be recovered from the descent set of the $i \pmod{k}$ -sub-word of $w(T)$, i.e. the word u defined by $u_j = w(T)_{jk+i}$. With this motivation, we define the k -descent set, denoted Des_k , of a word w by

$$\text{Des}_k(w) = \{(x, y) \mid \tilde{c}(y) - \tilde{c}(x) = k \text{ and } w(x) > w(y)\}. \quad (4.6)$$

Let T, U be two k -tuples of tableaux with content reading words $w(T), w(U)$. We have established the following relationship between the tableaux and the reading words.

$$\tilde{s}(T) = \tilde{s}(U) \iff \text{inv}_k(w(T)) = \text{inv}_k(w(U)), \quad (4.7)$$

$$\text{shape}(T) = \text{shape}(U) \iff \text{Des}_k(w(T)) = \text{Des}_k(w(U)). \quad (4.8)$$

Example 4.2. Suppose T is the 4-tuple of tableaux

$$\begin{array}{ccc} \begin{array}{|c|c|} \hline 8 & 12 \\ \hline 3 & 7 \\ \hline \end{array} & \begin{array}{|c|c|} \hline 9 & \\ \hline 1 & 13 \\ \hline \end{array} & \begin{array}{|c|} \hline 10 \\ \hline 4 & 6 \\ \hline 2 & 5 \\ \hline \end{array} \\ T^{(0)} & T^{(1)} & T^{(2)} \quad T^{(3)} \end{array}$$

with shifted contents $(0, 1, 2, 3)$. Then the content reading word of T , say $w = w(T)$, beginning with content -5 , is given by

$$w = 10, 8, 9, \emptyset, 4, \frac{12}{3}, 1, \emptyset, \frac{6}{2}, 7, 13, \emptyset, 5, 11.$$

Since T is standard, let us abuse notation by representing a position of w by the entry which it contains. Then the set of 4-inversions is

$$\text{Inv}_4(w) = \left\{ \begin{array}{ccccc} (10, 8), & (9, 4), & (4, 1), & (12, 2), & (7, 5), \\ (10, 9), & (9, 3), & (12, 1), & (3, 1), & (13, 5), \\ (8, 4), & (4, 3), & (12, 6), & (3, 2), & (13, 11) \end{array} \right\},$$

and so $\text{inv}_4(w) = 15$. Finally, the 4-descent set of w is

$$\text{Des}_4(w) = \{(10, 4), (8, 3), (9, 1), (4, 2), (12, 7), (6, 5)\}.$$

The vertices for which we wish to establish a dual equivalence graph will be the content readings words of standard k -tuples of tableaux. That is

$$V^{(k)} = \{w(T) \mid T \in \text{SYT}_k\}. \quad (4.9)$$

By equation (4.8), preserving the shape of the k -ribbon tableaux amounts to preserving the k -descent set of the content reading words. Refining $V^{(k)}$, for each (skew) shape μ , define

$$V_\mu^{(k)} = \{w(T) \mid T \in \text{SYT}_k(\mu)\}. \quad (4.10)$$

If we can construct colored edge sets E_i on $V^{(k)}$ such that whenever $\{w(T), w(U)\} \in E_i$, we have $\text{Des}_k(w(T)) = \text{Des}_k(w(U))$, then E_i will restrict to a collection of i -colored edges for $V_\mu^{(k)}$.

Define a signature function $\sigma : V^{(k)} \longrightarrow \{\pm 1\}^{n-1}$ by

$$\sigma(w)_j = \begin{cases} +1 & \text{if } \tilde{c}(j) < \tilde{c}(j+1) \\ -1 & \text{if } \tilde{c}(j) > \tilde{c}(j+1) \end{cases} \quad (4.11)$$

for $j = 1, \dots, n-1$, where $\tilde{c}(j)$ is the content of the cell containing the letter j .

Remark 4.3. Since we are considering only words arising from standard k -ribbon tableaux, if i and j occur with the same content in w , say $i < j$, then they appear along the same diagonal in a piece of the k -quotient.

$$\begin{array}{|c|} \hline j \\ \hline i \\ \hline \end{array} \rightsquigarrow \begin{array}{|c|c|} \hline a & j \\ \hline i & b \\ \hline \end{array}$$

In this case, cells containing a and b as depicted above must exist, and since $i < a, b < j$, we must have $j - i \geq 3$. Therefore equation (4.11) defines σ for all j .

Since standardizing a k -tuple of tableaux using the content reading word preserves the k -inversion number, the LLT polynomial $G_\mu^{(k)}$ defined in equation (2.13) may be expressed in terms of quasi-symmetric functions as

$$\tilde{G}_\mu^{(k)}(x; q) = \sum_{T \in \text{SYT}_k(\mu)} q^{\tilde{s}(T)} Q_{n, \sigma(T)}(x) = \sum_{v \in V_\mu^{(k)}} q^{\tilde{s}(v)} Q_{n, \sigma(v)}(x). \quad (4.12)$$

Therefore our goal is to define colored edges E_i on the vertex set $V^{(k)}$ in such a way that $(V^{(k)}, \sigma, E)$ is a dual equivalence graph and whenever $\{w(T), w(U)\} \in E_i$, we have $\text{Des}_k(w(T)) = \text{Des}_k(w(U))$ and $\text{inv}_k(w(T)) = \text{inv}_k(w(U))$. For each (skew) shape μ , this makes $\mathcal{G}_\mu^{(k)} = (V_\mu^{(k)}, \sigma, E)$ into a dual equivalence graph for which the k -inversion number is constant on connected components. Then, by equation (4.12) and Corollary 3.12, we conclude that $\tilde{G}_\mu^{(k)}(x; q)$ is a symmetric function, and we obtain a combinatorial interpretation of the Schur expansion as enumerating connected components of the graph.

4.2 Constructing the edges

The following involutions are the basic ingredients in constructing colored edges on $V^{(k)}$ which preserve the k -inversion number and k -descent set. We regard the distance between two entries of a word as the absolute value of the difference between their contents, with the extension $\text{dist}(a_1, \dots, a_l) = \max_{i,j} \{\text{dist}(a_i, a_j)\}$. Recall from Remark 4.3 that none of $i-1, i, i+1$ may occur with the same content.

Definition 4.4. Define involutions d_i and \tilde{d}_i on words which admit an i -neighbor by

$$\begin{aligned} d_i(\cdots i \cdots i \pm 1 \cdots i \mp 1 \cdots) &= \cdots i \mp 1 \cdots i \pm 1 \cdots i \cdots, \\ \tilde{d}_i(\cdots i \cdots i \pm 1 \cdots i \mp 1 \cdots) &= \cdots i \pm 1 \cdots i \mp 1 \cdots i \cdots, \end{aligned}$$

where all other entries remain fixed. Define $D_i^{(k)}$ by

$$D_i^{(k)}(w) = \begin{cases} d_i(w) & \text{if } \text{dist}(i-1, i, i+1) > k \\ \tilde{d}_i(w) & \text{if } \text{dist}(i-1, i, i+1) \leq k \end{cases}.$$

Remark 4.5. Let w be a word which admits an i -neighbor. If $\text{dist}(i-1, i, i+1) > k$ in w , then $\text{Des}_k(w) = \text{Des}_k(d_i(w))$ and $\text{inv}_k(w) = \text{inv}_k(d_i(w))$. Similarly, if $\text{dist}(i-1, i, i+1) \leq k$ in w , then $\text{Des}_k(w) = \text{Des}_k(\tilde{d}_i(w))$ and $\text{inv}_k(w) = \text{inv}_k(\tilde{d}_i(w))$. Therefore

$$\text{Des}_k(w) = \text{Des}_k(D_i^{(k)}(w)) \quad \text{and} \quad \text{inv}_k(w) = \text{inv}_k(D_i^{(k)}(w)). \quad (4.13)$$

Furthermore, since d_i and \tilde{d}_i change the relative positions of i and $i \pm 1$ but do not change the positions of entries other than $i-1, i, i+1$, we have

$$\sigma(w)_j = -\sigma(D_i^{(k)}(w))_j \quad \text{for } j = i-1, i, \text{ and} \quad (4.14)$$

$$\sigma(w)_h = \sigma(D_i^{(k)}(w))_h \quad \text{for } h < i-2 \text{ and } i+1 < h. \quad (4.15)$$

Since $\text{dist}(i-1, i, i+1) \geq 2$ for every $w \in V^{(k)}$, $D_i^{(1)}$ is just the standard elementary dual equivalence on $i-1, i, i+1$. Therefore allowing $D_i^{(1)}$ to define i -colored edges in the obvious way recovers the standard dual equivalence graph \mathcal{G}_μ .

When $k = 2$ the situation is not much more complicated. Since it is possible to have $\text{dist}(i-1, i, i+1) = 2$, the full description of $D_i^{(2)}$ is needed. However, it is still relatively simple to see that allowing $D_i^{(2)}$ to define i -colored edges on $V^{(2)}$ gives a dual equivalence graph.

Theorem 4.6. *Define E_i to be the set of pairs $\{w, D_i^{(2)}(w)\}$ for words $w \in V^{(2)}$ which admit an i -neighbor. Then $\mathcal{G}_\mu^{(2)} = (V_\mu^{(2)}, \sigma, E)$ is a dual equivalence graph of type (n, n) for which the cospin statistic is constant on connected components.*

Proof. By equation (4.13), $w \in V_\mu^{(2)}$ if and only if $D_i^{(2)}(w) \in V_\mu^{(2)}$, so E_i restricts to $V_\mu^{(2)}$ and preserves cospin. By equations (4.14) and (4.15), $\mathcal{G}_\mu^{(2)}$ satisfies dual equivalence axioms 1 and 2.

If $D_i^{(2)}(w) = \tilde{d}_i(w)$, then $\{w, D_i^{(2)}(w)\} \notin E_{i \pm 1}$. In this case $i-1, i, i+1$ must have consecutive contents with i not between $i-1$ and $i+1$. Furthermore, by Remark 4.3, $i \pm 2$ may not have the same content as any of $i-1, i, i+1$, so $\sigma(w)_h = \sigma(D_i^{(2)}(w))_h$ for $h = i-2, i+1$. On the other hand, if $D_i^{(2)}(w) = d_i(w)$, then $\{w, D_i^{(2)}(w)\} \in E_{i \pm 1}$ if and only if $d_i(w) = d_{i \pm 1}(w)$. This occurs exactly when both $i \pm 2$ and $i \mp 1$ lie between i and $i \pm 1$ in w , thereby demonstrating axiom 3.

Let v, w, x be words such that $\{v, w\} \in E_{i-2}$, $\{w, x\} \in E_i$, and w does not admit an $i-1$ -neighbor. Dual equivalence axioms 1-3 imply that v admits an $i-1$ -neighbor and an i -neighbor, and x admits an $i-2$ -neighbor and an $i-1$ -neighbor. We claim that if $D_{i-1}^{(2)}(v) = D_i^{(2)}(v)$, then $D_{i-2}^{(2)}(x) = D_{i-1}^{(2)}(x)$. Since the number of words u such that $D_{i-1}^{(2)}(u) = D_i^{(2)}(u)$ is equal to the number of words y such that $D_{i-2}^{(2)}(y) = D_{i-1}^{(2)}(y)$, this will establish axiom 4.

Since $D_i^{(2)}(v) = D_{i-1}^{(2)}(v)$, $i-2$ and $i+1$ lie between i and $i-1$ in v . By symmetry, assume that i lies to the left of $i-1$ in v . Suppose first that $w = d_{i-2}(v)$. Then $i-2$ must lie to the right of $i-3, i-1, i, i+1$ in w and also in x . If $i-1$ lies left of $i-3$ in w , then, in x , $i-1$ lies left of $i-3$ and i . If $i-1$ lies right of $i-3$ in w , then $i-1$ is in the same position in v and w , so $i-1$ lies left of both $i-3$ and i in x . In both cases $i-3$ and i lie between $i-1$ and $i-2$ in x , so $D_{i-2}^{(2)}(x) = D_{i-1}^{(2)}(x)$. Now suppose that $w = \tilde{d}_{i-2}(v)$. As in the previous case, $i-2$ must lie to the right of $i-3, i-1, i, i+1$ in w and also in x . Since $i-3, i-2, i-1$ must occur with consecutive contents in v and w , $i+1$ must lie to the left of $i-1$ in w . Therefore $i-1$ will lie to the left of $i-3$ and i in x , so again $D_{i-2}^{(2)}(x) = D_{i-1}^{(2)}(x)$.

Finally, whenever $|i-j| \geq 3$, $\{i-1, i, i+1\} \cap \{j-1, j, j+1\} = \emptyset$, in which case $\mathcal{D}_i^{(2)}$ and $\mathcal{D}_j^{(2)}$ commute when both are defined on a word w . Therefore E_i and E_j satisfy axiom 5 for $|i-j| \geq 3$. \square

Corollary 4.7. *The LLT polynomial $\tilde{G}_\mu^{(2)}(x; q)$ is Schur positive.*

When $k \geq 3$, $D_i^{(k)}$ will not give the edges of a dual equivalence graph. For instance, if w has the pattern 2431 with $\text{dist}(1, 2, 3) \leq k$, then $D_2^{(k)}(w)$ contains the pattern 3412, resulting in a necessary double edge for E_2 and E_3 by dual equivalence axiom 3(b). This implies that the positions of entries other than $i-1, i, i+1$ may need to differ for vertices paired in E_i .

In general, we will construct edges E_i inductively using $D_i^{(k)}$, but taking into account i -edges forced on us by axiom 3(b). We distinguish between the cases when $\text{dist}(i-1, i, i+1) > k$ and $\text{dist}(i-1, i, i+1) \leq k$, the former being given by the original elementary dual equivalences. For the latter, we first categorize certain vertices which admit an i -neighbor. Then we establish a correspondence for these representatives of each i -package, using the terminology of Section 3.4. The goal is to use Proposition 3.29 to define E_i on $V^{(k)}$ in terms of these correspondences of i -packages when E_2, \dots, E_{i-1} have already been defined in such a way that $\mathcal{G}_\mu^{(k)}$ is a dual equivalence graph of type (i, N) for which cospin is constant on connected components.

Definition 4.8. For $3 < i < n$, let $v \in V^{(k)}$ be such that $\text{dist}(i-1, i, i+1) \leq k$ and v admits an i -neighbor and has i -type C. Then v is called *twisted* if $\tilde{d}_i(v)$ has i -type B but is not one of the two left vertices of Figure 3.3, and v is called *pivotal* if the position of $i-1$ differs in v and w , where $\{v, w\} \in C_{i-2}$.

Definition 4.9. Define the *category* of certain vertices of $V^{(k)}$ which have $\text{dist}(i-1, i, i+1) \leq k$ and admit an i -neighbor as follows:

category 0 \tilde{w} such that $\{\tilde{w}, \tilde{x}\} \in E_{i-1}$ and $\sigma(\tilde{w})_i = -\sigma(\tilde{x})_i$.

category 1 w such that $\tilde{d}_i(w)$ is category 0.

category 1' \tilde{v} such that $\{\tilde{d}_i(\tilde{v}), w\} \in E_{i-2}$ for some w in category 1.

category 2 u which is pivotal but not twisted or category 1.

category 2' \tilde{v} such that $\{\tilde{d}_i(\tilde{v}), u\} \in E_{i-2}$ for some u in category 2.

The vertices described in Definition 4.9 are called *categorized*. An i -package is *categorized* if it contains a categorized vertex, otherwise it is *uncategorized*. A vertex which lies on an uncategorized i -package is called *uncategorized*. Note that not all vertices are categorized or uncategorized. For instance, a vertex which is not categorized but which lies on the i -package of a categorized vertex is neither categorized nor uncategorized. We shall refer to these vertices as *implicitly categorized*. When $k \leq 2$, all vertices will be uncategorized. When $k = 3$, we will show that each vertex belongs to at most one category.




Figure 4.1: An illustration of categories and the involution φ when $k = 3$.

Refining Definition 4.9 slightly, for \tilde{w} category 0 such that $\{\tilde{w}, \tilde{x}\} \in E_{i-1}$, say that \tilde{w} is *crossed* if both $\tilde{d}_{i-1}(\tilde{w})$ and $\tilde{d}_{i-1}(\tilde{x})$ are category 0. For w category 1, say that w is *crossed* if $\tilde{d}_i(w)$ is crossed. Similarly, for \tilde{v} category 1' and w such that $\{\tilde{d}_i(\tilde{v}), w\} \in E_{i-2}$, say that \tilde{v} is *crossed* if w is crossed.

Define a map φ on categorized and uncategorized vertices by equation (4.16). For notational simplicity in the following definition, we identify E_h with the involution on vertices admitting an h -neighbor which sends each such vertex to its h -neighbor.

$$\varphi(w) = \begin{cases} E_{i-1}(w) & \text{for } w \text{ category 0,} \\ \tilde{d}_i \tilde{d}_{i-1} \tilde{d}_i(w) & \text{for } w \text{ category 1 and crossed,} \\ \tilde{d}_i E_{i-1} \tilde{d}_i(w) & \text{for } w \text{ category 1 and not crossed,} \\ \tilde{d}_i E_{i-2} \tilde{d}_i \tilde{d}_{i-1} \tilde{d}_i E_{i-2} \tilde{d}_i(w) & \text{for } w \text{ category 1' and crossed,} \\ \tilde{d}_i E_{i-2} \tilde{d}_i E_{i-1} \tilde{d}_i E_{i-2} \tilde{d}_i(w) & \text{for } w \text{ category 1' and not crossed,} \\ \tilde{d}_i(w) & \text{for } w \text{ category 2,} \\ \tilde{d}_i E_{i-2} \tilde{d}_i E_{i-2} \tilde{d}_i(w) & \text{for } w \text{ category 2',} \\ D_i^{(k)}(w) & \text{for } w \text{ uncategorized.} \end{cases} \quad (4.16)$$

The idea is to use φ to construct E_i on the i -package of each categorized or uncategorized

vertex. For this to be a viable approach, we must show that φ satisfies the hypotheses of Proposition 3.29 from Section 3.4.

By Remark 4.5 and induction on the color of edges, φ preserves the k -descent set and the k -inversion number. We must show that φ satisfies the following conditions for categorized or uncategorized vertices v and w :

1. φ is an involution;
2. v and w lie on the same i -package if and only if $\varphi(v)$ and $\varphi(w)$ lie on the same i -package;
3. \mathcal{P}'_w and $\mathcal{P}'_{\varphi(w)}$ are i -compatible at w and $\varphi(w)$.

Then we may define E_i to be the set of pairs $\{w, \varphi(w)\}$ for each categorized or uncategorized vertex w . By Proposition 3.29, E_i may be extended along i -packages, and doing so makes $(V_\mu^{(k)}, \sigma, E_2 \cup \dots \cup E_i)$ into a dual equivalence graph of type $(i+1, n)$ for which cospin is constant on connected components.

Note that this definition of E_i coincides with defining edges directly from $D_i^{(k)}$ when $k \leq 2$ since all vertices will be uncategorized. We believe that, in fact, the map φ constructed in equation (4.16) always satisfies the hypotheses of Proposition 3.29. If true, this would establish the following conjecture.

Conjecture 4.10. *The graph $\mathcal{G}_\mu^{(k)}$ is well-defined and is a dual equivalence graph of type (n, n) for which the cospin statistic is constant on connected components.*

4.3 The graph on 3-ribbon tableaux is a dual equivalence graph

In this section, we prove by induction that $\mathcal{G}_\mu^{(3)} = (V, \sigma, E)$ is well-defined and is a dual equivalence graph for which the 3-inversion number is constant on connected components. The base case for the induction follows from the fact $D_2^{(3)}$ defines E_2 .

Assume that $(V, \sigma, E_2 \cup \dots \cup E_{i-1})$ is a dual equivalence graph for which the 3-inversion number and 3-descent set are constant on connected components. We also assume that for $\{w, x\} \in E_h$ with $h < i$,

- (IH1) the positions of letters $> h + 1$ are the same in w as in x ;
- (IH2) if $\text{dist}(h-1, h, h+1) > 3$ in w , then $x = d_h(w)$;
- (IH3) if $\text{dist}(h-1, h, h+1) \leq 3$ in w , then, when moving from w to x , the positions of $h-1, h, h+1$ change according to Figure 4.2 as explained in Remark 4.11;
- (IH4) and if the positions of $i-1, i, i+1$ in x are the same as in $\tilde{d}_h(w)$, then $x = \tilde{d}_h(w)$.




Figure 4.2: Possible E_h edges when $\text{dist}(h-1, h, h+1) \leq 3$.

Remark 4.11. Figure 4.2 should be interpreted as follows. The four indicated letters occupy the same four positions on both sides of the figure, though the specific letter represented by “*” may differ on the two sides. The edges between sides indicate which patterns may be linked by an h -edge. Reversing the words throughout the diagram gives the remaining cases.

From the inductive hypotheses it follows that if w is category 0, then $\{w, \tilde{d}_{i-1}(w)\} \in E_{i-1}$. Therefore all category 0 vertices are crossed, and hence so, too, are all category 1 vertices. However, in this case the definition of φ for crossed and noncrossed vertices agree.

Since $D_i^{(k)}$ does not affect the position of letters $> i + 1$, by IH1, if E_i may be constructed as described in Section 4.2, then it also satisfies IH1. By IH1, the positions of $i - 1, i, i + 1$ are the same for every vertex on a reduced i -package. Furthermore, positions of $i, i + 1$ are the same for every vertex of an i -package, though the position of $i - 1$ may differ. The following proposition investigates to what extent this can occur on an i -package. Recall that C_{i-2} denotes those E_{i-2} edges connecting vertices having i -type C.

Proposition 4.12. *For $\{w, v\} \in C_{i-2}$, we have $\text{dist}(i-1, i, i+1) > 3$ in w if and only if $\text{dist}(i-1, i, i+1) > 3$ in v .*

Proof. We will show that when $\text{dist}(i-1, i, i+1) > 3$, if $i \pm 1$ lies between i and $i \mp 1$ in w , then the same holds in v . Furthermore, the positions of the outer two of $i-1, i, i+1$ are the same in v as in w . By symmetry, assume that i lies to the left of $i-1$ and $i+1$ in w .

If $i-1$ lies further right in v than in w , then by IH2 and IH3 applied to E_{i-2} , $i-2$ must lie between i and $i-1$ in v . If $i+1$ also lies between i and $i-1$ and $\text{dist}(i, i-1) > 3$, then, by IH2, $\{v, d_{i-1}(v)\} \in E_{i-1}$ which forces v to have i -type B. To avoid the contradiction, $i-1$ may not lie left of $i+1$ in v if $i-1$ lies between i and $i+1$ in w , and $i-1$ may not be further right in v than in w if $i+1$ lies between i and $i-1$ in w .

If $i-1$ lies further left in v than in w , then $i-1$ must remain to the right of i , otherwise v would not admit an i -neighbor. If $i-1$ lies between i and $i+1$ in w , then combining this with the previous assertions shows that the same holds in v . Similarly, if $i+1$ lies between i and $i-1$ in w , then $\text{dist}(i-1, i) > 3$ in w , so $i-2$ must lie between i and $i-1$ in w in order for $i-1$ to be further to the left in v . But then, by IH2, $\{w, d_{i-1}(w)\} \in E_{i-1}$ which forces w to have i -type B. Therefore $i-1$

may not be further left in v than in w if $i+1$ lies between $i-1$ and i in w . Therefore whenever $i+1$ lies between $i-1$ and i in w , the position of $i-1$ is the same in v as in w . \square

By Proposition 4.12, all vertices with $\text{dist}(i-1, i, i+1) > 3$ are uncategorized. Therefore, IH2 holds for E_i . Furthermore, φ , acting as d_i , is an involution among all such vertices. By Remark 4.5, for w with $\text{dist}(i-1, i, i+1) > 3$, d_i gives an isomorphism from the $(i-1, i-2)$ -restriction of \mathcal{P}_w to the $(i-1, i-2)$ -restriction of $\mathcal{P}_{d_i(w)}$ satisfying conditions c2 and c3 of i -compatibility (Definition 3.24). Therefore if w and v lie on the same i -package and $\text{dist}(i-1, i, i+1) > 3$ in w , then $\varphi(w)$ and $\varphi(v)$ lie on the same i -package. The following proposition shows that d_i also satisfies condition c4 of i -compatibility, which establishes the hypotheses of Proposition 3.29 in the case of vertices with $\text{dist}(i-1, i, i+1) > 3$.

Proposition 4.13. *If w admits an i -neighbor and has $\text{dist}(i-1, i, i+1) > 3$, then w has i -type C if and only if $d_i(w)$ has i -type C.*

Proof. Let W be the set of words admitting an $i-2$ -neighbor and an i -neighbor but not an $i-1$ -neighbor for which $\text{dist}(i-1, i, i+1) > 3$. Similarly let X denote those words admitting an $i-2$ -neighbor, an $i-1$ -neighbor and an i -neighbor, but which are not one of the left two vertices of i -type B, for which $\text{dist}(i-1, i, i+1) > 3$. Note that each word of W, X has i -type C or one of the right two vertices of i -type B. Therefore if we establish that d_i gives a correspondence between the subsets of words of i -type B, we have also established it for i -type C.

Let $w \in W$ and $x \in X$. Let v, y be such that $\{w, v\}, \{x, y\} \in E_{i-2}$. Then w has i -type B if and only if there exists u such that $\{v, u\} \in E_{i-1}$ and $\sigma(v)_i = -\sigma(u)_i$. On the other side, x has i -type B if and only if $\{x, y\} \in E_{i-1}$. We will give direct characterizations of the subsets of i -type B based on the positions of $i-3$ and $i-2$. By symmetry we may assume i is to the left of $i-1$ and $i+1$ in W and to the right in X .

First suppose that $i+1$ lies between $i-1$ and i in w and in x . Since $i-2$ must lie to the right of $i-1$ in w , the position of $i-1$ in v will be weakly right of the position in w . Therefore w has i -type B if and only if $i-3$ lies to the right of i in w . On the other hand, x has i -type B if and only if $i-2$ lies to the right of i and $i-3$ lies to the right of $i-1$ in x . Therefore d_i provides the desired correspondence between the subsets of W and X of vertices with i -type B.

Now suppose that $i-1$ lies between $i+1$ and i in w and in x . In this case, w has i -type B if and only if $i-2$ lies to the right of $i+1$ in w and $i-1$ lies to the right of $i+1$ in v . The requirement for x to have i -type B is that $i-2$ must lie to the right of i in x and $i-1$ must lie to the left of i in y . By IH1, the action of E_{i-2} regarded as an involution on words admitting an $i-2$ -neighbor is independent of the positions of letters $> i-1$. Therefore d_i will not change how E_{i-2} acts, and as such provides the desired correspondence. \square

We may now focus on the case when $\text{dist}(i-1, i, i+1) \leq 3$ in w . By Remark 4.3, none of $h-1, h, h+1$ may reside in the same letter. Taking this further, $h-2$ may reside with $h+1$ if and

only if $h - 1$ and h are exactly 3 letters away in opposite directions, hence $i - 2$ may not reside in the same letter as $i + 1$ in this case. Furthermore, $h - 3$ may reside with $h + 1$ if and only if two of $h - 2, h - 1, h$ are exactly 3 letters away in opposite directions. These observations allow us to make a complete characterization of when a vertex admitting an i -neighbor has i -type B or C based on the relative positions of $i - 3$ and $i - 2$ with $i - 1, i, i + 1$, see Figure 4.4. Similarly, Figure 4.3 gives the C_{i-2} edges for pivotal vertices. A careful analysis of these figures leads to the following observations.




Figure 4.3: Pivotal vertices when $k = 3$.

Remark 4.14. Recall that the action of E_h commutes with the action of \tilde{d}_i , meaning $\tilde{d}_i E_h(w) = E_h \tilde{d}_i(w)$, whenever $h \leq i - 3$. Furthermore, by inspecting Figure 4.4, the action of C_{i-2} commutes with the action of \tilde{d}_i at a vertex w , meaning $\tilde{d}_i C_{i-2}(w) = C_{i-2} \tilde{d}_i(w)$, if and only if neither w nor $C_{i-2}(w)$ is twisted.

Remark 4.15. Insofar as it affects the positions of $i - 3, i - 2, i - 1, i, i + 1$, φ may be explicitly defined for all categorized vertices as indicated on Figure 4.4. Consequently, it is clear that φ conforms with Figure 4.2 for all categorized and uncategorized vertices. Since the first case of a pivotal vertex in Figure 4.3 is determined by category 1 and the latter two by category 2, a quick check verifies the remaining inductive hypotheses, IH3 and IH4.

Lemma 4.16. *If w, v are categorized vertices lying on the same i -package, then either they are the same category or one is category 0 and the other category $1'$. In both cases, w and v lie on the same reduced i -package.*

Proof. First we claim that there are no pivotal vertices on the i -packages of category 0, $1'$ and $2'$ vertices, and consequently the position of $i - 1, i, i + 1$ is constant on the i -packages of these vertices. By Figure 4.3, the only instance when this is not clear is when w is category $1'$ with the pattern $i - 1 \ i + 1 \ i - 3 \ i \ i - 2$ and x is pivotal with the pattern $i - 3 \ i - 1 \ i + 1 \ i - 2 \ i$. Since x is pivotal and the positions of $i - 1, i, i + 1$ are the same for w and x , by Lemma 3.14, w and x must reside on the same reduced i -package. Furthermore, there exists a sequence of edges from w to x in \mathcal{P}'_w consisting of at most one E_{i-3} edge. Therefore, in w , $i - 4$ must be immediately right of $i - 2$ with E_{i-3} acting by \tilde{d}_{i-3} . Using Lemma 3.14 yet again, there is at most one E_{i-4} edge in the edge sequence from here to x , which makes it impossible to get $i - 3$ into the necessary position.

Consider w lying on an external i -package. Suppose first that w is category 1. Then v cannot belong to category $1'$ or $2'$ since the positions of $i - 1, i, i + 1$ for such vertices are not

compatible with the positions in vertices of \mathcal{P}_w . The only way for v to be category 2 is if there is a path from w to v crossing at most one E_{i-2} edge, by Lemma 3.14. However, Lemma 3.14 also means that the path crosses at most one E_{i-3} edge on each of \mathcal{P}'_w and \mathcal{P}'_v . As in the previous argument, given the necessary relative positions of $i-3, i-2, i-1, i, i+1$ in w and v , this is not possible. Similarly, if w is category 2 then v cannot be category 1' or 2' based on the positions of $i-1, i, i+1$.

Since Figure 4.4 shows that all twisted vertices lie on the i -packages of category 1 or 2 vertices, it follows from Lemma 4.16 that no twisted vertices lie on the i -package of a category 0, 1' or 2' vertex. Therefore, by Remark 4.14, \tilde{d}_i commutes with all edges of \mathcal{P}_w whenever w is category 1' or 2'. Hence a path from w to v in \mathcal{P}_w gives rise to a path from $\tilde{d}_i(w)$ to $\tilde{d}_i(v)$ in $\mathcal{P}_{\tilde{d}_i(w)}$. If w, v are category 1', 2', respectively, then $\tilde{d}_i(w), \tilde{d}_i(v)$ lie on the i -packages of category 1, 2 vertices, respectively, which contradicts the previous result.

By Lemma 3.14 and the preceding discussion, since the positions of $i-1, i, i+1$ and at least one of $i-2$ or $i-3$ are fixed for any given category, whenever two categorized vertices of the same category lie on the same i -package, they in fact lie on the same reduced i -package.

Turning our attention to internal i -packages, the positions of $i-1, i, i+1$ dictate that the only categories which could possibly appear, apart from category 0, are categories 1' or 2'. However, since \tilde{d}_i commutes with all edges of the internal i -package, having a category 2' vertex would imply that a category 2 vertex lies on the same i -package as a category 1 vertex, which is still a contradiction. Finally, by Remark 3.27, all vertices of an internal i -package lie on the same reduced i -package. \square

As a corollary, note that Lemma 4.16 proves that a given vertex belongs to at most one category. In particular, φ is well-defined.

Lemma 4.17. 1. If $\tilde{d}_i(w)$ is category 0, then $\sigma(w)_{i-2} = \sigma(\varphi(w))_{i-2}$.

2. If $\sigma(w)_{i-2} = -\sigma(\tilde{d}_i(w))_{i-2}$, then one of $w, \tilde{d}_i(w)$ is category 0.

Proof. If $\sigma(w)_{i-2} = -\sigma(\tilde{d}_i(w))_{i-2}$, then $\text{dist}(i-2, i-1, i, i+1) \leq 3$ with, say, w having the pattern $i \ i+1 \ i-2 \ i-1$. By Remark 4.3, E_{i-1} may not move $i-1$ or i into the position where $i+1$ resides since E_{i-1} preserves 3-descents. Also, if $i-2$ moves in with $i+1$, then i and $i-1$ must be exactly 3 letters away from $i+1$ in opposite directions. Figure 4.2 leaves no alternative but for E_{i-1} to act on $i-2, i-1, i$ by \tilde{d}_{i-1} , giving $\sigma(w)_i = -\sigma(\tilde{d}_{i-1}(w))_i$ and proving the second assertion.

Inspecting this situation further, if $\{\tilde{d}_i(w), \tilde{d}_i(x)\} \in E_{i-1}$ are category 0, then $\sigma(w)_{i-2} = -\sigma(\tilde{d}_i(w))_{i-2}$ if and only if $\sigma(x)_{i-2} = \sigma(\tilde{d}_i(x))_{i-2}$. By dual equivalence axiom 3 applied to E_{i-1} , $-\sigma(\tilde{d}_i(w))_{i-2} = \sigma(\tilde{d}_i(x))_{i-2}$. Putting these together proves the first assertion. \square

Proposition 4.18. *The map φ is an involution for each category and for uncategorized vertices.*

Proof. We first show that for any vertex w lying on a categorized i -package, $\tilde{d}_i(w)$ also lies on a categorized i -package. For w lying on an internal i -package, $\tilde{d}_i(w)$ lies on the i -package of a category 1 vertex by Remark 4.14 since internal i -packages may not have twisted vertices. Similarly, for w

lying on the i -package of a category $1'$ vertex, $\tilde{d}_i(w)$ lies on the i -package of a category 1 vertex since no twisted vertices lie on a category $1'$ i -package. If w is on a category 1 i -package, then by Lemma 3.14, we may choose a path in \mathcal{P}_w from the category 1 vertex to w which crosses at most one E_{i-2} edge. Then $\tilde{d}_i(w)$ lies on a category 0 i -package if this path does not pass through a twisted vertex, and, by Figure 4.4, $\tilde{d}_i(w)$ lies on a category $1'$ i -package otherwise.

Category 2 and $2'$ may be dealt with in a similar fashion. For w on a category $2'$ i -package, $\tilde{d}_i(w)$ lies on a category 2 i -package by Remark 4.14 since there are no twisted vertices on the i -package of w . For w on a category 2 i -package, $\tilde{d}_i(w)$ lies on a category 2 i -package if there is a path from w to the category 2 vertex which does not pass through a twisted vertex, and it lies on a category $2'$ i -package otherwise.

It follows now that φ is an involution for uncategorized vertices, so it remains only to show the result for categorized vertices. This is obvious for category 0, from which it follows that φ is also an involution for category 1. Lemma 4.17 implies that φ preserves the i -type of category 1 vertices (see Remark 3.19), and so φ is an involution for category $1'$ as well. Again, φ is clearly an involution on category 2, from which it follows that it is also an involution on category $2'$. \square

Proposition 4.19. *For categorized or uncategorized vertices w and v which lie on the same i -package, $\varphi(w)$ and $\varphi(v)$ also lie on the same i -package.*

Proof. First consider the case when w and v are uncategorized. Since no twisted or pivotal vertices lie on uncategorized i -packages, by Remark 4.14, \tilde{d}_i gives an isomorphism from the $(i-1, i-2)$ -restriction of \mathcal{P}_w to the $(i-1, i-2)$ -restriction of $\mathcal{P}_{\tilde{d}_i(w)}$. Hence $\phi(w)$ and $\phi(v)$ both lie on $\mathcal{P}_{\tilde{d}_i(w)}$.

Now assume w and v are both categorized. By Lemma 4.16, w and v lie on the same reduced i -package. For internal i -packages, it suffices to assume both w and v are category 0 since category $1'$ vertices are always connected to category 0 vertices by E_{i-3} . In this case, Lemma 3.26 ensures that any path from w to v in \mathcal{P}'_w gives rise to a path from $\varphi(w)$ to $\varphi(v)$ in $\mathcal{P}'_{\varphi(w)}$. We may now use the commutivity relations suggested by Figure 4.1 to push the result through for other categories. Since \tilde{d}_i gives an isomorphism between reduced i -packages, the analogous result holds for i -packages with a category 1 vertex. This being the case for category 1, it must hold for category $1'$ as well. The result for category 2 also follows from the fact that \tilde{d}_i gives an isomorphism between reduced i -packages, and again this may now be extended to category $2'$. \square

Proposition 4.20. *For w categorized or uncategorized with $\text{dist}(i-1, i, i+1) \leq 3$, \mathcal{P}'_w and $\mathcal{P}'_{\varphi(w)}$ are i -compatible at w and $\varphi(w)$.*

Proof. Lemma 3.26 proves the result in the case of internal i -packages, i.e. when w is category 0. By Lemma 4.17, \tilde{d}_i preserves σ_{i-2} for all uncategorized vertices. Since no twisted vertices lie on uncategorized i -packages, \tilde{d}_i establishes the i -compatibility of \mathcal{P}'_w and $\mathcal{P}'_{\varphi(w)}$ at w and $\varphi(w)$ for uncategorized vertices w . Therefore we may assume w is a categorized vertex lying on an external i -package. To simplify notation, let \tilde{v} denote $\tilde{d}_i(v)$ for any vertex v which admits an i -neighbor.

Recall that \tilde{d}_i is an isomorphism between the $(i-2, i-2)$ -restrictions of \mathcal{P}'_w and $\mathcal{P}'_{\tilde{w}}$. Since $\mathcal{P}_{\tilde{v}}$ and $\mathcal{P}_{\varphi(\tilde{v})}$ are i -compatible at \tilde{v} and $\varphi(\tilde{v})$ whenever \tilde{v} is category 0, for all category 1 and 2 vertices w , φ extends to an isomorphism between the $(i-2, i-2)$ -restrictions of \mathcal{P}'_w and $\mathcal{P}'_{\varphi(w)}$ satisfying condition c2 and the latter part of c3 in Definition 3.24. We claim that φ also preserves σ_{i-2} .

The claim for u in category 2 follows immediately from Lemma 4.16 and part 2 of Lemma 4.17. Consider w in category 1, say with $x = \varphi(w)$. By part 1 of Lemma 4.17, $\sigma(w)_{i-2} = \sigma(x)_{i-2}$. Let v and y be corresponding vertices on \mathcal{P}'_w and \mathcal{P}'_x . Then \tilde{v} and \tilde{y} correspond on $\mathcal{P}'_{\tilde{w}}$ and $\mathcal{P}'_{\tilde{x}}$, and so $\sigma(\tilde{v})_{i-2} = -\sigma(\tilde{y})_{i-2}$ if and only if $\{\tilde{v}, \tilde{y}\} \in E_{i-1} \cap E_i$. That is to say, \tilde{v} and \tilde{y} are category 0, and v and y are category 1. Therefore whenever v and y are not category 1, $\sigma(\tilde{v})_{i-2} = \sigma(\tilde{y})_{i-2}$, and so, by part 2 of Lemma 4.17, $\sigma(v)_{i-2} = \sigma(y)_{i-2}$.

By Proposition 3.20, φ extends to an isomorphism of connected components of $(V, \sigma, E_2 \cup \dots \cup E_{i-2})$ for categories 1 and 2. Using Lemma 4.17 in conjunction with \tilde{d}_i establishes that φ is an isomorphism between reduced i -packages for categories 1' and 2'. It remains only to show that, on reduced i -packages, if w has i -type C, then so does $\varphi(w)$.

For any category 1 vertex w , by Remark 4.14 and Proposition 4.19, v is twisted in \mathcal{P}'_w if and only if \tilde{v} belongs to category 1' in $\mathcal{P}'_{\tilde{w}}$, and v belongs to category 1 in \mathcal{P}'_w if and only if \tilde{v} belongs to category 0 in $\mathcal{P}'_{\tilde{w}}$. Letting $x = \varphi(w)$, it follows that corresponding vertices of \mathcal{P}'_w and \mathcal{P}'_x differ in i -type if and only if corresponding vertices $\mathcal{P}'_{\tilde{w}}$ and $\mathcal{P}'_{\tilde{x}}$ differ in i -type. Since \tilde{w} and \tilde{x} are category 0, corresponding vertices of $\mathcal{P}'_{\tilde{w}}$ and $\mathcal{P}'_{\tilde{x}}$ have the same i -type, so the same must hold for corresponding vertices of \mathcal{P}'_w and \mathcal{P}'_x . With $\mathcal{P}_w, \mathcal{P}_x$ i -compatible at w and x , the same argument may be extended to category 1'.

For w in category 2, no vertex of \mathcal{P}'_w is twisted or category 1' or 2'. Therefore corresponding vertices both or neither have i -type C. The previous argument may be used to extend this result to category 2'. \square

Define E_i to be the set of pairs $\{w, \varphi(w)\}$ for each categorized or uncategorized vertex w . By Propositions 4.18, 4.19 and 4.20, E_i satisfies the hypotheses of Proposition 3.29. Therefore E_i may be extended along i -packages, and doing so makes $(V_\mu^{(k)}, \sigma, E_2 \cup \dots \cup E_i)$ into a dual equivalence graph of type $(i+1, n)$ for which cospin is constant on connected components. Since the four inductive hypotheses have been established for E_i defined in this way, we have proved the following.

Theorem 4.21. *For $k \leq 3$, $\mathcal{G}_\mu^{(k)}$ is well-defined and is a dual equivalence graph of type (n, n) for which the cospin statistic is constant on connected components.*

Corollary 4.22. *For $k \leq 3$, let $C_\mu^{(k)}(\lambda)$ denote the set of connected components of $\mathcal{G}_\mu^{(k)}$ which are isomorphic to \mathcal{G}_λ . Then*

$$\tilde{G}_\mu^{(k)}(x; q) = \sum_\lambda \left(\sum_{\mathcal{C} \in C_\mu^{(k)}(\lambda)} q^{\tilde{s}(\mathcal{C})} \right) s_\lambda(x).$$

In particular, $\tilde{G}_\mu^{(k)}(x; q)$ is Schur positive.

At this point it makes sense to discuss the status of Conjecture 4.10 in the general case. Inductive hypotheses IH1 and IH2 are easily proven from the remarks in the proof for $k = 3$. With a weakened version of IH3, which is also easy to prove, Propositions 4.12 and 4.13 follow as presented. Therefore, for arbitrary k , the result is established for vertices with $\text{dist}(i-1, i, i+1) > k$. For vertices with $\text{dist}(i-1, i, i+1) \leq k$, finding a generalized version of Figure 4.4 is difficult, and as such Remarks 4.14 and 4.15 are only proven in special cases. However, Lemma 4.16 follows as presented if these Remarks can be established. Also, inductive hypotheses IH3 and IH4 follow from Remark 4.15, and as a corollary to IH3 we also obtain Lemma 4.17. Furthermore the proofs of Propositions 4.18, 4.19 and 4.20 will hold in general once these supporting Remarks and Lemmas have been established. Therefore by carefully analyzing the vertices of types B and C which have $\text{dist}(i-1, i, i+1) \leq k$, we hope that Remarks 4.14 and 4.15 may be established in general thereby leading to a proof of Conjecture 4.10 for all k .

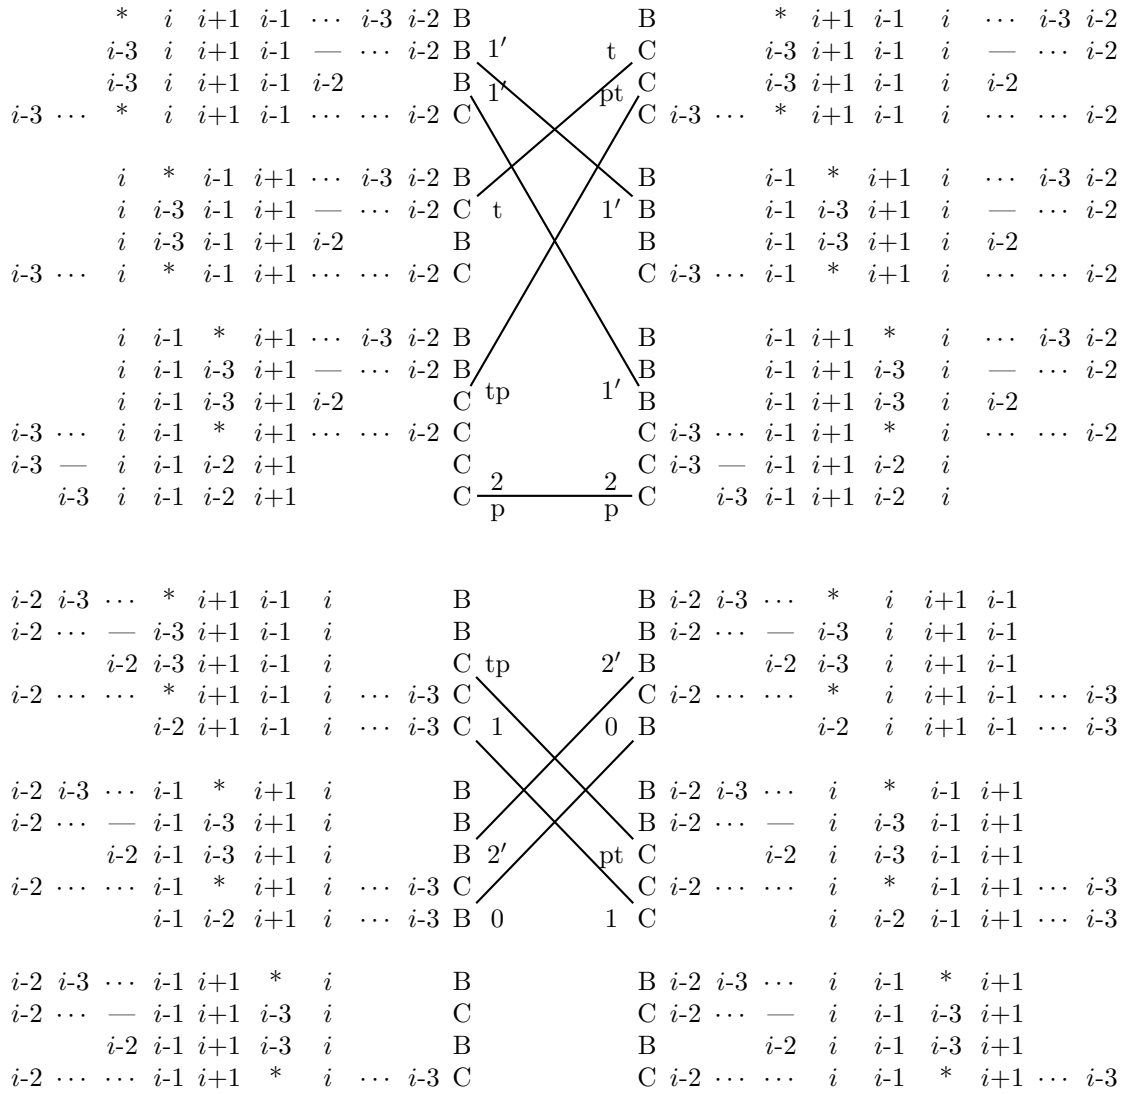


Figure 4.4: Vertices of i -type B or C, with categories indicated by number and twisted / pivoted indicated by t/p.

Chapter 5

Macdonald polynomials

In 1988, Macdonald [Mac88] found a remarkable new basis of symmetric functions in two parameters which specializes to Schur functions, complete homogeneous, elementary and monomial symmetric functions and Hall-Littlewood functions, among others. The transformed Macdonald polynomials $\tilde{H}_\mu(x; q, t)$ are uniquely characterized by certain orthogonality and triangularity conditions as follows.

Proposition 5.1 ([Hai99]). *The transformed Macdonald polynomials $\tilde{H}_\mu(x; q, t)$ are the unique functions satisfying the following conditions:*

1. $\tilde{H}_\mu(x; q, t) \in \mathbb{Q}(q, t)\{s_\lambda[X/(1-q)] \mid \lambda \geq \mu\}$;
2. $\tilde{H}_\mu(x; q, t) \in \mathbb{Q}(q, t)\{s_\lambda[X/(1-t)] \mid \lambda \geq \mu'\}$;
3. $\tilde{H}_\mu[1; q, t] = 1$.

The square brackets in Proposition 5.1 stand for *plethystic substitution*. In short, $s_\lambda[A]$ means s_λ applied as a Λ -ring operator to the expression A , where Λ is the ring of symmetric functions. For a thorough account of plethysm, see [Hai99].

It is clear from this definition that $\tilde{H}_\mu(x; q, t)$ is a symmetric function, therefore we may define the Kostka-Macdonald polynomials $\tilde{K}_{\lambda, \mu}(q, t)$, which give the Schur expansion of Macdonald polynomials, by

$$\tilde{H}_\mu(x; q, t) = \sum_{\lambda} \tilde{K}_{\lambda, \mu}(q, t) s_\lambda(x). \quad (5.1)$$

Garsia and Haiman [GH93] conjectured that the transformed Macdonald polynomials $\tilde{H}_\mu(x; q, t)$ could be realized as the bigraded characters of certain modules for the diagonal action of S_n on two sets of variables. This became known as the $n!$ Conjecture. By analyzing the algebraic geometry of the Hilbert scheme of n points in the plane, Haiman [Hai01] proved the $n!$ Conjecture and consequently established Macdonald positivity.

Theorem 5.2 ([Hai01]). *We have $\tilde{K}_{\lambda,\mu}(q, t) \in \mathbb{N}[q, t]$.*

Below we give Haglund's combinatorial formula for Macdonald polynomials expanded in terms of monomials, and show how the graphs constructed for LLT polynomials also apply to Macdonald polynomials. Therefore Conjecture 4.10 has the further consequence of a combinatorial formula for the Kostka-Macdonald polynomials $\tilde{K}_{\lambda,\mu}(q, t)$.

5.1 Haglund's formula for Macdonald polynomials

In 2004, Haglund [Hag04] conjectured a combinatorial formula for the monomial expansion of $\tilde{H}_\mu(x; q, t)$. This formula, which was proven by Haglund, Haiman and Loehr [HHL05a], does not give a combinatorial proof of $\tilde{K}_{\lambda,\mu}(q, t) \in \mathbb{N}[q, t]$. However, combining Theorem 2.3, Proposition 3.4 and equation (23) from [HHL05a], to give a combinatorial description of $\tilde{K}_{\lambda,\mu}(q, t)$ it suffices to give a description of the Schur expansion of certain LLT polynomials.

Recall from Section 2.1 that for a cell c in the diagram of λ , $a(c)$ is the number of cells to the east of c , and $l(c)$ is the number of cells north of c .

Let S be a filling of λ , that is $S : \lambda \rightarrow \mathbb{N}$. A *descent* of S is an ordered pair of cells (c, d) of λ such that c lies immediately north of d and the entry of c is greater than the entry of d . Denote by $\text{Des}(S)$ the set of all descents of S , i.e.

$$\text{Des}(S) = \{((i, j), (i, j - 1)) \in \lambda \mid S(i, j) > S(i, j - 1)\}. \quad (5.2)$$

Then define the *major index* of S , denote $\text{maj}(S)$, by

$$\text{maj}(S) \stackrel{\text{def}}{=} |\text{Des}(S)| + \sum_{(c,d) \in \text{Des}(S)} l(c). \quad (5.3)$$

An ordered pair of cells (c, d) is called *attacking* if they lie in the same row and c is west of d , or if c is in the row immediately north of d and c lies strictly east of d . An *inversion pair* of S is an attacking pair (c, d) such that the entry of c is greater than the entry of d . Denote by $\text{Inv}(S)$ the set of inversion pairs of S , i.e.

$$\text{Inv}(S) = \left\{ ((i, j), (g, h)) \in \lambda \mid \begin{array}{l} j = h \text{ and } i < g \text{ or } j = h + 1 \text{ and } g < i, \\ \text{and } S(i, j) > S(g, h) \end{array} \right\}. \quad (5.4)$$

Then define the *inversion number* of S , denoted $\text{inv}(S)$, by

$$\text{inv}(S) \stackrel{\text{def}}{=} |\text{Inv}(S)| - \sum_{(c,d) \in \text{Des}(S)} a(c). \quad (5.5)$$

Note that if $(c, d) \in \text{Des}(S)$, for every cell e of the arm of c , the entry of e is either bigger than the entry of d or smaller than the entry of c (or both). In the former case, (e, d) will form an inversion pair, and in the latter case, (c, e) will form an inversion pair. Therefore $\text{inv}(S)$ is a nonnegative integer.

Example 5.3. Let S be the following filling of $(5, 4, 4, 1)$.

5					
11	14	9	2		
6	3	4	10		
8	1	13	7	12	

As in Example 4.2, let us abuse notation by representing a cell of S by the entry which it contains. Then the inversion pairs of S are given by

$$\text{Inv}(S) = \left\{ \begin{array}{l} (11, 9), (14, 2), (9, 6), (6, 4), (10, 1), (13, 7), \\ (11, 2), (14, 6), (9, 3), (4, 1), (8, 1), (13, 12) \\ (14, 9), (9, 2), (6, 3), (10, 8), (8, 7), \end{array} \right\}.$$

Similarly, the descents of S are given by

$$\text{Des}(S) = \{(11, 6), (14, 3), (3, 1), (9, 4), (10, 7)\}.$$

Therefore the statistics associated to S are

$$\begin{aligned} \text{maj}(S) &= 5 + (1 + 0 + 1 + 0 + 1) = 8, \\ \text{inv}(S) &= 17 - (3 + 2 + 2 + 1 + 0) = 9. \end{aligned}$$

Theorem 5.4 ([HHL05a]). *The transformed Macdonald polynomials $\tilde{H}_\mu(x; q, t)$ may be expressed in terms of monomials as*

$$\tilde{H}_\mu(x; q, t) = \sum_{S: \mu \rightarrow \mathbb{N}} q^{\text{inv}(S)} t^{\text{maj}(S)} x^S. \quad (5.6)$$

This formula may be related to LLT polynomials as follows. Let D be a possible descent set for μ , that is, $D = \text{Des}(S)$ for some filling S of μ . For $i = 1, \dots, \mu_1$, let $\mu_D^{(i-1)}$ be the ribbon obtained from the i th column of μ by putting the cell (i, j) immediately east of $(i, j+1)$ if $((i, j+1), (i, j)) \in D$ and immediately south of $(i, j+1)$ otherwise. Let μ_D be the μ_1 -tuple of partitions $(\mu_D^{(0)}, \dots, \mu_D^{(\mu_1-1)})$. Then each filling S of shape μ with $\text{Des}(S) = D$ may be regarded as a μ_1 -ribbon tableau of shape μ_D . For example, the filling S of $(5, 4, 4, 1)$ in Example 5.3 corresponds to the following 5-tuple.

			14								
5	11			3	9			2	10		
		6	8	1	4	13			7	12	

If we define the shifted contents of $(\mu_D^{(0)}, \dots, \mu_D^{(\mu_1-1)})$ so that the southeasternmost cell of $\mu_D^{(i)}$ has shifted content i , then the inversion pairs of S correspond precisely with the μ_1 -inversions of the μ_1 -tuple as defined in equation (4.1).

Since the major index statistic depends only on the descent set, for a given descent set D we may define $\text{maj}(D)$ by $\text{maj}(D) = \text{maj}(S)$ for any filling S with $\text{Des}(S) = D$. Similarly, define

$a(D) = |\text{Inv}(S)| - \text{inv}(S)$ for any filling S with $\text{Des}(S) = D$. By Proposition 4.1, for every k -tuple of partitions ν , there exists a constant, say $d(\nu)$, such that for every semi-standard filling T of shape ν

$$\text{cospin}(T) + d(\nu) = \text{inv}_k(T).$$

Therefore we may rewrite equation (5.6) in terms of LLT polynomials as

$$\tilde{H}_\mu(x; q, t) = \sum_D q^{d(\mu_D) - a(D)} t^{\text{maj}(D)} \tilde{G}_{\mu_D}^{(\mu_1)}(x; q). \quad (5.7)$$

For each possible descent set D , there exists a filling S such that $\text{Des}(S) = D$ and $\text{cospin}(T) = 0$ for the corresponding μ_1 -tuple T . In this case, $|\text{Inv}(S)| = d(\mu_D)$. Since $|\text{Inv}(S)| \geq a(D)$ for every filling S such that $\text{Des}(S) = D$, we must have $d(\mu_D) \geq a(D)$, and consequently $d(\mu_D) - a(D)$ is a nonnegative integer. Therefore, by equation (5.7), giving a combinatorial formula for $\tilde{K}_{\lambda, \mu_D}^{(\mu_1)}(q)$ also gives a combinatorial formula for $\tilde{K}_{\lambda, \mu}(q, t)$.

5.2 A combinatorial formula for Kostka-Macdonald polynomials

Rather than use equation (5.7), we will directly describe a dual equivalence graph for fillings of a shape μ in order to obtain a combinatorial formula for $\tilde{K}_{\lambda, \mu}(q, t)$. The graphs we construct are just special cases of the constructions of Chapter 4.

For a filling S of the diagram μ , define the *spaced row reading word* of S , denoted $r(S)$, to be the row reading word of S augmented with \emptyset 's in each cell of $(\mu_1^{\mu_1})/\mu$. For example, the spaced row reading of S from Example 5.3 is

$$r(S) = 5 \ \emptyset \ \emptyset \ \emptyset \ \emptyset \ 11 \ 14 \ 9 \ 2 \ \emptyset \ 6 \ 3 \ 4 \ 10 \ \emptyset \ 8 \ 1 \ 13 \ 7 \ 12.$$

Define the *descent signature* $\sigma(S) \in \{\pm 1\}^{n-1}$ by

$$\sigma(S)_i = \begin{cases} +1 & \text{if } i \text{ appears to the left of } i+1 \text{ in } r(S) \\ -1 & \text{if } i+1 \text{ appears to the left of } i \text{ in } r(S) \end{cases}. \quad (5.8)$$

Note that if the filling S of μ with $\text{Des}(S) = D$ corresponds to the semi-standard k -tuple of tableaux T of shape μ_D , then the spaced row reading word of S is exactly the content reading word of T , and so we also have $\sigma(S) = \sigma(T)$ by equation (4.11).

Expressed in terms of quasi-symmetric functions, equation (5.6) becomes

$$\tilde{H}_\mu(x; q, t) = \sum_{S: \mu \rightarrow [n]} q^{\text{inv}(S)} t^{\text{maj}(S)} Q_{n, \sigma(S)}(x). \quad (5.9)$$

Note that here the sum is over all standard fillings S of μ .

Recall from Section 4.1 the k -descent set of a word defined in equation (4.6), and the set of k -inversions of a word defined in equation (4.4). For S a filling of the diagram of μ , we have

$$\text{Des}_{\mu_1}(r(S)) = \text{Des}(S), \quad (5.10)$$

$$\text{Inv}_{\mu_1}(r(S)) = \text{Inv}(S). \quad (5.11)$$

Furthermore, the arm and leg statistics for cells of S may be recovered from letters of $r(S)$ as follows. Let x be a nonempty letter of $r(S)$ which occurs at position $a \cdot k + b$. Then define $a(x)$ to be the number of nonempty letters $r_{a \cdot k + (b+1)}, \dots, r_{a \cdot k + (k-1)}$, and define $l(x)$ to be the number of nonempty letters $r_{(a-1) \cdot k + b}, \dots, r_{0 \cdot k + b}$. Then if a cell c of S corresponds to the letter x of $r(S)$, we have

$$a(x) = a(c) \quad \text{and} \quad l(x) = l(c).$$

Define the major index and inversion statistics for spaced words analogous to equations (5.3) and (5.5) by

$$\text{maj}_k(r) = |\text{Des}_k(r)| + \sum_{(x,y) \in \text{Des}_k(r)} l(x), \quad (5.12)$$

$$\text{inv}_k(r) = |\text{Inv}_k(r)| - \sum_{(x,y) \in \text{Des}_k(r)} a(x). \quad (5.13)$$

Then for S a filling of μ , we have

$$\text{maj}_{\mu_1}(r(S)) = \text{maj}(S), \quad (5.14)$$

$$\text{inv}_{\mu_1}(r(S)) = \text{inv}(S). \quad (5.15)$$

For each partition μ , $|\mu| = n$, define the set W_μ by

$$W_\mu = \{r(S) \mid S : \mu \xrightarrow{\sim} [n]\}. \quad (5.16)$$

The edges E_i constructed in Section 4.2 restrict to edges on W_μ . In particular, using the notation of Section 5.1, letting D range over descent sets of fillings of μ we have

$$W_\mu = \bigcup_D V_{\mu_D}^{(\mu_1)}$$

Since E_i preserves inv_k and Des_k for vertices of $V^{(k)}$, E_i will preserve inv and maj for vertices of W_μ . Therefore, letting $\mathcal{H}_\mu = (W_\mu, \sigma, E)$, we have the following extension of Conjecture 4.10.

Theorem 5.5. *Assuming Conjecture 4.10 holds, let \mathcal{H}_μ be the dual equivalence graph on the spaced row reading words of standard fillings of μ . For each partition λ , let $C_\mu(\lambda)$ denote the set of connected components of \mathcal{H}_μ which are isomorphic to \mathcal{G}_λ . Then*

$$\tilde{K}_{\lambda, \mu}(q, t) = \sum_{\mathcal{C} \in C_\mu(\lambda)} q^{\text{inv}(\mathcal{C})} t^{\text{maj}(\mathcal{C})}.$$

In particular, by Theorem 4.21, we have a combinatorial formula for $\tilde{K}_{\lambda, \mu}(q, t)$ when μ is a partition with at most 3 columns.

Bibliography

- [CL95] C. Carré and B. Leclerc, *Splitting the square of a Schur function into its symmetric and antisymmetric parts*, Journal of Algebraic Combinatorics **4** (1995), 201–231.
- [Ges84] I. Gessel, *Multipartite p -partitions and inner products of skew Schur functions*, Contemp. Math. **34** (1984), 289–317.
- [GH93] A. Garsia and M. Haiman, *A graded representation model for Macdonald’s polynomials*, Proc. Natl. Acad. Sci. **90** (1993), no. 8, 3607–3610.
- [GH07] I. Grojnowski and M. Haiman, *Affine Hecke algebras and positivity of LLT and Macdonald polynomials*, preprint, 2007.
- [Hag04] J. Haglund, *A combinatorial model for the Macdonald polynomials*, Proc. Natl. Acad. Sci. **101** (2004), no. 46, 16127–16131.
- [Hai92] M. Haiman, *Dual equivalence with applications, including a conjecture of Proctor*, Discrete Mathematics **99** (1992), 79–113.
- [Hai99] M. Haiman, *Macdonald polynomials and geometry*, New perspectives in geometric combinatorics, MSRI Publications **37** (1999), 207–254.
- [Hai01] M. Haiman, *Hilbert schemes, polygraphs, and the Macdonald positivity conjecture*, Journal of the American Mathematical Society **14** (2001), 941–1006.
- [HHL05a] J. Haglund, M. Haiman, and N. Loehr, *Combinatorial theory of Macdonald polynomials i: Proof of Haglund’s formula*, Proc. Natl. Acad. Sci. **102** (2005), no. 8, 2690–2696.
- [HHL⁺05b] J. Haglund, M. Haiman, N. Loehr, J. Remmel, and A. Ulyanov, *A combinatorial formula for the characters of the diagonal coinvariants*, Duke Math. J. **126** (2005), 195–232.
- [JK81] G. James and A. Kerber, *The representation theory of the symmetric group*, Encyclopedia of Mathematics and its Applications, vol. 16, Addison-Wesley Publishing Company, Reading, Massachusetts, 1981.

- [KMS95] M. Kashiwara, T. Miwa, and E. Stern, *Decomposition of q -deformed Fock spaces*, *Selecta Mathematica (N.S.)* **1** (1995), no. 4, 787–805.
- [LLT97] A. Lascoux, B. Leclerc, and J.Y. Thibon, *Ribbon tableaux, Hall-Littlewood functions, quantum affine algebras, and unipotent varieties*, *J. Math. Phys.* **38** (1997), 1041–1068.
- [LT00] B. Leclerc and J.Y. Thibon, *Littlewood-Richardson coefficients and Kazhdan-Lusztig polynomials*, *Adv. Stud. Pure Math.* **28** (2000), 155–220.
- [Mac88] I. G. Macdonald, *A new class of symmetric functions*, *Actes du 20e Seminaire Lotharingien* **372** (1988), 131–171.
- [Mac95] I. G. Macdonald, *Symmetric functions and Hall polynomials*, second edition ed., Clarendon Press, Oxford University Press, New York, 1995.
- [SSW03] M. Shimozono, A. Schilling, and D. White, *Branching formula for q -Littlewood-Richardson coefficients.*, *Adv. in Appl. Math.* **30** (2003), no. 1-2, 258–272.
- [SW85] D. Stanton and D. White, *A Schensted algorithm for rim hook tableaux*, *Journal of Combinatorial Theory* **40** (1985), 211–247.
- [vL05] M. van Leeuwen, *Spin-preserving Knuth correspondences for ribbon tableaux*, *Electronic Journal of Combinatorics Research Paper* **10** (2005), 65pp.