Dual Equivalence Graphs and their Applications

Austin Roberts

A dissertation submitted in partial fulfillment of the
requirements for the degree of

Doctor of Philosophy

University of Washington

2014

Reading Committee:

Sara Billey, Chair

William M. McGovern

Isabella Novik

Program Authorized to Offer Degree:
Department of Mathematics

University of Washington

Abstract

Dual Equivalence Graphs and their Applications

Austin Roberts

Chair of the Supervisory Committee:
Professor Sara Billey
Department of Mathematics

In 2007 Sami Assaf introduced dual equivalence graphs as a method for demonstrating that a quasisymmetric function is Schur positive. The method involves the creation of a graph whose vertices are weighted by Ira Gessel's fundamental quasisymmetric functions in such a way that the sum of the weights of a connected component is a single Schur function. The graphs are termed dual equivalence graphs, and this dissertation is the compilation of works that focus on the further development of the theory of said graphs. This work further includes applications to Macdonald polynomials, Hall-Littlewood polynomials, and Lascoux-Leclerc-Thibon polynomials. In joint work with Sara Billey, Zach Hamaker, and Benjamin Young, we also give a generalization of dual equivalence graphs to the Coxeter-Knuth graph of Lie type B and illustrate the relationship of these graphs to a newly defined type B Little bump. For the sake of completeness, we also include an appendix providing a proof of the original axiomatization of dual equivalence graphs as described by Sami Assaf.

TABLE OF CONTENTS

ACKNOWLEDGMENTS

There are many people that have helped to stimulate the ideas and develop the results contained in this dissertation. We would like to express great gratitude to David Anderson, Sami Assaf, Jonathan Browder, Nantel Bergeron, Avinash Dalal, Adriano Garsia, William McGovern, Leigh Mercer, Jenifer Morse, James Haglund, Zachary Hamaker, Isabella Novik, Brendan Pawlowski, Jair Taylor, Benjamin Young, and Mike Zabrocki for all of the conversation, ideas, and time that have guided this work forward.

Particular thanks belong to Sara Billey for her invaluable guidance throughout the process of creating this dissertation. From the original development of research interests, to navigating the mathematical community, through to the development of high writing standards, none of this could have been done without her.

Finally, I would like to thank those people that sat in coffee shops with me while I put endless hours into this work: Tominda Adkins, Nick Brown, Francesca Cannata-Bowman, Kaia Chessen, Lindsay Erickson, Toby Johnson, Stephen Lewis, Chris McMurdie, James Pfeiffer, Erik Slivken, Simon Spicer, Olga Trichtenko, Ashley Wearly, and many others; I surely would have have lost myself without them.

Chapter 1

THE INTRODUCTION

This dissertation is primarily concerned with the use of dual equivalence graphs as a tool for finding Schur expansions of symmetric functions and as a tool for describing Coxeter-Knuth classes. Each of the following three chapters is taken from a previous paper, [Roberts, 2014], [Roberts, 2013a], and [Billey et al.], respectively. The appendix was used in partial completion of a Master's thesis. As such, each chapter has its own introduction and notation. This preliminary introduction will thus attempt a less formal overview of the topics to be covered.

1.1 Symmetric Functions and the Schur Basis

A symmetric function f is a formal power series in infinitely many variables $X = x_1, x_2, x_3, \ldots$ with finite degree and the following additional symmetry property. If $\pi \in S_m$ is any permutation, we may let π act on f via its action on the indices of $x_1, x_2, \ldots, x_m$. We say that f is a symmetric function if f is invariant under the action of all permutations π in all finite symmetric groups S_m. For example,

$$f(X) = \sum_{i<j} x_i^2 x_j + \sum_{i<j} x_i x_j^2$$

is a symmetric function. The vector space of symmetric functions is denoted by Λ. We will often assume that this vector space is over the field $\mathbb{Q}$ or $\mathbb{R}$, but that distinction will be unimportant for this treatment.

The vector space Λ may be regarded as a graded ring by using the degree of monomials as a grading and power series multiplication. We may then try to find a basis for the vector space Λ_n of homogenous functions whose terms all have degree n. The dimension of this subspace is equal to the number of partitions of size n. That is, the number of weakly decreasing sequences of positive integers whose sum is equal to n. Partitions are denoted by $\lambda = (\lambda_1, \lambda_2, \ldots, \lambda_k)$, where $\lambda_i \geq \lambda_{i+1} > 0$, and $|\lambda| = \sum \lambda_i = n$. We will also say that λ is a partition of n, denoted $\lambda \vdash n$.

The vector space Λ_n has several important bases. We will primarily be concerned with the basis of *Schur functions* $\{s_\lambda\}$, where $\lambda \vdash n$. While there are many useful ways of expressing s_λ, we will focus on two that are related to the combinatorics of tableaux. Given a partition λ, the *diagram* of λ is formed by placing boxes on the integer Cartesian plane, with λ_i consecutive boxes in row i starting on the vertical axis, as shown in Figure 1.1. We will often refer to boxes of λ by identifying λ with its diagram.

Figure 1.1: The diagrams for (4,3,2,2) and (5,3,3,1), from left to right.

A *semistandard Young tableaux* of shape λ is a function that assigns a positive integer to each box of λ such that the rows are weakly increasing left to right and strictly increasing up columns, as shown in Figure 1.2. The set of all semistandard Young tableaux of shape λ is denoted SSYT(λ). Given $T \in \mathrm{SSYT}(\lambda)$ with α_i many i's, we let $x^T = \prod x_i^{\alpha_i}$. We may then define the Schur function s_λ by

$$s_\lambda := \sum_{T \in \mathrm{SSYT}(\lambda)} x^T. \tag{1.1.1}$$

Proofs that the Schur functions are symmetric and form a basis for Λ, as well as broader treatment of symmetric functions, can be found in [Fulton, 1997], [Sagan, 1991], and [Stanley, 2001], amongst others.

For combinatorial purposes, it is often useful to replace the infinite set SSYT(λ) with the finite set of standard Young tableaux, SYT(λ), the cardinality of which is denoted f^λ. Here, a standard Young tableau is a bijective function from $\lambda \vdash n$ to $[n] = \{1, 2, \ldots, n\}$ whose entries are strictly increasing up columns and across rows from left to right, as shown in

Figure 1.2: On the left, a semistandard tableau in SSYT((4, 3, 2)). On the right, a standard tableau in SYT((4, 3, 2)).

Figure 1.2. We may then give an equivalent definition for s_λ by

$$s_\lambda := \sum_{T \in \mathrm{SYT}(\lambda)} F_{\sigma(T)}(X). \tag{1.1.2}$$

where $F_{\sigma(T)}(X)$ is Ira Gessel's fundamental quasisymmetric function associated to T. The fundamental quasisymmetric functions can, in turn, be defined as

$$F_\sigma(X) := \sum_{\substack{i_1 \leq \cdots \leq i_n \\ i_j = i_{j+1} \Rightarrow \sigma_j = +1}} x_{i_1} \cdots x_{i_n},$$

where the signature function σ gives a sequence of $+$'s and $-$'s, usually defined via inverse descent sets of permutations. Numerous symmetric functions have known expansions in terms of $F_{\sigma(T)}(X)$, including the polynomials in Section 1.3.1, plethysms of Schur functions in [Loehr and Warrington, 2012], and k-Schur functions in [Assaf and Billey, 2012]. For this reason, we will focus much attention on using (1.1.2) to expand such symmetric functions into the basis of Schur functions.

1.2 The Role of Dual Equivalence Graphs

As mentioned, this dissertation is primarily concerned with the theory and applications of dual equivalence graphs. Dual equivalence graphs, in turn, are primarily concerned with describing expansions of symmetric functions into the basis of Schur functions.

The technique of applying dual equivalence graphs can be described as follows. Suppose that a symmetric function can be defined as a sum of fundamental quasisymmetric functions $f(X) = \sum F_\sigma(X)$. In particular, this sum is usually indexed by some combinatorial set, for instance, the set of standard Young tableaux of some partition shape λ. The first step is to define a set of involutions E_i on the indexing set. There is then a *signed colored graph* $\mathcal{G}$ with vertices defined by the indexing set of the sum, edge sets labeled by i as defined by each E_i, and an additional label by a $+/-$ string, called the signature, given to each vertex corresponding to σ in the above sum. The prototypical example using standard Young tableaux as vertices can be seen in Figure 1.3. The question is, can such a graph be used to find a Schur expansion? It was proven in [Assaf, 2013], and improved upon in Chapter 2 of this dissertation, that if a signed colored graph $\mathcal{G}$ satisfies a certain set of axioms, then each component of $\mathcal{G}$ corresponds to a single Schur function, as uniquely determined by the isomorphism type of the component. A signed colored graph satisfying these axioms is termed a *dual equivalence graph*. In particular, if $V_1, V_2, \ldots, V_k$ are the vertex sets for each

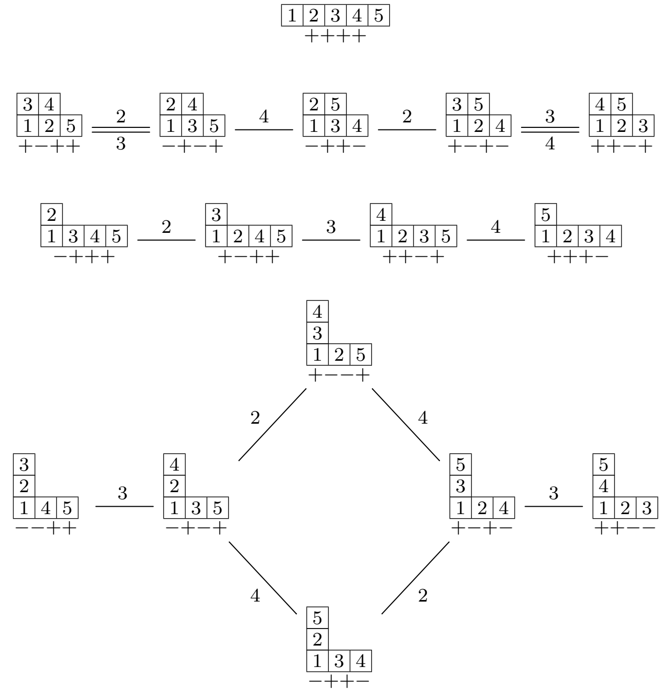

Figure 1.3: Examples of dual equivalence graphs with standard Young tableaux as vertices.

component of a dual equivalence graph $\mathcal{G}$, then

$$f(X) = \sum_{i=1}^{k} \sum_{v \in V_i} F_{\sigma(v)}(X) = \sum_{i=1}^{k} s_{\lambda(i)}, \tag{1.2.1}$$

where $s_{\lambda(i)}$ is the Schur function corresponding to the isomorphism type of the component with vertex set V_i. Applications of this technique can be found in Chapters 2 and 3, as well as in [Assaf et al., 2012], [Assaf and Billey, 2012], [Assaf, 2013], [Hamaker and Young, 2013], and [Chmutov, 2013].

1.3 Applications of Schur Expansions

The question remains, why do we care about Schur expansions? One answer is that they have proven to be both interesting and challenging on a strictly mathematical level, providing many difficult problems open to a wide range of approaches. It is worth, however, taking a moment to justify the pursuit of Schur expansions externally, mentioning a few of the topics where they have emerged in important roles, though a full treatment of each topic is beyond the scope of this dissertation.

1. S_n and $GL_n(\mathbb{C})$ modules:

   Schur functions arise as the characters of irreducible polynomial $GL_n(\mathbb{C})$ representations. Hence, if f is the character of some $GL_n(\mathbb{C})$ representation, then coefficients of the Schur expansion of f give the multiplicity of each of the irreducible subrepresentations.

   Similarly, there exists a homomorphism sending each irreducible S_n representation to a unique Schur function. If we consider any symmetric function f that can be expressed as a nonnegative integer sum of Schur functions, then f corresponds to a unique S_n representation (up to isomorphism) under the inverse of said morphism. For more information, see [Fulton, 1997] or [Sagan, 1991].

2. Enumeration of reduced words:

   Each permutation $\pi \in S_n$ can be written as a sequence of adjacent transpositions that, when composed, multiply to π. The minimal length of such a sequence is termed the *length* of π. In this case, the sequence is termed a *reduced word* for π. The question is, how many reduced words does π have? Richard Stanley answered this question by associating each π to a sum of fundamental quasisymmetric functions, with each term associated to a reduced word for π. Said function can be expressed as a sum of Schur functions, say $\sum c_\lambda s_\lambda$. Since s_λ is the sum of f^λ many fundamental quasisymmetric functions, the number of reduced words of π can then be computed as $\sum c_\lambda f^\lambda$. For more information, see [Stanley, 1984].

3. Schubert calculus:

   There are several ways that one can generalize Schur functions. One important way connects Schur functions to cohomology classes of Grassmannian manifolds. More generally, the set of Schur functions are a subset of the set of Schubert polynomials, each of which corresponds to the cohomology class of a Schubert variety inside of a flag variety. In this setting, multiplication of Schubert polynomials corresponds to the cup product of cohomology classes, which in turn corresponds to the intersection of varieties. For more information, see [Fulton, 1997].

4. Less combinatorial topics:

   We conclude the list by pointing out a few works with a more probabilistic, or physics oriented focus. First, Schur processes assign probability measures to sequences of partitions via evaluating each Schur function at some fixed real values of the variables X. Similar statements may be made about any nonnegative integer sum of Schur functions. For more information, see [Borodin and Corwin, 2011] and [Borodin and Gorin, 2012]. Symmetric functions and their expansions into various bases arise in the study of gauge theory, as can be seen in [Zeidler, Eberhard, 2011] and [Gadde et al., 2013]. A particular product of Schur functions known as a plethysm also arises in the study of atomic spectroscopy in [Wybourne, 1970].

1.3.1 The Polynomials of Chapters 2 and 3

We pause to give special attention to the functions most directly addressed in this dissertation: Macdonald polynomials, Lascoux-Leclerc-Thibon (LLT) polynomials, and Hall-Littlewood polynomials. Each of these polynomials is either a q-symmetric function or a q, t-symmetric function. By this we mean that each term has some $q^a t^b$ as a factor, where q and t are indeterminates and a and b are nonnegative integers.

The Macdonald polynomials were introduced in [Macdonald, 1988] and are often defined as the set of q, t-symmetric functions that satisfy certain orthogonality and triangularity conditions. The modified Macdonald polynomials were shown to be Schur positive by Mark Haiman via representation-theoretic and geometric means in [Haiman, 2001]. Macdonald polynomials also specialize to several well known functions, including Hall-Littlewood polynomials, Jack polynomials, and Schur functions. While finding a combinatorial description of the Schur expansion remains an important open problem, there are several cases for which such a descriptions is known. The Schur expansion of modified Macdonald polynomials indexed by shapes with strictly less than three columns was described in [Haglund et al., 2005a], which in turn drew on the earlier work in [Carré and Leclerc, 1995], [van Leeuwen, 2001]. The first combinatorial description of the two column case was given in [Fishel, 1995], but others were subsequently given in [Zabrocki, 1998], [Lapointe and Morse, 2003], and [Assaf, 2008/09]. In addition, an algorithm for finding the Schur expansion of Macdonald polynomials indexed by shapes with at most four cells in the first row and at most two cells in the second row was given in [Zabrocki, 1999]. Finding a combinatorial interpretation for the three column case is still an open problem, though there is a conjectured formula in [Haglund, 2004]. We will provide a description in the case where the first row has less than four cells and the second row has less than three in Chapter 2, along with that of a family of generalized Macdonald polynomials in Chapter 3. For details on the combinatorics of Macdonald polynomials, see [Haglund, 2008].

Lascoux-Leclerc-Thibon polynomials were first introduced in [Lascoux et al., 1997] as a q-analogue to products of Schur functions and were later given a description in terms of tuples of skew tableaux in [Haglund et al., 2005b]. These functions are also related to the Fock space representation of $U_q\left(\widehat{\mathfrak{sl}}_n\right)$. Importantly for our purposes, it was shown in [Haglund et al., 2005a] that modified Macdonald polynomials may be expressed in terms of LLT polynomials. In some cases, nice Schur expansions for LLT polynomials are already known. In particular, the set of LLT polynomials indexed by two skew shapes was described in [Carré and Leclerc, 1995] and [van Leeuwen, 2001]. In Chapter 2, we expand this result to a family of LLT polynomials strictly containing this case.

As just noted, Macdonald polynomials specialize to Hall-Littlewood polynomials, specifically by letting $q = 0$. Hall-Littlewood polynomials, in turn, specialize to the Schur functions, Schur Q-functions, and the monomial symmetric functions. They were first studied by Philip Hall in relation to the Hall algebra in [Hall, 1957] and later by D.E. Littlewood in [Littlewood, 1961]. It should be noted that the earliest known work on Hall-Littlewood polynomials actually dates back to the lectures of Ernst Steinitz in [Steinitz, 1901]. Hall-Littlewood polynomials have proven to be a rich mathematical topic, with recent combinatorial work including (but certainly not limited to) [Nakayashiki and Yamada, 1997], [Carbonara, 1998], [Dalal and Morse, 2012], and [Loehr et al., 2013]. Expanding modified Hall-Littlewood polynomials into Schur functions can be achieved via the charge statistic, as found in [Lascoux and Schützenberger, 1978], though we will present a new expansion in Chapter 3 as a sum over a subset of the Yamanouchi words. Equivalently, our result gives a new combinatorial rule for the coefficients of the Kostka-Foulkes polynomial in one variable t. The coefficient of this polynomial also gives the unipotent character χ^λ of $GL_n(\mathbb{F}_t)$ on a unipotent element with Jordan canonical form specified by the indexing partition μ of the Hall-Littlewood polynomial. For more details on the topic of Hall-Littlewood polynomials, see [Macdonald, 1995]. For connections between LLT and Hall-Littlewood polynomials, see [Lascoux et al., 1997].

1.4 Organization and Chapter Topics

1.4.1 Chapter 2

In this chapter, we improve on Assaf's axiomatization of dual equivalence graphs, giving locally testable criteria that are more easily verified by computers. We further advance the theory of dual equivalence graphs by describing a broader class of graphs that correspond to an explicit Schur expansion in terms of Yamanouchi words. Along the way, we demonstrate several symmetries in the structure of dual equivalence graphs. We then apply these techniques to give explicit Schur expansions for a family of Lascoux-Leclerc-Thibon polynomials. This family properly contains the previously known case of polynomials indexed by two skew shapes, as was described in a 1995 paper by Christophe Carré and Bernard Leclerc. As an immediate corollary, we gain an explicit Schur expansion for a family of modified Macdonald polynomials in terms of Yamanouchi words. This family includes all polynomials indexed by shapes with at most three cells in the first row and at most two cells in the second row, providing an extension to the combinatorial description of the two column case described in 2005 by James Haglund, Mark Haiman, and Nick Loehr.

1.4.2 Chapter 3

This chapter uses the theory of dual equivalence graphs to give explicit Schur expansions for several families of symmetric functions. We begin by giving a combinatorial definition of the modified Macdonald polynomials and modified Hall-Littlewood polynomials indexed by any diagram $\delta \subset \mathbb{Z} \times \mathbb{Z}$, written as $\widetilde{H}_\delta(X;q,t)$ and $\widetilde{H}_\delta(X;0,t)$, respectively. We then give an explicit Schur expansion of $\widetilde{H}_\delta(X;0,t)$ as a sum over a subset of the Yamanouchi words, as opposed to the expansion using the charge statistic given in 1978 by Lascoux and Schüztenberger. We further define the symmetric function $R_{\gamma,\delta}(X)$ as a refinement of $\widetilde{H}_\delta(X;0,t)$ and similarly describe its Schur expansion. We then analyze $R_{\gamma,\delta}(X)$ to determine the leading term of its Schur expansion. We also provide a conjecture towards the Schur expansion of $\widetilde{H}_\delta(X;q,t)$. To achieve these results, we use a construction from the 2007 work of Sami Assaf to associate each Macdonald polynomial with a signed colored graph $\mathcal{H}_\delta$. In the case where a subgraph of $\mathcal{H}_\delta$ is a dual equivalence graph, we provide the Schur expansion of its associated symmetric function, yielding several corollaries.

1.4.3 Chapter 4

Dual equivalence is typically defined on permutations, and it may be defined on words via a standardization operator. Of crucial importance in this definition is the interplay of the R-S-K correspondence, Knuth equivalence, and dual equivalence. We may, however, consider what the appropriate action of dual equivalence is when acting on a reduced word of a root system. In Lie types A, B, and D, there are insertion algorithms that play the role of the R-S-K insertion algorithm, while Coxeter-Knuth equivalence plays the role of Knuth equivalence. In type A, significant work has already been done in [Hamaker and Young, 2013]. This work is generalized to the type B case in Chapter 4, as part of joint work with Sara Billey, Zachary Hamaker, and Benjamin Young. In the type A case, the results largely mirror the known results for standard dual equivalence graphs. In types B, however, the setting changes from tableaux on straight shapes to shifted tableaux, as discussed in [Haiman, 1992] and [Kraskiewcz, 1989].

Chapter 4 is also dedicated to the development of a Little bump in the type B setting and its interaction with Kraśkiewcz insertion. In this setting, the Little bump provides a map from a reduced word of a signed permutation to the reduced word of a unique isotropic Grassmanian permutation, while Kraśkiewcz insertion realizes each reduced word as a pair of shifted tableaux. As in the type A case, the little bump acts as an isomorphism on Coxeter-Knuth graphs, providing a straight forward proof that each such graph is isomorphic to a dual equivalence graph upon appropriately labeling vertices and edges.

1.4.4 Appendix

In the appendix, we provide a proof of Sami Assaf's original axiomatization of dual equivalence graphs. We include this work for the sake of completion, since Chapter 2 uses it as a stepping stone for further results.

Chapter 2

DUAL EQUIVALENCE GRAPHS REVISITED AND THE EXPLICIT SCHUR EXPANSION OF A FAMILY OF LLT POLYNOMIALS

2.1 Introduction

Dual equivalence was developed and applied by Mark Haiman in [Haiman, 1992] as an extension of work done by Donald Knuth in [Knuth, 1970]. Sami Assaf then introduced the theory of dual equivalence graphs in her Ph.D. dissertation [Assaf, 2007] and subsequent preprint [Assaf, 2013]. In these papers, she is able to associate a number of symmetric functions to dual equivalence graphs and each component of a dual equivalence graph to a Schur function, thus demonstrating Schur positivity. More recently, variations of dual equivalence graphs are given for k-Schur functions in [Assaf and Billey, 2012] and for the product of a Schubert polynomial with a Schur polynomial in [Assaf et al., 2012].

A key connection between dual equivalence graphs and symmetric functions is the ring of quasisymmetric functions. The quasisymmetric functions were introduced by Ira Gessel in [Gessel, 1984] as part of his work on P-partitions. Currently there are a number of functions that are easily expressed in terms of Gessel's *fundamental quasisymmetric functions* that are not easily expressed in terms of Schur functions. For example, such an expansion for plethysms is described in [Loehr and Warrington, 2012], for Lascoux-Leclerc-Thibon (LLT) polynomials in [Haglund et al., 2005b], for Macdonald polynomials in [Haglund et al., 2005a], and conjecturally for the composition of the nabla operator with an elementary symmetric function in [Haglund et al., 2005b]. An expressed goal of developing the theory of dual equivalence graphs is to create a tool for turning such quasisymmetric expansions into explicit Schur expansions.

Previously, dual equivalence graphs were defined by five *dual equivalence axioms* that are locally testable and one that is not. One of the main results of this paper is to give an equivalent definition using only local conditions, as stated in Theorem 2.3.17. Many graphs, while not satisfying all of these axioms, correspond to Schur positive expansions. In particular, those admitting a morphism onto a dual equivalence graph, as described in Definition 2.2.6, are necessarily Schur positive. In Theorems 2.3.13 and 2.3.14 we give a classification of the set of graphs admitting such a morphism and obeying the first dual equivalence axiom. In particular, Theorem 2.3.13 gives an explicit Schur expansion for the symmetric functions associated to such graphs in terms of standardized Yamanouchi words.

The paper concludes by applying the above results to LLT polynomials in Theorem 2.4.3. LLT polynomials were first introduced in [Lascoux et al., 1997] as a q-analogue to products of Schur functions and were later given a description in terms of tuples of skew tableaux in [Haglund et al., 2005b]. Corollary 2.4.4 then applies the results of [Haglund et al., 2005a] to give an explicit combinatorial description for a family of modified Macdonald polynomials. First introduced in [Macdonald, 1988], Macdonald polynomials are often defined as the set of q,t-symmetric functions that satisfy certain orthogonality and triangularity conditions, as is well described in [Macdonald, 1995]. Part of the importance of Macdonald polynomials derives from the fact that they specialize to a wide array of well known functions, including Hall-Littlewood polynomials and Jack polynomials (see [Macdonald, 1995] for details). In [Haiman, 2001], Mark Haiman used geometric and representation-theoretic techniques to prove that Macdonald polynomials are Schur positive.

In some cases, nice Schur expansions for LLT and Macdonald polynomials are already known. In particular, the set of LLT polynomials indexed by two skew shapes was described in [Carré and Leclerc, 1995] and [van Leeuwen, 2001], and modified Macdonald polynomials indexed by shapes with strictly less than three columns was described in [Haglund et al., 2005a] (which in turn drew on the earlier work in [Carré and Leclerc, 1995], [van Leeuwen, 2001]). The first combinatorial description of the two column case was given in [Fishel, 1995], but others were subsequently given in [Zabrocki, 1998], [Lapointe and Morse, 2003], and [Assaf, 2008/09]. In addition, an algorithm for finding the Schur expansion of Macdonald polynomials indexed by shapes with at most four cells in the first row and at most two cells in the second row was given in [Zabrocki, 1999]. Finding a combinatorial interpretation for the three column case is still an open problem, though there is a conjectured formula in [Haglund, 2004].

This paper is broken into sections as follows. Section 2.2 reviews the necessary material on partitions, tableaux, the Robinson-Schensted-Knuth correspondence, and symmetric functions, before giving the necessary background on dual equivalence graphs. Section 2.3 is dedicated to further developing the theory of dual equivalence graphs, culminating in a new axiomatization for dual equivalence graphs in Theorem 2.3.17. Section 2.4 applies the results of Section 2.3 to LLT polynomials and Macdonald polynomials. The graph structure given to LLT polynomials in [Assaf, 2013] is reviewed before Theorem 2.4.11 classifies the set of LLT polynomials that correspond to dual equivalence graphs. Theorem 2.4.3 states that said set of LLT polynomials have a Schur expansion indexed by standardized Yamanouchi words. This set strictly contains the set of LLT polynomials indexed by two skew shapes. Corollary 2.4.4 then gives a Schur expansion for modified Macdonald polynomials indexed by partition shapes with strictly less than four boxes in the first row and strictly less than three boxes in the second row.

2.2 Preliminaries

This section is dedicated to introducing the key notation and definitions that underlie the rest of the paper. Particular attention is given to known results about dual equivalence graphs.

2.2.1 Tableaux

A *partition* λ is a weakly decreasing finite sequence of nonnegative integers $\lambda_1 \geq \ldots \geq \lambda_k \geq 0$. If $\sum \lambda_i = n$, we say that λ is a partition of n and write $\lambda \vdash n$. Partitions are often expressed in terms of diagrams where λ_i is the number of boxes, or *cells*, in the i^{th} row, from bottom to top, as in the left diagram of Figure 2.1. It is sometimes useful to treat a diagram as a subset of the integer Cartesian plane with the bottom left corner of the diagram at the origin. Given a partition λ, the *conjugate partition* of λ, denoted $\tilde{\lambda}$, is defined by $\tilde{\lambda}_i := |\{j : \lambda_j \geq i\}|$. The diagram of $\tilde{\lambda}$ is obtained by reflecting the diagram of λ over the the main diagonal $x = y$ in the Cartesian plane.

If the diagram of ρ is contained in the diagram of λ, equivalently $\rho_i \leq \lambda_i$ for all i, then we may consider the *skew diagram* λ/ρ defined by omitting the boxes of ρ from λ, as in the

Figure 2.1: The diagrams for (4,3,2,2) and (4,3,1,1)/(2,2)

right skew diagram of Figure 2.1. Here, λ/ρ is referred to as the *shape* of the skew diagram. The number of cells of λ/ρ is called the *size of* λ/ρ and is denoted by $|\lambda/\rho|$. If we need to distinguish the shape of a partition from a skew shape, we will refer to it as a *straight shape.* We say that μ is a *subdiagram* of λ/ρ if some translation of μ is contained in λ/ρ when considered as subsets of the Cartesian plane.

A *filling* assigns a positive integer to each cell of a partition or skew shape, usually written inside of the cell. Any filling of $\lambda \vdash n$ that assigns each value in $[n] = \{1, \ldots, n\}$ exactly once is termed a *bijective filling.* We will primarily be concerned with *standard Young tableaux*, or *tableaux* for short, which are bijective fillings that are also increasing up columns and across rows from left to right (see Figure 2.2). The set of all standard Young tableaux of shape λ is denoted $\mathrm{SYT}(\lambda)$. The union of all $\mathrm{SYT}(\lambda)$ over $\lambda \vdash n$ is denoted $\mathrm{SYT}(n)$. Similarly, the set of all skew tableaux of shape λ/ρ is denoted $\mathrm{SYT}(\lambda/\rho)$. In general, all tableaux will be assumed to be fillings of straight shapes unless stated otherwise.

The notion of a standard Young tableau extends to fillings of skew shapes, creating *skew tableaux*, as seen in the right side of Figure 2.2. The well known process of *jeu de taquin* slides gives a map from any skew tableau to a straight tableau via a sequence of slide operations that move cells either west or south depending on the values in the filling. For more information, see [Fulton, 1997, Part I], [Sagan, 1991, Ch. 3], or [Stanley, 2001, Ch. 7]. Given a (possibly skew) tableau T, define $\mathrm{sh}(T)$ to be the shape of the underlying diagram of T.

The *content* of a cell x, denoted $c(x)$, is $j - i$, where j is the column of x and i is the row of x in Cartesian coordinates. In other words, each diagonal going southwest to northeast has the same content, with the uppermost diagonal having the smallest content. In a standard Young tableau, the 1-cell is located at the origin of the Cartesian plane, and so has content

0. A connected skew shape having at most one cell of each content is called a *ribbon.*

Define the *content reading word* of a tableau as the word retrieved by reading off each entry from lowest content to highest, moving northeast along each diagonal, as in Figure 2.2. We also define the *row reading word* of a tableau by reading across rows from left to right, starting with the top row and working down. The content reading word and row reading word of a standard Young tableau are necessarily permutations.

Figure 2.2: On the left, a tableau with content reading word 438162957 and row reading word 483691257. On the right, a skew tableau with content reading word 3214 and row reading word 3214.

The *signature* of a word (or permutation in one-line notation) is a string of 1's and -1's, or $+$'s and $-$'s for short, where there is a $+$ in the i^{th} position if and only if i comes before $i+1$ in the word. If $\sigma_i(w) = -1$, then i is referred to as an *inverse descent of* w. Notice that a word is one entry longer than its signature. We may then define the signature of a tableau T, denoted $\sigma(T)$, as the signature of the content reading word of T. For example, the signatures of the tableau in Figure 2.2 are $+--+-+-+$ and $--+$, respectively. Given a fixed tableau, the row reading word and the content reading word always have the same signature (see [Stanley, 2001, Ch.7] for details).

2.2.2 Knuth Equivalence and the R-S-K Correspondence

While we assume familiarity with the Robinson-Schensted-Knuth (R-S-K) correspondence and jeu de taquin, we will use this section as a refresher and to set notation. For a full treatment, see [Fulton, 1997, Ch. 2-4], [Sagan, 1991, Ch. 3], or [Stanley, 2001, Ch. 7].

The R-S-K correspondence gives a bijection between permutations in S_n and pairs of standard Young tableaux (P, Q), where P and Q have the same shape $\lambda \vdash n$. The first tableau is called the *insertion tableau* and the latter is termed the *recording tableau.* For the duration of this paper, $P\colon S_n \to \mathrm{SYT}(n)$ and $Q\colon S_n \to \mathrm{SYT}(n)$ will be the functions taking a permutation to its insertion tableau and recording tableau, respectively. These two functions are related by

$$Q(w) = P(w^{-1}). \tag{2.2.1}$$

A detailed proof of this fact can be found in [Fulton, 1997, Ch. 4.1]. We then write $\mathrm{sh}(w)$ to mean $\mathrm{sh}(P(w))$.

For each tableau T with entries in $[n]$, the set of permutations in S_n sent to T by P is termed a *Knuth equivalence class.* Two words in the same Knuth equivalence class are said to be *Knuth equivalent.* The equivalence relations of Knuth classes are generated by the *fundamental Knuth equivalences*, denoted K_j for $1 < j < n$. Each K_j is defined as an involution that fixes all entries of $w \in S_n$ except for those with indices $j-1, j$, and $j+1$. Its action on these three entries can be written as,

$$\begin{aligned} K_j(\ldots xyz \ldots) &= (\ldots xyz \ldots), & K_j(\ldots zyx \ldots) &= (\ldots zyx \ldots), \\ K_j(\ldots yxz \ldots) &= (\ldots yzx \ldots), & K_j(\ldots xzy \ldots) &= (\ldots zxy \ldots), \end{aligned} \tag{2.2.2}$$

where $x < y < z$. In words, if the $j-1, j$, and $j+1$ entries are not strictly increasing or strictly decreasing, then switch the location of the two extreme values.

A number of important constructions yield words from the same Knuth class. Given a tableau T with row reading word w, then $P(w) = T$. The same can be shown to be true for the content reading word of T, demonstrating that row and content reading words are Knuth equivalent. In fact, the row and content reading words of a skew tableau are also Knuth equivalent. Further, the row reading words (as well as content reading words) of two skew tableaux related by a sequence of jeu de taquin slides are Knuth equivalent. In particular, if v and w are the row reading words of skew tableaux that are related by jeu de taquin, then $\mathrm{sh}(v) = \mathrm{sh}(w)$. It also follows that the row reading words of distinct tableaux (on straight shapes) are in different Knuth classes and that there is exactly one such word per class (see [Fulton, 1997, Ch. 2.1] for the details of the proof).

Next, we comment on the relationship between $\mathrm{sh}(w)$ and subwords of w, as is well presented in [Fulton, 1997, Ch. 3]. If $\mathrm{sh}(w) = \lambda$, then the longest increasing subword of

w has length λ_1, and the longest decreasing subword of w has length $\tilde{\lambda}_1$. For instance, if $w = 15342$, then $\mathrm{sh}(w) = (3, 1, 1)$, the longest increasing subword is 134, and the longest decreasing words are 532 and 542. In particular, if two words are Knuth equivalent, then both of their longest increasing subwords have the same length. Furthermore, if w and v are Knuth equivalent words in S_n, we may consider the restrictions of w and v to the consecutive values in some set $S = \{a, a+1, \ldots, b\}$, where $1 \leq a < b \leq n$. Call these two subwords w_S and v_S, respectively. Then w_S and v_S are Knuth equivalent, and so the longest increasing subwords of w_S and v_S both have the same length. The proof of this last fact can be found in [Fulton, 1997, Lem. 3].

Lastly, we define a particularly nice Knuth class. Let U_λ denote the tableau of shape $\lambda \vdash n$ formed by filling cells with values 1 through n row by row from bottom to top. Define $\mathrm{SYam}(\lambda)$ to be the set of $w \in S_n$ such that $P(w) = U_\lambda$. There is, however, a more direct way of deriving this set. A *Yamanouchi word* has entries in the positive integers such that when read backwards there are always more 1's than 2's, more 2's than 3's, and more i's than $i+1$'s. For instance, 25432431121 is a Yamanouchi word, but 231321 is not. The set $\mathrm{Yam}(\lambda)$ consists of all Yamanouchi words where 1 occurs λ_1 times, 2 occurs λ_2 times, and so on. We may *standardize* a word in $\mathrm{Yam}(\lambda)$ by replacing all of the 1's with $1, \ldots, \lambda_1$ in increasing reading order, all of the 2's with $\lambda_1 + 1, \ldots, \lambda_2$ in reading order, et cetera. We call the resulting words *standardized Yamanouchi words*. It is a simple exercise to verify that the set of standardized Yamanouchi subwords derived from $\mathrm{Yam}(\lambda)$ is precisely $\mathrm{SYam}(\lambda)$.

2.2.3 Symmetric Functions

The ring of symmetric functions has several well-known bases with ties to tableaux, as is well laid out in [Stanley, 2001, Ch. 7], [Fulton, 1997, Part I], or [Sagan, 1991, Ch. 4]. Of primary importance is the basis of Schur functions, denoted $\{s_\lambda\}$. We will take the unorthodox approach of defining these functions using a result of Ira Gessel. While less immediately intuitive than standard approaches, this definition contains the only properties that we need. First, a preliminary definition:

Definition 2.2.1. Given any signature $\sigma \in \{\pm 1\}^{n-1}$, define the *fundamental quasisymmet-*

ric function $F_\sigma(X) \in \mathbb{Z}[x_1, x_2, \ldots]$ by

$$F_\sigma(X) := \sum_{\substack{i_1 \leq \cdots \leq i_n \\ i_j = i_{j+1} \Rightarrow \sigma_j = +1}} x_{i_1} \cdots x_{i_n}.$$

The set of fundamental quasisymmetric functions of degree n forms a homogeneous basis for the vector space of degree n quasisymmetric functions. The *ring of quasisymmetric functions* is created by allowing formal multiplication as power series. The extent that we need this ring to motivate our results is limited to a few facts. The first is the promised definition of Schur functions.

Definition 2.2.2. [Gessel, 1984] Given any skew shape λ/ρ, define

$$s_{\lambda/\rho}(X) := \sum_{T \in \mathrm{SYT}(\lambda/\rho)} F_{\sigma(T)}(X), \tag{2.2.3}$$

where $s_{\lambda/\rho}$ is termed a Schur function if λ/ρ is a straight shape and a skew Schur function in general.

While it is not obvious from this definition that Schur functions are symmetric or that the Schur functions indexed by straight shapes form a basis for the ring of symmetric functions, what we have gained from this definition is a clear connection to the signatures of tableaux. Further, the quasisymmetric definition of Schur functions is always a finite sum, unlike the more common sum over all semistandard tableaux of a given shape.

The important Lascoux-Leclerc-Thibon (LLT) polynomials and Macdonald polynomials (as introduced in [Lascoux et al., 1997] and [Macdonald, 1988], respectively) may also be expressed using the sum of fundamental quasisymmetric polynomials. We now present these combinatorial definitions, as they will be needed in Section 2.4. The LLT polynomials, denoted $G_{\boldsymbol{\nu}}(X;q)$, were originally described in terms of ribbon tableaux in [Lascoux et al., 1997]. We will instead use the equivalent definition given in [Haglund et al., 2005b, Cor. 5.2.4], which defines $G_{\boldsymbol{\nu}}(X;q)$ by using a k-tuple of skew shapes $\boldsymbol{\nu}$.

Given a k-tuple of skew shapes $\boldsymbol{\nu} = (\nu^{(0)}, \ldots, \nu^{(k-1)})$, we write $|\boldsymbol{\nu}| = n$ if $\sum_{i=0}^{k-1} |\nu^{(i)}| = n$. A standard filling $\mathbf{T} = (T^{(0)}, \ldots, T^{(k-1)})$ of $\boldsymbol{\nu}$ is a bijective filling of the diagram of $\boldsymbol{\nu}$ with entries in $[n]$ such that for all $0 \leq i < k$, each $T^{(i)}$ is strictly increasing up columns and across rows from left to right. Denote the set of standard fillings of $\boldsymbol{\nu}$ as $\mathrm{SYT}(\boldsymbol{\nu})$. Define the *shifted content* of a cell x in $\nu^{(i)}$ as,

$$\tilde{c}(x) = k \cdot c(x) + i, \tag{2.2.4}$$

where $c(x)$ is the content of x in $\nu^{(i)}$. The *shifted content word* of $\mathbf{T}$ is defined as the word retrieved from reading off the values in the cells from lowest shifted content to highest, reading northeast along diagonals of constant shifted content. We may then define $\sigma(\mathbf{T})$ as the signature of the shifted content word of $\mathbf{T}$. For an example, see Figure 2.3.

Figure 2.3: On the left, the shifted contents of a pair of skew diagrams. On the right, a standard filling of the same tuple with shifted content word 453826179 and signature $---+++-+$.

Letting $\mathbf{T}(x)$ denote the entry in cell x, the set of *k-inversions of* $\mathbf{T}$ is

$$\mathrm{Inv}_k(\mathbf{T}) := \{(x, y) \mid k > \tilde{c}(y) - \tilde{c}(x) > 0 \text{ and } \mathbf{T}(x) > \mathbf{T}(y)\}. \tag{2.2.5}$$

The *k-inversion number of* $\mathbf{T}$ is defined as

$$\mathrm{inv}_k(\mathbf{T}) := |\mathrm{Inv}_k(\mathbf{T})|. \tag{2.2.6}$$

If w is the shifted content word of $\mathbf{T} \in \mathrm{SYT}(\boldsymbol{\nu})$, and $\boldsymbol{\nu}$ is a k-tuple, then the *$\boldsymbol{\nu}$-inversion number of* w is defined as

$$\mathrm{inv}_{\boldsymbol{\nu}}(w) := \mathrm{inv}_k(\mathbf{T}). \tag{2.2.7}$$

As an example, let $\mathbf{T}$ be as in Figure 2.3. Denoting cells with their values in $\mathbf{T}$, $\mathrm{Inv}_2(\mathbf{T})$ is comprised of the pairs (5,3), (3,2), and (8,2). Hence, $\mathrm{inv}_2(\mathbf{T})$=$\mathrm{inv}_{\boldsymbol{\nu}}(453826179) = 3$.

Now define the set of LLT polynomials by

$$G_{\boldsymbol{\nu}}(X;q) := \sum_{\mathbf{T}\in \mathrm{SYT}(\boldsymbol{\nu})} q^{\mathrm{inv}_k(\mathbf{T})} F_{\sigma(\mathbf{T})}(X). \tag{2.2.8}$$

Though LLT polynomials are known to be symmetric, with proofs in [Lascoux et al., 1997, Thm. 6.1] and [Haglund et al., 2005a, Theorem 3.3], it is still challenging to expand them in terms of Schur functions. A partial solution to this problem is given in Section 2.4.

We now move on to the definition of the modified Macdonald polynomials $\widetilde{H}_{\mu/\rho}(X;q,t)$. We will use [Haglund et al., 2005a, Theorem 2.2] to give a strictly combinatorial definition. To do this, we will first need to define several functions.

Given any skew shape μ/ρ with each cell represented by a pair (i,j) in Cartesian coordinates, let $\mathrm{TR}(\mu/\rho)$ be the set of tuples of ribbons $\boldsymbol{\nu} = (\nu^{(0)}, \ldots, \nu^{(k-1)})$, such that $\nu^{(i)}$ has a cell with content j if and only if $(i,-j)$ is a cell in μ/ρ. There is then a bijection between standard fillings of shapes in $\mathrm{TR}(\mu/\rho)$ and bijective fillings of μ/ρ given by turning each ribbon into a column of μ/ρ as demonstrated in Figure 2.4.

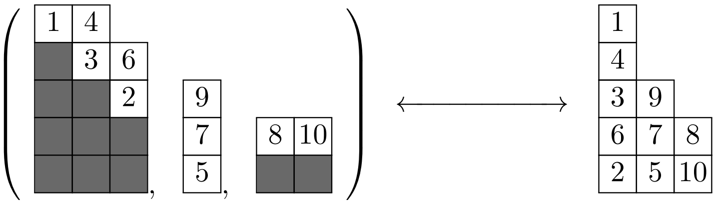

Figure 2.4: An example of the bijection between standard fillings of shapes in $\mathrm{TR}(\mu/\rho)$ and bijective fillings of μ/ρ.

Consider any k-tuple of ribbon shapes $\boldsymbol{\nu} = (\nu^{(0)}, \ldots, \nu^{(k-1)})$ with some cell $x \in \nu^{(i)}$. Define the *arm of* x and the *leg of* x, denoted $a(x)$ and $l(x)$ respectively, by

$$a(x) := |\{\nu^{(j)} : j > i \text{ and there exists some } y \in \nu^{(j)} \text{ such that } c(x) = c(y)\}|. \tag{2.2.9}$$

$$l(x) := |\{y : y \in \nu^{(i)} \text{ and } c(y) < c(x)\}|. \tag{2.2.10}$$

Here, $c(x)$ and $c(y)$ always refer to the content within the skew tableaux containing x and y, respectively. As an example, if x is the cell containing a six in the left diagram of Figure 2.4, then $a(x) = 2$ and $l(x) = 3$.

We define a *descent* of $\boldsymbol{\nu}$ to be any cell x in some $\nu^{(i)}$ that has a cell directly below it and define the descent set of $\boldsymbol{\nu}$ as

$$\mathrm{Des}(\boldsymbol{\nu}) := \{x \in \boldsymbol{\nu} : x \text{ is a descent of } \boldsymbol{\nu}\}. \tag{2.2.11}$$

For example, the descent set of the tuple of ribbon tableaux in Figure 2.4 is the set of cells with values in $\{4, 6, 7, 9\}$. Given a standard filling $\mathbf{T}$ of $\boldsymbol{\nu}$, our final three statistics can then be defined as

$$a(\boldsymbol{\nu}) := \sum_{x \in \mathrm{Des}(\boldsymbol{\nu})} a(x), \tag{2.2.12}$$

$$\mathrm{inv}(\mathbf{T}) := \mathrm{inv}_k(\mathbf{T}) - a(\boldsymbol{\nu}), \tag{2.2.13}$$

$$\mathrm{maj}(\mathbf{T}) := \mathrm{maj}(\boldsymbol{\nu}) := \sum_{x \in \mathrm{Des}(\boldsymbol{\nu})} 1 + l(x). \tag{2.2.14}$$

Using the left diagram in Figure 2.4 as an example again, we have $a(\boldsymbol{\nu}) = 3$, $\mathrm{inv}(\mathbf{T}) = 4-3 = 1$, and $\mathrm{maj}(\mathbf{T}) = 9$. A simple proof that $\mathrm{inv}(\mathbf{T})$ is always nonnegative can be found in [Haglund et al., 2005a, Sec. 2].

We are now able to define the modified Macdonald polynomials and show their relationship with LLT polynomials:

$$\widetilde{H}_{\mu/\rho}(X; q, t) := \sum_{\substack{\boldsymbol{\nu} \in \mathrm{TR}(\mu/\rho) \\ \mathbf{T} \in \mathrm{SYT}(\boldsymbol{\nu})}} q^{\mathrm{inv}(\mathbf{T})} t^{\mathrm{maj}(\mathbf{T})} F_{\sigma(\mathbf{T})} = \sum_{\boldsymbol{\nu} \in \mathrm{TR}(\mu/\rho)} q^{-a(\boldsymbol{\nu})} t^{\mathrm{maj}(\boldsymbol{\nu})} G_{\boldsymbol{\nu}}(X; q). \tag{2.2.15}$$

By using this definition, results about LLT polynomials can be easily translated into results about Macdonald polynomials.

Lastly, we will have use for the following symmetry of modified Macdonald polynomials.

It follows from results in [Macdonald, 1995] (see also [Haglund, 2008, Eq. 2.30]) that

$$\widetilde{H}_{\mu/\rho}(X;q,t) = \widetilde{H}_{\tilde{\mu}/\tilde{\rho}}(X;t,q). \tag{2.2.16}$$

2.2.4 Dual Equivalence Graphs

We now provide the necessary definitions and results from [Assaf, 2013]. We begin by recalling Mark Haiman's dual to the fundamental Knuth equivalences defined in (2.2.2).

Definition 2.2.3. Given a permutation in S_n expressed in one-line notation, define an *elementary dual equivalence* as an involution d_i that interchanges the values $i-1, i$, and $i+1$ as

$$\begin{aligned}
d_i(\dots i-1 \dots i \dots i+1 \dots) &= (\dots i-1 \dots i \dots i+1 \dots),\\
d_i(\dots i+1 \dots i \dots i-1 \dots) &= (\dots i+1 \dots i \dots i-1 \dots),\\
d_i(\dots i \dots i-1 \dots i+1 \dots) &= (\dots i+1 \dots i-1 \dots i \dots),\\
d_i(\dots i-1 \dots i+1 \dots i \dots) &= (\dots i \dots i+1 \dots i-1 \dots).
\end{aligned} \tag{2.2.17}$$

Two words are *dual equivalent* if one may be transformed into the other by successive elementary dual equivalences.

As an example, 21345 is dual equivalent to 51234 because $d_4(d_3(d_2(21345))) = d_4(d_3(31245)) = d_4(41235) = 51234$. Notice that if i is between $i-1$ and $i+1$, then d_i acts as the identity. It follows immediately from (2.2.2) and (2.2.17) that d_i is related to K_i by

$$d_i(w) = (K_i(w^{-1}))^{-1}. \tag{2.2.18}$$

By (2.2.1) and (2.2.18), $Q(w) = P(w^{-1}) = P(K_i(w^{-1})) = Q((K_i(w^{-1}))^{-1}) = Q(d_i(w))$. Thus, Q is constant on dual equivalence classes.

We may also let d_i act on the entries of a tableau T by applying them to the row reading word of T. It is not hard to check that the result is again a tableau of the same shape. The transitivity of this action is described in the following theorem.

Theorem 2.2.4 ([Haiman, 1992, Prop. 2.4])**.** *Two standard Young tableaux on partition shapes are dual equivalent if and only if they have the same shape.*

If we rewrite Theorem 2.2.4 in terms of permutations, it states that dual equivalence classes are precisely the set of permutations w satisfying $Q(w) = T$ for some fixed tableau T.

The same action of d_i on tableaux is defined by using the content reading word instead of the row reading word. To see this, recall that the row reading word of a tableau is Knuth equivalent to the content reading word of the same tableau. Given any $w \in S_n$, it follows from [Haiman, 1992, Lemma 2.3] that for all $1 < i, j < n$,

$$Q(K_j \circ d_i(w)) = Q(K_j(w)) = Q(d_i \circ K_j(w)). \tag{2.2.19}$$

Applying (2.2.18) and (2.2.19), yields

$$\begin{array}{rclclcl} P(d_i \circ K_j(w)) & = & Q((d_i \circ K_j(w))^{-1}) & = & Q(K_i((K_j(w))^{-1})) & & \\ & = & Q(K_i \circ d_j(w^{-1})) & = & Q(d_j \circ K_i(w^{-1})) & & \\ & = & Q(d_j((d_i(w))^{-1})) & = & Q((K_j \circ d_i(w))^{-1}) & = & P(K_j \circ d_i(w)). \end{array}$$

In particular,

$$P(d_i \circ K_j(w)) = P(K_j \circ d_i(w)) = P(d_i(w)). \tag{2.2.20}$$

Thus, the fact that the row reading word and content reading word of a tableau are in the same Knuth class implies that they determine the same action of d_i on a tableau.

By definition, d_i is an involution, and so we define a graph on standard Young tableaux by letting each nontrivial orbit of d_i define an edge colored by i. By Theorem 2.2.4, the graph on SYT(n) with edges labeled by $1 < i < n$ has connected components with vertices in SYT(λ) for each $\lambda \vdash n$. We may further label each vertex with its signature to create a *standard dual equivalence graph* that we will denote $\mathcal{G}_\lambda$ (see Figure 2.5).

Definition 2.2.2 and Theorem 2.2.4 determine the connection between Schur functions and dual equivalence graphs as highlighted in [Assaf, 2013, Cor. 3.10]. Given any standard

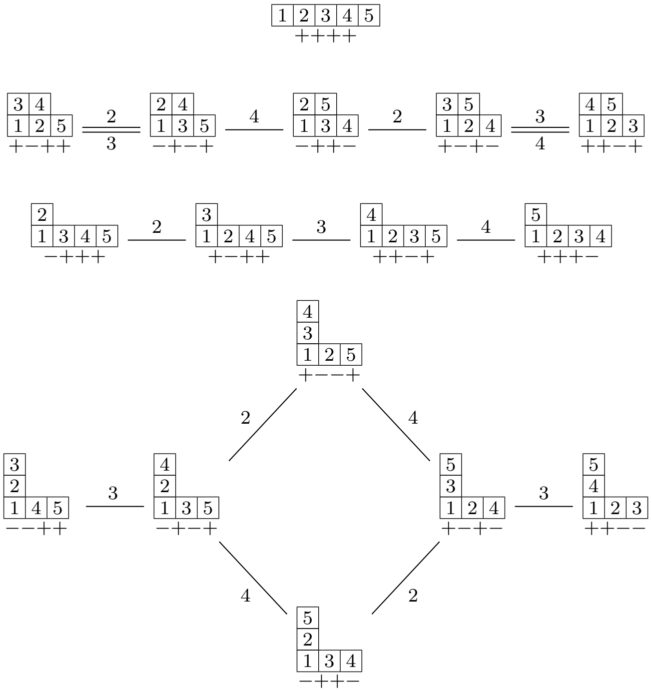

Figure 2.5: The standard dual equivalence graphs on partitions of 5 up to conjugation.

dual equivalence graph $\mathcal{G}_\lambda = (V, \sigma, E)$,

$$\sum_{v \in V} F_{\sigma(v)} = s_\lambda. \tag{2.2.21}$$

Here, $\mathcal{G}_\lambda$ is an example of the following broader class of graphs.

Definition 2.2.5. An *edge colored graph* consists of the following data:

1. *a finite vertex set V,*
2. *a collection E_i of unordered pairs of distinct vertices in V for each $i \in \{m+1, \dots, n-1\}$, where m and n are positive integers.*

A *signed colored graph* is an edge colored graph with the following additional data:

3. *a signature function $\sigma \colon V \to \{\pm 1\}^{N-1}$ for some positive integer $N \geq n$.*

We denote a signed colored graph by $\mathcal{G} = (V, \sigma, E_{m+1} \cup \cdots \cup E_{n-1})$ or simply $\mathcal{G} = (V, \sigma, E)$. If a signed colored graph has $m = 1$, as described above, then it is said to have *type* (n, N) and is termed an (n, N)-signed colored graph.

For our purposes, whenever V is a set of permutations or tableaux, it will be assumed that σ is the signature function defined in Section 2.2.1. To be explicit, we will sometimes refer to this definition of the signature function as *given by inverse descents.*

Signed colored graphs of different types may often be related by restricting some of the data. For example, if $\mathcal{G}$ is an (n, N)-signed colored graph, $M \leq N$, and $m \leq n$, then the (m, M)*-restriction* of $\mathcal{G}$ is the result of excluding E_i for $i \geq m$ and projecting each signature onto its first $M - 1$ coordinates. The (m, M)-component of a vertex v of $\mathcal{G}$ is the connected component containing v in the (m, M)-restriction of $\mathcal{G}$.

In order to describe which signed colored graphs have the same structure as a standard dual equivalence graph, we first need to define isomorphisms.

Definition 2.2.6. A map $\phi \colon \mathcal{G} \to \mathcal{H}$ between edge colored graphs $\mathcal{G} = (V, E_{m+1} \cup \ldots \cup E_{n-1})$ and $\mathcal{H} = (V', E'_{m+1} \cup \ldots \cup E'_{n-1})$ is called a *morphism* if it preserves i-edges. That is, $\{v, w\} \in E_i$ implies $\{\phi(v), \phi(w)\} \in E'_i$ for all $v, w \in V$ and all $m < i < n$.

A map $\phi \colon \mathcal{G} \to \mathcal{H}$ between signed colored graphs $\mathcal{G} = (V, \sigma, E_{m+1} \cup \ldots \cup E_{n-1})$ and $\mathcal{H} = (V', \sigma', E'_{m+1} \cup \ldots \cup E'_{n-1})$ is called a *morphism* if it is a morphism of edge colored graphs that also preserves signatures. That is, $\sigma'(\phi(v)) = \sigma(v)$.

In both cases, a morphism is an *isomorphism* if it admits an inverse morphism.

Though the term morphism is given two different definitions above, the specific definition should be clear from the context.

The next proposition can be thought of as stating that standard dual equivalence graphs are unique up to isomorphism and have trivial automorphism groups.

Proposition 2.2.7 ([Assaf, 2013] Proposition 3.11)**.** *If $\phi \colon \mathcal{G}_\lambda \to \mathcal{G}_\mu$ is an isomorphism of signed colored graphs, then $\lambda = \mu$, and ϕ is the identity morphism.*

Notice that in a standard dual equivalence graph, a vertex v is included in an i-edge if and only if $\sigma(v)_{i-1} = -\sigma(v)_i$, motivating the following definition.

Definition 2.2.8. Let $\mathcal{G} = (V, \sigma, E)$ be a signed colored graph. We say that $w \in V$ *admits an i-neighbor* if $\sigma(w)_{i-1} = -\sigma(w)_i$.

Before moving on to an abstract generalization of the structure inherent in any standard dual equivalence graph, recall that a *complete matching* is a simple graph such that every vertex is contained in exactly one edge.

Definition 2.2.9. A signed colored graph $\mathcal{G} = (V, \sigma, E_{m+1} \cup \ldots \cup E_{n-1})$ is a *dual equivalence graph* if the following axioms hold:

(ax1): For $m < i < n$, each E_i is a complete matching on the vertices of V that admit an i-neighbor.

(ax2): If $\{v, w\} \in E_i$, then $\sigma(v)_i = -\sigma(w)_i, \sigma(v)_{i-1} = -\sigma(w)_{i-1}$, and $\sigma(v)_h = \sigma(w)_h$ for all $h < i-2$ and all $h > i+1$.

(ax3): For $\{v, w\} \in E_i$, if σ_{i-2} is defined, then v or w (or both) admits an $(i-1)$-neighbor, and if σ_{i+1} is defined, then v or w (or both) admits an $(i+1)$-neighbor.

(ax4): For all $m+1 < i < n$, any component of the edge colored graph $(V, E_{i-2} \cup E_{i-1} \cup E_i)$ is isomorphic to a component of the restriction of some $\mathcal{G}_\lambda = (V', \sigma', E')$ to $(V', E'_{i-2} \cup E'_{i-1} \cup E'_i)$, where E_{i-2} is omitted if $i = m+2$ (see Figures 2.6 and 2.7).

(ax5): For all $1 < i, j < n$ such that $|i - j| > 2$, if $\{v, w\} \in E_i$ and $\{w, x\} \in E_j$, then there exists $y \in V$ such that $\{v, y\} \in E_j$ and $\{x, y\} \in E_i$.

(ax6): For all $m < i < n$, any two vertices of a connected component of $(V, \sigma, E_{m+1} \cup \cdots \cup E_i)$ may be connected by some path crossing at most one E_i edge.

A dual equivalence graph that is also an (n, N)-signed colored graph is said to have *type* (n, N) and is termed an (n, N)-*dual equivalence graph.*

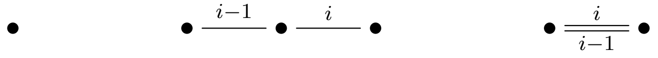

Figure 2.6: Allowable $E_{i-1} \cup E_i$ components of Axiom 4

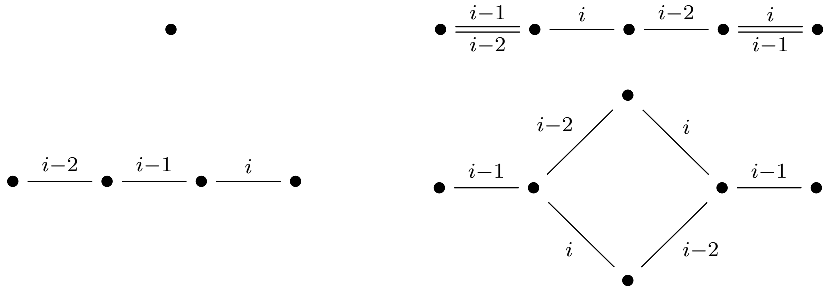

Figure 2.7: The allowable $E_{i-2} \cup E_{i-1} \cup E_i$ components of Axiom 4.

Remark 2.2.10. The following are immediate consequences of Definition 2.2.9:

1. A connected component of a dual equivalence graph is also a dual equivalence graph.

2. For any $m \leq n$ and $M \leq N$, the (m, M) restriction of an (n, N)-dual equivalence graph is an (m, M)-dual equivalence graph.

3. If a signed colored graph has type (n, n) or if $m + 1 < i < n - 1$, then Axiom 3 is implied by Axioms 1, 2, and 4 on components of two consecutive colors. In the presence of Axioms 1 and 2, Axiom 3 can be restated in terms of signatures as follows. For $\{v, w\} \in E_i$, if $\sigma(v)_{i-2} = -\sigma(w)_{i-2}$, then $\sigma(v)_{i-2} = -\sigma(v)_{i-1}$ whenever $i > 2$, and if $\sigma(v)_{i+1} = -\sigma(w)_{i+1}$, then $\sigma(v)_{i+1} = -\sigma(v)_i$ whenever σ_{i+1} is defined. This is the original definition of Axiom 3 used in [Assaf, 2013].

4. It is an instructional exercise to check that if Axioms 1, 2, and 6 are obeyed, then Axiom 4 on 2 consecutive colors implies Axiom 4 on 3 consecutive colors.

5. A signed colored graph $\mathcal{G} = (V, \sigma, E_{m+1} \cup \ldots \cup E_{n-1})$ satisfies Axioms 1, 2, 3, and 4, if and only if for any $m < i < n$, each component of $(V, \sigma, E_{i-2} \cup E_{i-1} \cup E_i)$ is isomorphic to a component of the restriction of some $\mathcal{G}_\lambda = (V', \sigma', E')$ to $(V', \sigma', E'_{i-2} \cup E'_{i-1} \cup E'_i)$, where E'_{i-2} or E'_{i-1} is omitted if $i \leq m + 2$ or $i = m + 1$, respectively. While this fact could be used to shorten the axiomatization, in practice it is often necessary to check Axioms 1, 2, 3, and 4 separately.

The next two theorems link the definition of dual equivalence graphs with that of standard dual equivalence graphs.

Theorem 2.2.11 ([Assaf, 2013, Theorem 3.5])**.** *For any* $\lambda \vdash n$*,* $\mathcal{G}_\lambda$ *is an* (n, n)*-dual equivalence graph.*

The converse is also true.

Theorem 2.2.12 ([Assaf, 2013, Theorem 3.9])**.** *Every connected component of an* (n, n)*-dual equivalence graph is isomorphic to a unique* $\mathcal{G}_\lambda$*.*

The key to proving Theorem 2.2.12 is building an appropriate morphism. Some of the same techniques will prove useful in this paper, and so we lay them out now.

Definition 2.2.13. Fix any partitions $\lambda \subset \mu$ with $|\lambda| = n$ and $|\mu| = N$ and a skew tableau A of shape μ/λ with entries $n+1, \ldots, N$. Define the set of standard Young tableaux *augmented by* A, denoted $\mathrm{ASYT}(\lambda, A)$, as the set of $T \in \mathrm{SYT}(\mu)$ such that the restriction of T to μ/λ is A. Further, define a signed colored graph $\mathcal{G}_{\lambda,A}$ on $\mathrm{ASYT}(\lambda, A)$ with signature given by inverse descents and edges given by the the nontrivial orbits of d_i for $1 < i < n$.

Remark 2.2.14. The graph $\mathcal{G}_{\lambda,A}$ is isomorphic to an (n, N)-component of $\mathcal{G}_\mu$, and every (n, N)-component of $\mathcal{G}_\mu$ is isomorphic to some $\mathcal{G}_{\lambda,A}$, as is clear from the definition of $\mathcal{G}_{\lambda,A}$. By Part 1 of Remark 2.2.10 and Theorem 2.2.11, $\mathcal{G}_{\lambda,A}$ is therefore a dual equivalence graph with (n, n)-restriction isomorphic to $\mathcal{G}_\lambda$. Applying Theorem 2.2.12, every (n, N)-component of an (N, N)-dual equivalence graph is isomorphic to $\mathcal{G}_{\lambda,A}$ for some $\lambda \vdash n$ and some A such that $|A| = N - n$.

With the notion of augmentation it is possible, in some sense, to reverse the process of restriction.

Proposition 2.2.15 ([Assaf, 2013, Lemma 3.13])**.** *Let* $\mathcal{G} = (V, \sigma, E)$ *be a connected* (n, N)*-dual equivalence graph, and let* ϕ *be a morphism from the* (n, n)*-restriction of* $\mathcal{G}$ *to* $\mathcal{G}_\lambda$ *for some partition* λ *of* n*. Then* ϕ *extends to an isomorphism* $\tilde{\phi}\colon \mathcal{G} \to \mathcal{G}_{\lambda,A}$*, where* A *is a skew tableau such that* $|\lambda| + |A| = N$*. Furthermore, the position of the cell containing* $n+1$ *in* A *is unique.*

Because of the uniqueness statement in Proposition 2.2.15, we can unambiguously refer to *the unique extension* of a connected $(n, n+1)$-dual equivalence graph to a connected $(n+1, n+1)$-dual equivalence graph. That is, if an $(n, n+1)$-dual equivalence graph is as in Proposition 2.2.15, then the unique extension is isomorphic to $\mathcal{G}_\mu$, where μ is the union of λ and the cell of A containing $(n+1)$.

Definition 2.2.16. Let $\mathcal{G}$ be a signed colored graph of type $(n+1, n+1)$. Two distinct components of the $(n, n+1)$-restriction of $\mathcal{G}$ that are connected by an n-edge in $\mathcal{G}$ are said to be *neighbors* in $\mathcal{G}$.

While not explicitly stated in [Assaf, 2013], the following is an immediate consequence of the proof of Theorem 2.2.12.

Corollary 2.2.17 ([Assaf, 2013, Theorem 3.14])**.** *Let $\mathcal{G}$ be a connected $(n+1, n+1)$-signed colored graph satisfying Axioms 1–5 whose $(n, n+1)$-restriction is a dual equivalence graph. Let $\mathcal{C}$ be any component of the $(n, n+1)$-restriction of $\mathcal{G}$, and let the unique extension of $\mathcal{C}$ be isomorphic to $\mathcal{G}_\mu$. Then $\mathcal{C} \cup (\bigcup \mathcal{B})$ is isomorphic to the $(n, n+1)$ restriction of $\mathcal{G}_\mu$, where the union is over all $\mathcal{B}$ that are neighbors of $\mathcal{C}$ in $\mathcal{G}$. Furthermore, there exists a morphism $\phi\colon \mathcal{G} \to \mathcal{G}_\mu$.*

2.3 The Structure of Dual Equivalence Graphs

The main results of this section are the classification of graphs satisfying Axiom 1 that admit a morphism onto a dual equivalence graph in Theorems 2.3.13 and Theorem 2.3.14, the improved axiomatization of dual equivalence graphs given in Theorem 2.3.17, and the more specific criterion for satisfying the dual equivalence axioms given in Corollary 2.3.20. In the process, a number of smaller results about the structure of dual equivalence graphs are highlighted.

2.3.1 Symmetries of Dual Equivalence Graphs

We begin by giving notation for a useful signed colored graph. Let $\mathcal{G}_n$ denote the (n, n)-signed colored graph with vertices indexed by the permutations in S_n, signature function given by inverse descents, and i-edges given by the nontrivial orbits of d_i for each $1 < i < n$.

The following lemma is a natural extension of [Haiman, 1992, Lemma 2.3]. It lays out a fundamental relationship between the dual equivalence and Knuth equivalence maps defined in (2.2.2) and (2.2.17).

Lemma 2.3.1. *Given any* $w \in S_n$ *and any* $1 < i < n$, $1 < j < n$, *then* $K_j \circ d_i(w) = d_i \circ K_j(w)$ *and* $\sigma(w) = \sigma(K_j(w))$. *In particular,* K_j *defines an automorphism of* $\mathcal{G}_n$.

Proof: The fact that K_j preserves inverse descent sets follows from its definition in (2.2.2). Thus, $\sigma(w) = \sigma(K_j(w))$. Now we prove that K_j commutes with d_i. Recall that the R-S-K correspondence provides a bijection that sends $w \in S_n$ to a pair of tableaux $(P(w), Q(w))$. Consider the effect that applying d_i and K_j to w has on the pair $(P(w), Q(w))$. By (2.2.19) and (2.2.20),

$$(P(d_i \circ K_j(w)), Q(d_i \circ K_j(w))) = (P(K_j \circ d_i(w)), Q(K_j \circ d_i(w))). \tag{2.3.1}$$

Applying the inverse R-S-K correspondence to S_n yields $K_j \circ d_i(w) = d_i \circ K_j(w)$, as shown in Figure 2.8.

To prove the last part of the lemma, notice that the above argument demonstrates that K_j defines a morphism on $\mathcal{G}_n$. Because K_j is its own inverse, the morphism must be an isomorphism. □

$$\begin{tikzcd}
w \arrow[r, "d_i"] \arrow[drr, leftrightarrow, "R-S-K"] \arrow[dd, "K_j"] & d_i(w) \arrow[drr, leftrightarrow, "R-S-K"] \arrow[dd, "K_j"] & & \\
& & (P(w), Q(w)) \arrow[r, "d_i"] \arrow[dd, "K_j"] & (P(d_i(w)), Q(w)) \arrow[dd, "K_j"] \\
K_j(w) \arrow[r, "d_i"] \arrow[drr, leftrightarrow, "R-S-K"] & K_j \circ d_i(w) \arrow[drr, leftrightarrow, "R-S-K"] & & \\
& & (P(w), Q(K_j(w))) \arrow[r, "d_i"] & (P(d_i(w)), Q(K_j(w)))
\end{tikzcd}$$

Figure 2.8: The commutativity of d_i, K_j, and the R-S-K correspondence acting on a word $w \in S_n$.

Knuth equivalences, jeu de taquin, standard skew tableaux, and dual equivalence graphs are all intimately related, as is demonstrated in the next theorem.

Theorem 2.3.2. *The function* $P\colon S_n \to \mathrm{SYT}(n)$ *induces a surjective morphism from* $\mathcal{G}_n$ *to* $\cup_{\lambda \vdash n}\mathcal{G}_\lambda$*. This morphism restricts to an isomorphism from any given component of* $\mathcal{G}_n$ *to* $\mathcal{G}_\lambda$ *for some* $\lambda \vdash n$*. In particular, if* λ/ρ *is a fixed skew shape and* V *is the set of row reading words of skew tableaux in* $\mathrm{SYT}(\lambda/\rho)$*, then the restriction of* $\mathcal{G}_n$ *to* V *is a dual equivalence graph.*

Proof : We begin by showing that P induces a morphism. As mentioned in Section 2.2.2, each $w \in S_n$ is in a Knuth class with the reading word of some $T \in \mathrm{SYT}(\lambda)$. Thus, there is some sequence of fundamental Knuth equivalences, call it K_w, that takes w to the row reading word of T. Recall from Sections 2.2.1, 2.2.2, and 2.2.4 that $P(w) = P(K_w(w)) = T$, $d_i(P(K_w(w)) = P(d_i \circ K_w(w))$, and $\sigma(P(w)) = \sigma(K_w(w))$. By Lemma 2.3.1, $\sigma(K_w(w)) = \sigma(w)$, so P preserves signatures. Treating $P(w)$ as a vertex in $\mathcal{G}_\lambda = (V, \sigma, E)$, P takes the vertices in some i-edge $\{w, d_i(w)\}$ in $\mathcal{G}_n$, to

$$\begin{aligned}\{P(w), P(d_i(w))\} &= \{P(K_w(w)), P(K_w \circ d_i(w))\}\\ &= \{P(K_w(w)), P(d_i \circ K_w(w))\}\\ &= \{P(K_w(w)), d_i(P(K_w(w)))\} \in E_i,\end{aligned}$$

again by Lemma 2.3.1. Therefore, P induces a morphism from $\mathcal{G}_n$ to $\cup\mathcal{G}_\lambda$. This morphism is surjective because the R-S-K correspondence guarantees every $T \in \mathrm{SYT}(n)$ is the image of some $w \in S_n$ under the action of P.

To prove the second statement in the theorem, it suffices to restrict the domain to any component $\mathcal{C}$ in $\mathcal{G}_n$ and then explicitly create an inverse morphism from $\mathcal{G}_\lambda$ to $\mathcal{C}$. Because Q is constant on dual equivalence classes, Q evaluates to some fixed tableau U on all of $\mathcal{C}$. For any $w \in \mathcal{C}$, $P(w) = T$ if and only if the inverse R-S-K correspondence sends (T, U) to w. The desired inverse morphism is thus given by sending the vertex T to the word corresponding to (T, U) in the R-S-K correspondence. This action takes $P(w)$ to w and $d_i(P(w)) = P(d_i(w))$ to $d_i(w)$, so it must preserve edges. The same analysis as the previous paragraph demonstrates that this action preserves signatures, so we have defined the desired inverse morphism.

For the last statement in the theorem, one can observe from the definition in (2.2.17) that $\mathrm{SYT}(\lambda/\rho)$ is closed under dual equivalence for any fixed shape λ/ρ. Hence, restricting $\mathcal{G}_n$ to the vertex set V is a restriction to a collection of connected components of $\mathcal{G}_n$. Thus, P takes each of these components to a standard dual equivalence graph, completing the proof. □

In light of Theorem 2.3.2, it makes sense to extend the notation $\mathcal{G}_\lambda$ to the skew case $\mathcal{G}_{\lambda/\rho}$. That is, $\mathcal{G}_{\lambda/\rho}$ is the dual equivalence graph with vertices in $\mathrm{SYT}(\lambda/\rho)$, edges given by nontrivial orbits of d_i, and signatures given by inverse descents. As with standard dual equivalence graphs, the actions of d_i and σ are defined via the row reading words of skew tableaux.

Remark 2.3.3. In the definition of $\mathcal{G}_{\lambda/\rho}$, both σ and E_i are defined via row reading words. Thus, $\mathcal{G}_{\lambda/\rho}$ is isomorphic to the signed colored graph induced by sending each vertex to its row reading word. As with standard dual equivalence graphs, both σ and E_i may be equivalently defined via content reading words. Thus, $\mathcal{G}_{\lambda/\rho}$ is also isomorphic to the signed colored graph induced by sending each vertex to its content reading word.

Theorem 2.3.2 leads to a number of simple corollaries. The first is a well-known fact (see [Fulton, 1997], [Sagan, 1991], or [Stanley, 2001]), while the rest help to illuminate the structure of dual equivalence graphs.

Corollary 2.3.4 (Littlewood-Richardson Rule)**.** *The skew Schur functions $s_{\nu/\lambda}$ are Schur positive. Moreover, for all λ/ρ,*

$$s_{\lambda/\rho} = \sum_{\mu \vdash |\lambda|-|\rho|} c^{\lambda}_{\rho,\mu}\, s_\mu,$$

where $c^{\lambda}_{\rho,\mu} = |\{w \in \mathrm{SYam}(\mu) : w \text{ is the row reading word of a skew tableau in } \mathrm{SYT}(\lambda/\rho)\}|$.

Proof : By (2.2.21) and Theorem 2.3.2, we can interpret $c^{\lambda}_{\rho,\mu}$ as the number of connected components isomorphic to $\mathcal{G}_\mu$ in $\mathcal{G}_{\lambda/\rho}$. Applying Theorem 2.3.2, any component that is isomorphic to $\mathcal{G}_\mu$ has exactly one vertex v_μ whose row reading word is mapped to U_μ by

P. Recall that $\mathrm{SYam}(\mu)$ is defined to be the set of words w such that $P(w) = U_\mu$. Thus, $c^\lambda_{\rho,\mu} = |\{w \in \mathrm{SYam}(\mu) : w \text{ is the row reading word of a filling in } \mathrm{SYT}(\lambda/\rho)\}|$. □

The next definition and corollary describe the effect of removing the signs and edges with the lowest labels from a signed colored graph.

Definition 2.3.5. Fix some signed colored graph $\mathcal{G} = (V, \sigma, E_{m+1} \cup \cdots \cup E_{n-1})$ and positive integer h. Let $\mathcal{H} = (V, \sigma', E')$ be the signed colored graph defined by $\sigma'_i = \sigma_{i+h}$ and $E'_i = E_{i+h}$ whenever σ_{i+h} and E_{i+h} are defined in $\mathcal{G}$, respectively. Then $\mathcal{H}$ is termed *the h-upward restriction of $\mathcal{G}$*. If $h = 1$, then $\mathcal{H}$ is termed *the upward restriction of $\mathcal{G}$*. The restriction of $\mathcal{G}$ to $(V, \sigma, E_{m+1} \cup \cdots \cup E_{n-2})$ is termed *the downward restriction of $\mathcal{G}$*.

Corollary 2.3.6. *If $\mathcal{G}$ is a dual equivalence graph, then the h-upward restriction of $\mathcal{G}$ is a dual equivalence graph.*

Proof : Without loss of generality, we may assume that $\mathcal{G} = (V, \sigma, E_{m+1} \cup \ldots \cup E_{n-1})$ is connected. Let $\mathcal{H}$ be the h-upward restriction of $\mathcal{G}$. It follows immediately from Definition 2.3.5 that $\mathcal{H}$ obeys Axioms 1, 2, 3, 4, and 5, so we need only demonstrate Axiom 6. If $h < m$, then Definitions 2.2.9 and 2.3.5 guarantee that Axiom 6 holds for $\mathcal{G}$ if and only if it holds for $\mathcal{H}$. We may then restrict to the case where $\mathcal{G}$ has type (n, n) by simply taking the $(m-1)$-upward restriction. It thus suffices to consider $\mathcal{G} = \mathcal{G}_\lambda$ for some $\lambda \vdash n$, by Theorem 2.2.12. Then $\mathcal{H}$ can be obtained by removing the 1 through h boxes from each vertex in $\mathcal{G}_\lambda$ and subtracting h from the values in all of the remaining boxes, creating some dual equivalence graph $\cup \mathcal{G}_{\lambda/\rho}$, where the union is over all ρ contained in λ such that $\rho \vdash h$. By Theorem 2.3.2, this union is a dual equivalence graph. □

The upward and the downward restrictions of a dual equivalence graph are structurally related, as is made precise in the following definition and corollary.

Definition 2.3.7. Let $\mathcal{G} = (V, \sigma, E)$, and let $\mathcal{H} = (V, \sigma', E')$ be signed colored graphs such that

1. σ and σ' are maps onto $\{\pm 1\}^{N-1}$,
2. $\sigma_i = \sigma'_{N-i}$ for all $1 \leq i < N$,

3. $E_i = E'_{N+1-i}$ whenever E_i or E'_{N+1-i} is defined.

Then $\mathcal{H}$ is termed the *color reversal* of $\mathcal{G}$.

Corollary 2.3.8. *Let $\mathcal{G}$ be an (n,n)-dual equivalence graph, and let $\mathcal{H}$ be the color reversal of $\mathcal{G}$. Then $\mathcal{G}$ is isomorphic to $\mathcal{H}$.*

Proof : It suffices to only consider connected graphs. Applying Theorem 2.2.12 allows us to further reduce to the case where $\mathcal{G}$ is an arbitrary $\mathcal{G}_\lambda$. Let μ/ρ be the skew shape given by rotating λ by 180°, and let $\mathcal{G}_{\mu/\rho}$ be the dual equivalence graph on $\mathrm{SYT}(\mu/\rho)$. The reader can check that $\mathcal{G}_{\mu/\rho}$ is isomorphic to the color reversal of $\mathcal{G}_\lambda$ by simply rotating any filling of λ and then reversing the order of the numbers in the filling as in Figure 2.9. In particular, $\mathcal{G}_{\mu/\rho}$ is connected. To show that $\mathcal{G}_\lambda \cong \mathcal{G}_{\mu/\rho}$, recall that $\mathcal{G}_{\mu/\rho}$ is isomorphic to the signed colored graph induced by sending its vertices to their row reading words with edges given by d_i and signature given by inverse descents. Applying Theorem 2.3.2, $\mathcal{G}_{\mu/\rho} \cong \mathcal{G}_\lambda$ if the row reading word of any vertex v—and thus all vertices—in $\mathcal{G}_{\mu/\rho}$ has $\mathrm{sh}(v) = \lambda$.

Let $T \in \mathrm{SYT}(\mu/\rho)$ be the skew tableau obtained by right justifying all of the rows of U_λ and then top justifying all of the columns, as in the right side of Figure 2.9. This transformation from U_λ to T is achieved by jeu de taquin, which preserves the shape of row reading words, as mentioned in Section 2.2.2. Therefore, the row reading word of T has shape λ, completing the proof. $\square$

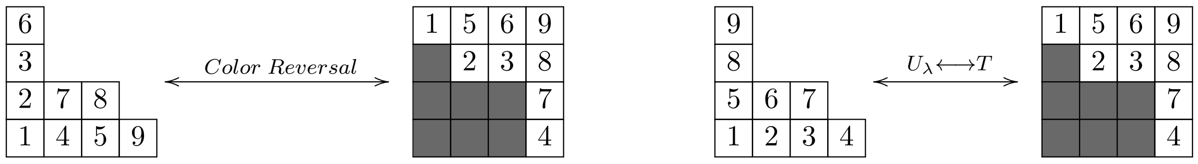

Figure 2.9: At left, a tableau and its color reversal. At right, U_λ for $\lambda = (4,3,1,1)$ and its corresponding skew tableau $T \in \mathrm{SYT}(\mu/\rho)$.

Remark 2.3.9. Because color reversal acts as an isomorphism between dual equivalence graphs, it induces an isomorphism between standard dual equivalence graphs. While not proven here, it can be shown that this induced isomorphism acts on tableau via the Schützenburger evacuation function as defined in [Stanley, 2012b, Sec. 3.20].

2.3.2 Morphisms

In this section, we set out to describe various properties of morphisms between signed colored graphs. We restrict our attention to graphs satisfying at least Axiom 1. For this reason, we may abuse notation and treat each E_i as a function defined on those vertices admitting an i-edge. That is, we write $E_i(v) = w$ to mean that v is contained in an i-edge with $w \neq v$.

Lemma 2.3.10. *Let $\mathcal{G} = (V, \sigma, E)$ and $\mathcal{H} = (V', \sigma', E')$ be nonempty signed colored graphs satisfying Axiom 1. Also, suppose that there exists a morphism $\phi\colon \mathcal{G} \to \mathcal{H}$. Then the following hold.*

1. *$\phi(E_i(v)) = w$ if and only if $E'_i(\phi(v)) = w$.*
2. *If $\mathcal{H}$ is connected, then ϕ is surjective.*
3. *If ϕ is surjective and either $\mathcal{G}$ or $\mathcal{H}$ obey Axiom 2 or 3, they both do.*
4. *If ϕ is surjective and $\mathcal{G}$ obeys Axiom 4, 5, or 6, then $\mathcal{H}$ does as well.*
5. *If $\mathcal{H}$ is connected, then ϕ is an m to 1 map on vertices for some positive integer m.*
6. *If ϕ is a bijection from V to V', then ϕ is an isomorphism from $\mathcal{G}$ to $\mathcal{H}$.*

Proof: We begin with Part 1 and continue in order. Morphisms preserve signatures, so if either v or $\phi(v)$ does not admit an i-edge, then Axiom 1 implies that neither is contained in an i-edge. Now suppose that $\{u, v\}$ and $\{\phi(v), w\}$ are i-edges in $\mathcal{G}$ and $\mathcal{H}$, respectively. By Axiom 1, these are the only i-edges containing v and $\phi(v)$, so by the definition of morphisms, $\phi(u) = w$. That is, $\phi(E_i(v)) = E'_i(\phi(v))$. Thus, Part 1 holds.

For Part 2, we apply Part 1. Choose any $v \in \mathcal{G}$. All of the vertices connected to $\phi(v)$ by an edge have a preimage in $\mathcal{G}$ by Part 1. Since $\mathcal{H}$ is connected, we may then induct to reach any vertex in $\mathcal{H}$ by repeating this process.

Axioms 2 and 3 are concerned with the relationship between signatures and edges. By the definition of a morphism, ϕ preserves signatures, and by Part 1, $\phi(E_i(v)) = w$ if and only if $E'_i(\phi(v)) = w$ for each i where E'_i is defined. Hence, if either graph obeys Axiom 2 or 3, they both do.

Proving Part 4 is a matter of repeatedly applying Part 1 to show that $\mathcal{H}$ inherits the desired properties from $\mathcal{G}$. As we will not need these properties for later results, we leave the details to the reader.

For Part 5, let $\phi^{-1}(x)$ be the set of vertices in $\mathcal{G}$ mapped to x by ϕ. If $E'_i(x) = y$, then Part 1 and Axiom 1 imply that the vertices in $\phi^{-1}(x)$ share an i-edge with distinct vertices in $\phi^{-1}(y)$ and vice versa. Hence, $|\phi^{-1}(x)| = |\phi^{-1}(y)|$. Because $\mathcal{H}$ is connected, we may induct from x to any vertex in $\mathcal{H}$ to show that the fiber over every vertex in $\mathcal{H}$ has the same cardinality. That is, ϕ is an m to one map on vertices.

Part 6 requires the existence of an inverse morphism, which follows from Part 1 and bijectivity. □

Remark 2.3.11. Let $\phi\colon \mathcal{G} \to \mathcal{H}$ be as in Lemma 2.3.10 and suppose that $\mathcal{H}$ is connected. Then in the language of algebraic graph theory, Part 1 implies that ϕ is a covering map. Here we use the definition of a covering map on graphs - a surjective map that sends vertices to vertices and induces a bijection between edges containing v and edges containing $\phi(v)$ for each vertex v in $\mathcal{G}$ - though the topological definition of a covering map can be made to apply as well. In this context, Parts 5 and 6 of Lemma 2.3.10 are well known properties of covering maps. For more details, see [Godsil and Royle, 2004, Sec. 6.8].

Corollary 2.3.12. *If ϕ is any morphism from a connected dual equivalence graph to a connected signed colored graph satisfying Axiom 1, then ϕ is an isomorphism.*

Proof: By Part 6 of Lemma 2.3.10, we need only show that ϕ is bijective. By Parts 2 and 5 of Lemma 2.3.10, ϕ is a surjective m to one map, so it suffices to show that $m = 1$. Taking restrictions if necessary, we may then assume that ϕ is a map between connected signed colored graphs of type (n, n). Applying Theorem 2.2.12 allows us to assume that the domain of ϕ is some $\mathcal{G}_\lambda$. Now notice that in $\mathcal{G}_\lambda$ there is only one vertex with the signature of U_λ, implying that $m = 1$. Thus, ϕ is a bijection. □

In light of (2.2.21) and Part 5 of Lemma 2.3.10, it is natural to look for signed colored graphs that admit a morphism onto a union of standard dual equivalence graphs. The next theorem describes a class of signed colored graphs that admit such a morphism and gives a formula for their Schur expansion.

Theorem 2.3.13. *Let $\mathcal{G} = (V, \sigma, E)$ be an (n, n)-signed colored graph satisfying the following properties:*

1. *$\mathcal{G}$ obeys Axiom 1,*
2. *The vertices in V are indexed by a subset of S_n,*
3. *The signature function σ is given by inverse descent sets of permutations,*
4. *$E_i(v)$ is Knuth equivalent to $d_i(v)$ for all $1 < i < n$ and all $v \in V$ admitting an i-edge.*

Then $P\colon V \to \mathrm{SYT}(n)$ induces a morphism $\phi\colon \mathcal{G} \to \bigcup_{\lambda \vdash n} \mathcal{G}_\lambda$. Furthermore,

$$\sum_{v \in V} F_{\sigma(v)} = \sum_{\lambda \vdash n} |\{V \cap \mathrm{SYam}(\lambda)\}| \cdot s_\lambda.$$

Proof : To show that P induces a morphism on $\mathcal{G}$, first notice that Theorem 2.3.2 implies that $\sigma(P(w)) = \sigma(w)$ for all $w \in V$, so ϕ preserves signatures. Now choose any $w \in V$ admitting an i-edge, and let K_w be a composition of fundamental Knuth equivalences such that $E_i(w) = K_w \circ d_i(w)$. Because P is constant on Knuth classes, $d_i(P(w)) = d_i(P(K_w(w)))$. By Lemma 2.3.1 and Theorem 2.3.2, d_i commutes with K_w and P. Thus,

$$d_i(P(K_w(w))) = P(d_i \circ K_w(w)) = P(K_w \circ d_i(w)) = P(E_i(w)).$$

Hence, ϕ preserves edges and signatures, satisfying the definition of a morphism.

To verify the second part of the theorem, we may restrict the domain to $\phi^{-1}(\mathcal{G}_\lambda)$. Applying (2.2.21), we need only show that $|\{V \cap \mathrm{SYam}(\lambda)\}|$ is equal to the value of the index m of this restriction, as stated in Part 5 of Lemma 2.3.10. By definition, $\mathrm{SYam}(\lambda)$ is the set of words mapped to U_λ by P, so $m = |\mathrm{SYam}(\lambda) \cap V|$, completing the proof. □

Theorem 2.3.13 is close to a complete description of morphisms from (n, n)-signed colored graphs satisfying Axiom 1 onto dual equivalence graphs. This is made precise in the following theorem.

Theorem 2.3.14. *Let $\mathcal{G}$ be an nonempty (n,n)-signed colored graph satisfying Axiom 1 and admitting a morphism $\phi\colon \mathcal{G} \to \mathcal{G}_\lambda$. Then $\mathcal{G}$ is isomorphic to the (n,n)-restriction of some $\mathcal{H} = (V', \sigma', E')$ satisfying the following properties:*

1. *$\mathcal{H}$ obeys Axiom 1,*
2. *The vertices in V' are indexed by a subset of S_N for some $N \geq n$,*
3. *The signature function $\sigma\colon V \to \{\pm 1\}^{N-1}$ is given by inverse descent sets of permutations,*
4. *$E_i(v')$ is Knuth equivalent to $d_i(v')$ for all $1 < i < n$ and all $v' \in V'$ admitting an i-edge.*

Proof: By Parts 2 and 5 of Lemma 2.3.10, ϕ is a surjective m to 1 map. Choose any $\mu \supset \lambda$ such that $|\mathrm{SYT}(\mu)| \geq m$. Here, $N = |\mu|$. Let A be any skew tableau of shape μ/λ with the values $n+1, \ldots, N$. Let $\tilde{\mathcal{G}} = (V, \sigma, E)$ be the (n, N)-signed colored graph with the same vertex and edge set as $\mathcal{G}$ and with $\sigma(v)$ defined for all $v \in V$ as the signature of $\phi(v)$ augmented by A. It is clear that $\mathcal{G}$ is the (n,n)-restriction of $\tilde{\mathcal{G}}$, so we will construct $\mathcal{H}$ to be isomorphic to $\tilde{\mathcal{G}}$.

We will define $\mathcal{H}$ by finding an appropriate relabeling of the vertices of $\tilde{\mathcal{G}}$. By the construction of $\tilde{\mathcal{G}}$, we may extend ϕ to a morphism $\tilde{\phi} : \tilde{\mathcal{G}} \to \mathcal{G}_{\lambda,A}$. Because $|\mathrm{SYT}(\mu)|$ is greater than or equal to the index m of $\tilde{\phi}$, there exists an injective map on the vertices in V, sending each vertex v to $(\tilde{\phi}(v), T_v)$ for some tableau T_v of shape μ. Since $\tilde{\phi}(v)$ and T_v are the same shape, the inverse R-S-K correspondence takes each pair to a unique permutation in S_N. Let f be the injective function taking vertices of V to these permutations in S_N. We claim that if $\mathcal{H} = (V', \sigma', E')$ is the signed colored graph induced by letting f relabel the vertices of $\tilde{\mathcal{G}}$, then $\mathcal{H}$ satisfies all of the desired properties in the statement of the theorem.

First, $\tilde{\mathcal{G}} \cong \mathcal{H}$, because $\tilde{\mathcal{G}}$ and $\mathcal{H}$ only differ by the labeling of their vertex sets. In particular, $\mathcal{H}$ inherits Axiom 1 from $\tilde{\mathcal{G}}$. Choose any $v \in V$ and $v' \in V'$ such that $f(v) = v'$. To see that σ' agrees with the signature given by inverse descents, notice that $\tilde{\phi}(v) = P(v')$, by the construction of f. By the definition of a morphism, $\tilde{\phi}$ preserves signature, and by Theorem 2.3.2, the signature of $P(v')$ is equal to the signature of v' given by inverse descents. Thus σ' is given by inverse descents. Lastly, we show that if $\{v, w\}$ is an i-edge in $\tilde{\mathcal{G}}$, then $d_i(v')$ is in the Knuth class of $E'_i(v') = f(w)$. Applying Part 1 of Lemma 2.3.10 and Theorem 2.3.2 gives

$$P(f(w)) = \tilde{\phi}(w) = d_i(\tilde{\phi}(v')) = d_i(P(v')) = P(d_i(v')),$$

and so $f(w)$ is in the Knuth class of $d_i(v')$. □

Remark 2.3.15. Theorems 2.3.13 and 2.3.14 can both be extended to statements about (n, N)-signed colored graphs with morphisms to augmented dual equivalence graphs.

2.3.3 Local Conditions for Axiom 6

Out of the six dual equivalence axioms, Axiom 6 is the only one that cannot be checked by testing local criteria. In this section we will show that an equivalent axiomatization is given by strengthening Axiom 4 and omitting Axiom 6.

Definition 2.3.16. A signed colored graph $\mathcal{G} = (V, \sigma, E_{m+1} \cup \ldots \cup E_{n-1})$ is said to obey Axiom 4^+ if for all $m + 1 < i < n$, any component of the edge colored graph $(V, E_{i-3} \cup E_{i-2} \cup E_{i-1} \cup E_i)$ is isomorphic to a component of the restriction of some $\mathcal{G}_\lambda = (V', \sigma', E')$ to $(V', E'_{i-3} \cup E'_{i-2} \cup E'_{i-1} \cup E'_i)$, where E_{i-3} or E_{i-2} is omitted if $i \leq m + 3$ or $i \leq m + 2$, respectively.

We now state the main result of this section. The proof is postponed until after a necessary lemma.

Theorem 2.3.17. *A signed colored graph satisfies Axioms 1, 2, 3,* 4^+*, and 5 if and only if it is a dual equivalence graph.*

Remark 2.3.18.

1. We may readily classify the set of edge colored graphs described in Definition 2.3.16, i.e., the set of edge colored graphs that arise as components of the restriction of some

$\mathcal{G}_\lambda = (V, \sigma, E)$ to $(V, E_{i-3} \cup E_{i-2} \cup E_{i-1} \cup E_i)$ for all choices $1 < i < |\lambda|$. Applying Corollary 2.3.6, each such edge colored graph is the result of choosing an appropriate $\lambda \vdash 6$, restricting $\mathcal{G}_\lambda = (V', \sigma', E')$ to the edge colored graph (V', E'), and adding a fixed nonnegative integer to each edge label.

2. A signed colored graph $\mathcal{G} = (V, \sigma, E_{m+1} \cup \ldots \cup E_{n-1})$ satisfies Axioms 1, 2, 3, and 4^+ if and only if, for any $m < i < n$, each component of $(V, \sigma, E_{i-3} \cup E_{i-2} \cup E_{i-1} \cup E_i)$ is isomorphic to a component of the restriction of some $\mathcal{G}_\lambda = (V', \sigma', E')$ to $(V', \sigma', E_{i-3} \cup E_{i-2} \cup E_{i-1} \cup E_i)$, where E_{i-3}, E_{i-2}, or E_{i-1} is omitted if $i \leq m+3, i \leq m+2$ or $i = m+1$, respectively.

Lemma 2.3.19. *Let $\mathcal{G}$ be a connected $(n+1, n+1)$-signed colored graph satisfying Axioms 1, 2, 3, 4, and 5 whose downward restriction and upward restriction are both dual equivalence graphs. Let $\mathcal{C}$ be any component of the downward restriction of $\mathcal{G}$. Additionally, suppose that for every pair of distinct $(n, n+1)$-components $\mathcal{A}$ and $\mathcal{B}$ that are neighbors of $\mathcal{C}$ in $\mathcal{G}$, there exists a component of the upward restriction of $\mathcal{G}$ whose vertex set intersects $\mathcal{A}$ nontrivially and intersects $\mathcal{B}$ nontrivially. Then $\mathcal{G}$ is a dual equivalence graph.*

Proof: Consider any such $\mathcal{A}$ and $\mathcal{B}$ whose vertices are nontrivially intersected by some component $\mathcal{D}$ of the upward restriction of $\mathcal{G}$. Because $\mathcal{D}$ obeys Axiom 6, there exists a path in $\mathcal{G}$ from a vertex in $\mathcal{A}$ to a vertex in $\mathcal{B}$ crossing a single n-edge, which would be labeled as an $(n-1)$-edge in $\mathcal{D}$. Hence, $\mathcal{A}$ and $\mathcal{B}$ are neighbors in $\mathcal{G}$. Because $\mathcal{A}$ and $\mathcal{B}$ were chosen arbitrarily, all neighbors of $\mathcal{C}$ are pairwise neighbors of each other in $\mathcal{G}$.

Next we show that every pair of vertices in $\mathcal{G}$ can be connected by a path containing at most one n-edge. By Corollary 2.2.17, every $(n, n+1)$-component of $\mathcal{G}$ has the same number of neighbors in $\mathcal{G}$. Thus, if $\mathcal{B}$ is any neighbor of $\mathcal{C}$ in $\mathcal{G}$, each neighbor of $\mathcal{B}$ is either $\mathcal{C}$ or a neighbor of $\mathcal{C}$ in $\mathcal{G}$. That is, $\mathcal{C} \cup (\bigcup \mathcal{B})$ contains all of the vertices of $\mathcal{G}$, where the union is over all $\mathcal{B}$ that neighbor $\mathcal{C}$ in $\mathcal{G}$. Thus, all $(n, n+1)$-components of $\mathcal{G}$ are pairwise neighbors. In particular, any two vertices in $\mathcal{G}$ can be connected by a path crossing at most one n-edge.

This property of paths in $\mathcal{G}$, along with the hypothesis that the $(n, n+1)$-restriction of $\mathcal{G}$ is a dual equivalence graph, guarantees that $\mathcal{G}$ satisfies Axiom 6. By assumption, $\mathcal{G}$ also satisfies Axioms 1, 2, 3, 4, and 5, so $\mathcal{G}$ is a dual equivalence graph. $\square$

Proof of Theorem 2.3.17:

We begin with the forward implication by showing that any signed colored graph satisfying Axioms $1, 2, 3, 4^+$, and 5 must also satisfy Axiom 6. Because Axiom 6 is concerned with edge sets, it suffices to only consider signed colored graphs of type (n, N), as in the proof of Corollary 2.3.6. We now proceed by induction on n. Axiom 4^+, when considered with Axioms 1, 2, and 3, implies the theorem holds for (n, N)-signed colored graphs with $n \leq 6$. Now suppose that the result holds for all (n, N)-signed colored graphs, and consider any $(n+1, n+1)$-signed colored graph obeying axioms 1, 2, 3, 4^+, and 5. Call this graph $\mathcal{G}$.

Choose any component $\mathcal{C}$ of the downward restriction of $\mathcal{G}$. We will show that $\mathcal{G}$ and $\mathcal{C}$ satisfy the hypotheses of Lemma 2.3.19. By assumption, $\mathcal{G}$ satisfies Axioms 1-5, and by induction, the upward and downward restrictions of $\mathcal{G}$ are dual equivalence graphs. It then remains to be shown that if $\mathcal{A}$ and $\mathcal{B}$ are any distinct neighbors of $\mathcal{C}$ in $\mathcal{G}$, then there exists a component $\mathcal{D}$ of the upward restriction of $\mathcal{G}$ that intersects the vertices of $\mathcal{A}$ and $\mathcal{B}$ nontrivially.

Next, we label the vertices in $\mathcal{G}$ by tableaux. By Theorem 2.2.12, $\mathcal{C}$ is isomorphic to some $\mathcal{G}_\lambda$, and by Proposition 2.2.15 there is some $\mathcal{G}_\mu$ isomorphic to the unique extension of $\mathcal{C}$. Furthermore, Corollary 2.2.17 guarantees the existence of a morphism from $\mathcal{G}$ to $\mathcal{G}_\mu$. Label the vertices of $\mathcal{G}$, as well as the vertices of the downward restrictions of $\mathcal{G}$, by the set of tableaux in $\mathrm{SYT}(\mu)$ as given by the image of this morphism. By Corollary 2.3.12, this morphism restricts to an isomorphism from any given component of the downward restriction of $\mathcal{G}$ to some component of the downward restriction of $\mathcal{G}_\mu$.

We may now associate the isomorphism types of $\mathcal{A}, \mathcal{B}$, and $\mathcal{C}$ to the position of the cell containing $n+1$ in fillings of μ, as in Figure 2.10. It follows from Theorem 2.2.4 that the number of $(n, n+1)$-components in $\mathcal{G}_\mu$ is equal to the number of Northeast corners of μ, with the isomorphism type of each $(n, n+1)$-component determined by which Northeast corner is filled by $n+1$. By Corollary 2.2.17, $\mathcal{A}$ can be described as the unique neighbor of $\mathcal{C}$ whose tableaux have $n+1$ in a particular Northeast corner of μ. Since $\mathcal{A}, \mathcal{B}$, and $\mathcal{C}$ have distinct isomorphism types, they each have different Northeast corners c_1, c_2, and c_3, respectively, filled by $n+1$.

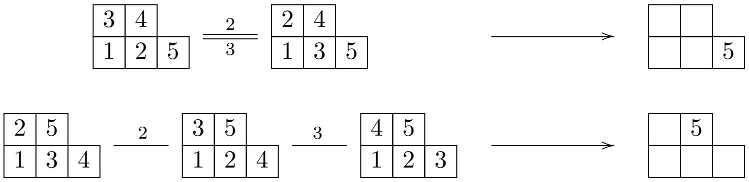

Figure 2.10: The $(4,4)$-components of $\mathcal{G}_{(3,2)}$, at left, are represented via the cells containing 5, at right. Signatures are omitted.

To apply Lemma 2.3.19, it suffices to show that there is a tableau $T \in \mathrm{SYT}(\mu)$ from which $n+1$ may be moved into each of these three Northeast corners while staying in the vertex set of a component of the upper restriction, i.e., without crossing a 2-edge. It is possible to use jeu de taquin to describe which cells $n+1$ may be moved to without crossing a 2-edge. Let $S \in \mathrm{SYT}(\mu)$ be some vertex of $\mathcal{G}_\mu$. As in the proof of Corollary 2.3.6, the set of vertices in the same component of the upward restriction of $\mathcal{G}_\mu$ as S is described by the component of S after omitting the 1-cell (for notational simplicity, we choose not to relabel the boxes). We may then perform jeu de taquin to retrieve some unique tableau $J(S)$ such that $\mathrm{sh}(J(S))$ is a straight shape. Any sequence of jeu de taquin slides acts on the row reading words of a filling by a sequence of fundamental Knuth equivalences, and so commutes with the action of d_i by Lemma 2.3.1. Theorem 2.2.4 guarantees that $n+1$ may then be moved to any Northeast corner of $\mathrm{sh}(J(S))$. In particular, if c is some Northeast corner of μ, x is the value assigned to c by S, and x is also in a Northeast corner of $J(S)$, then there is a path in $\mathcal{G}_\mu$ with no 2-edges that connects S to a tableau with the value $n+1$ in c. Therefore, to apply Lemma 2.3.19, we need only find $T \in \mathrm{SYT}(\mu)$ such that the values in c_1, c_2, and c_3 are in Northeast corners of $J(T)$.

Having reduced the problem to a matter of jeu de taquin, we now consider two cases. First, suppose μ has at least four Northeast corners. Given c_1, c_2, and c_3, we are free to lose some fourth corner c_4. By filling all boxes weakly southwest of c_4 with as low of values as possible, the fourth corner will always be moved, leaving the three given corners unchanged, as in the left half of Figure 2.11.

Figure 2.11: At left, a filling chosen to lose the corner containing 6 to jeu de taquin. At right, a filling of a nonstaircase shape that does not lose any Northeast corners to jeu de taquin.

As a second case, suppose μ is not a staircase shape, i.e. μ is not of shape $(\mu_1, \mu_1 - 1, \ldots, 2, 1)$. Then we claim that there exists some $T \in \mathrm{SYT}(\mu)$ such that every value in a Northeast corner of T is in a Northeast corner of $J(T)$. Notice that μ must contain some rectangle whose Northeast corner is not a Northeast corner of μ but is either the most east cell in its row or the most north cell in its column. It is easy to check that any filling of μ such that all values inside of the rectangle are less than all values outside of the rectangle will suffice. See the right half of Figure 2.11 for an example.

We have shown that $\mathcal{G}$ is a dual equivalence graph if μ is not a staircases or has at least four Northeast corners. There are only three staircase shapes that have strictly less than four Northeast corners, all of which have size less than or equal to 6. All shapes with size less than or equal to 6 are contained in the base case, so $\mathcal{G}$ is a dual equivalence graph.

Assuming our inductive hypotheses, we have demonstrated that any $(n+1, n+1)$-signed colored graph that satisfies Axioms 1, 2, 3, 4^+, and 5 also satisfies Axiom 6. An $(n+1, N)$-signed colored graph obeys Axiom 6 if and only if its $(n+1, n+1)$-restriction obeys Axiom 6, completing the inductive step. Hence, any signed colored graph obeying Axioms 1, 2, 3, 4^+, and 5 is a dual equivalence graph.

The reverse implication follows more quickly. If $\mathcal{G} = (V, \sigma, E)$ is a dual equivalence graph, then we need only show that $\mathcal{G}$ obeys Axiom 4^+. As in Part 1 of Remark 2.3.18, we may assume that $\mathcal{G}$ has type (n, n). By Theorem 2.2.12, $\mathcal{G}$ is isomorphic to some $\mathcal{G}_\lambda$, so axiom 4^+ follows immediately. Hence, every dual equivalence graph satisfies Axioms 1, 2, 3, 4^+, and 5, completing the proof. $\square$

We have actually proven a slightly stronger—and sometimes easier to check—condition. In the previous proof, Axiom 4^+ was only invoked when considering the staircase with six cells, while the usual Axiom 4 could have been used for the staircase with three cells. If we assume Axioms 1, 2, 3, 4, and 5, then the staircase with six cells can only break Axiom 6 in a specific set of graphs. The (5,6)-components of $\mathcal{G}_{(3,2,1)}$ have three distinct isomorphism types. By Corollary 2.2.17, each of these components is connected to two neighboring components. If Axiom 6 is not satisfied, then Lemma 2.3.19 implies these two components cannot be neighbors of each other. Rather, the (5,6)-components must form a loop. The smallest example is presented in Figure 2.12. In the figure, there are two copies of each (5,6)-component. For each positive integer $m \geq 2$, there is then a unique graph with m isomorphic copies of each (5,6)-components, as in Figure 2.13. We may then omit signatures and relabel edges with elements in $\{i-3, i-2, i-1, i\}$, as mentioned in Remark 2.3.18 for a full description of how Axiom 4^+ can break in the presence of Axioms 1, 2, 3, 4, and 5.

Let $\mathcal{F}$ be the set of edge colored graphs with edge sets $E_{i-3} \cup \ldots \cup E_i$ that satisfy Axioms 1, 2, 3, 4, and 5 but not Axiom 6. The following corollary reformulates Theorem 2.3.17 in terms of $\mathcal{F}$.

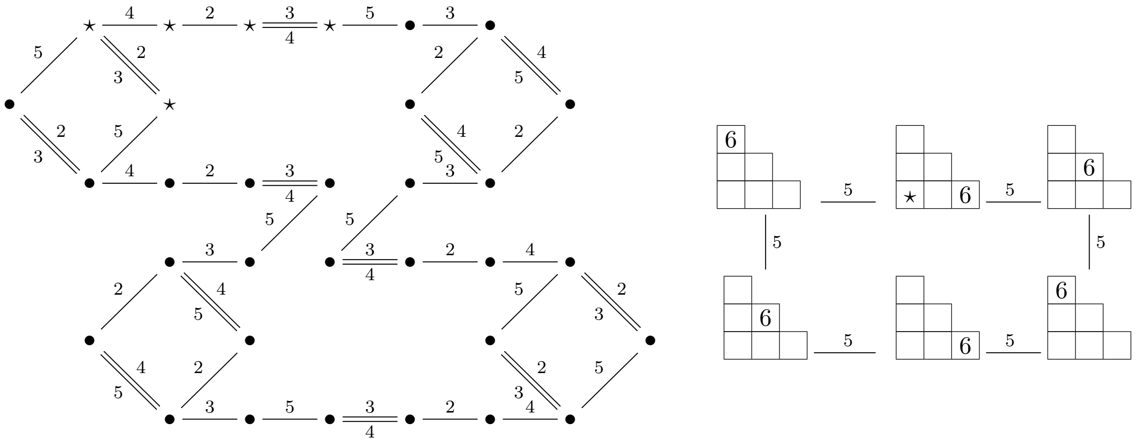

Figure 2.12: Two representations of the smallest graph in $\mathcal{F}$ with edge labels in $\{2, 3, 4, 5\}$. Starred vertices on the left correspond to the starred shape on the right.

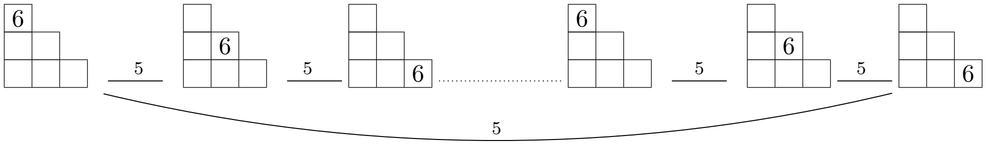

Figure 2.13: A generic element of $\mathcal{F}$ in the case where edge labels are in $\{2,3,4,5\}$.

Corollary 2.3.20. *Let $\mathcal{G} = (V, \sigma, E_{m+1} \cup \ldots \cup E_{n-1})$ be a signed colored graph satisfying Axioms 1, 2, 3, 4, and 5. Then $\mathcal{G}$ is a dual equivalence graph if and only if for all $m+4 < i < n$, the restriction of $\mathcal{G}$ to the edge colored graph $(V, E_{i-3} \cup E_{i-2} \cup E_{i-1} \cup E_i)$ has no components isomorphic to an element of $\mathcal{F}$.*

Remark 2.3.21. For any edge colored graph in $\mathcal{F}$, every vertex shares an edge with at least two other vertices. We may then give yet another characterization of dual equivalence graphs. Let $\mathcal{G} = (V, \sigma, E_{m+1} \cup \ldots \cup E_{n-1})$ be a signed colored graph obeying Axioms 1, 2, 3, 4, and 5. Choose $\mathcal{C}$ to be any component of the restriction of $\mathcal{G}$ to the edge colored graph $(V, E_{i-3} \cup E_{i-2} \cup E_{i-1} \cup E_i)$ such that $m+3 < i < n$ and the vertices of $\mathcal{C}$ all have at least two adjacent vertices in $\mathcal{C}$. Then $\mathcal{G}$ is a dual equivalence graph if and only if $\mathcal{C}$ is not in $\mathcal{F}$ for any choice of $\mathcal{C}$. This characterization of dual equivalence graphs is used in the computer verification of Theorem 2.4.11.

2.4 LLT and Macdonald Polynomials

In this section we show that a family of LLT polynomials can be generated as a sum over vertices of a dual equivalence graph, and we provide a simple Schur expansion for polynomials in this family. Along the way, we recall Assaf's set of signed colored graphs associated to LLT polynomials and classify which ones are dual equivalence graphs.

2.4.1 The Schur Expansion of $G_{\boldsymbol{\nu}}(X;q)$ when diam($\boldsymbol{\nu}$) ≤ 3

Recall the notation and vocabulary introduced in Section 2.2.3 on LLT polynomials and Macdonald polynomials. To state the main theorem of this section, we will also need the following definition.

Definition 2.4.1. Given a k-tuple of skew shapes $\boldsymbol{\nu}$, let $S(\boldsymbol{\nu})$ be the set of distinct shifted contents of the cells in $\boldsymbol{\nu}$. Define the *diameter* of $\boldsymbol{\nu}$, denoted $\operatorname{diam}(\boldsymbol{\nu})$, as

$$\operatorname{diam}(\boldsymbol{\nu}) := \max\{|R| : R \subset S(\boldsymbol{\nu}) \text{ and } |x-y| \leq k \text{ for all } x, y \in R\}.$$

See Figure 2.14 for an example.

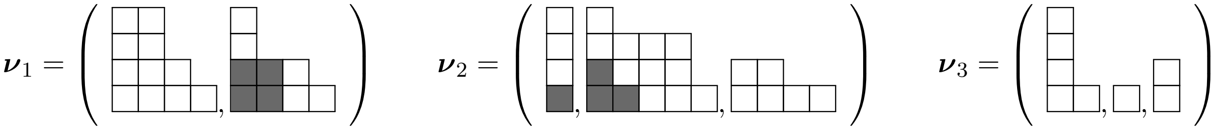

Figure 2.14: The tuples $\boldsymbol{\nu}_1$ and $\boldsymbol{\nu}_2$ have diameter 3. The tuple $\boldsymbol{\nu}_3$ has diameter 4.

Remark 2.4.2. If $\boldsymbol{\nu} = (\nu^{(0)}, \ldots, \nu^{(k-1)})$ is a tuple of skew shapes, it follows from Definition 2.4.1 that $\operatorname{diam}(\boldsymbol{\nu}) \leq k+1$. There are many examples where this bound is sharp. For instance, if $\boldsymbol{\nu}$ is a k-tuple of straight shapes where $\nu^{(0)}$ has as least two columns and $|\nu^{(i)}| \geq 1$ for all $1 \leq i \leq k-1$, then $\operatorname{diam}(\boldsymbol{\nu}) = k+1$. In Figure 2.14, $\boldsymbol{\nu}_3$ is an example of such a tuple.

To further ease the presentation of the following results, we define two sets,

$$\mathbf{T}(\boldsymbol{\nu}, \lambda) := \{\mathbf{T} \in \operatorname{SYT}(\boldsymbol{\nu}) : \text{the shifted content word of } \mathbf{T} \text{ is in } \operatorname{SYam}(\lambda)\}, \quad (2.4.1)$$

$$T(\mu/\rho, \lambda) := \{T \in \operatorname{SYT}(\mu/\rho) : \text{the content reading word of } T \text{ is in } \operatorname{SYam}(\lambda)\}. (2.4.2)$$

We are now able to present the main theorem of this section.

Theorem 2.4.3. *Let $\boldsymbol{\nu}$ be any k-tuple of skew shapes with* $\operatorname{diam}(\boldsymbol{\nu}) \leq 3$. *Then*

$$G_{\boldsymbol{\nu}}(X;q) = \sum_{\lambda \vdash |\boldsymbol{\nu}|} \sum_{\mathbf{T} \in \mathbf{T}(\boldsymbol{\nu}, \lambda)} q^{\operatorname{inv}_k(\mathbf{T})} s_\lambda.$$

The proof of Theorem 2.4.3 is postponed until Section 2.4.2. As mentioned in Remark 2.4.2, the set of $\boldsymbol{\nu}$ such that $\operatorname{diam}(\boldsymbol{\nu}) \leq 3$ properly contains the set of $\boldsymbol{\nu}$ that are 2-tuples.

The next corollary follows immediately by applying Theorem 2.4.3 to the definition of modified Macdonald polynomials in (2.2.15). We also use the fact that tuples of ribbons in $\mathrm{TR}(\mu/\rho)$ have diameter less than or equal to three if and only if μ/ρ does not contain (3,3) or (4) as a subdiagram. This fact is easily shown by noticing that the cells of every tuple of ribbons in $\mathrm{TR}(\mu/\rho)$ have the same set of distinct shifted contents.

Corollary 2.4.4. *Let μ/ρ be a skew shape not containing $(3,3)$ or (4) as a subdiagram. Then*

$$\widetilde{H}_{\mu/\rho}(X;q,t) = \sum_{\lambda\vdash|\mu/\rho|} \sum_{T\in T(\mu/\rho,\lambda)} q^{\mathrm{inv}(T)} t^{\mathrm{maj}(T)} s_\lambda.$$

In particular, Corollary 2.4.4 applies to all $\widetilde{H}_\mu(X;q,t)$ with $\mu_1 \leq 3$ where $\mu_2 \leq 2$.

Remark 2.4.5. The conditions on $\boldsymbol{\nu}$ and μ/ρ in Theorem 2.4.3 and Corollary 2.4.4, respectively, are sharp in the following sense. Let $\lambda = (2,2)$ and $\boldsymbol{\nu} = ((2),(1),(1))$ or $((1),(1),(1),(1))$. In particular, $\mathrm{diam}(\boldsymbol{\nu}) = 4$, and $\boldsymbol{\nu}$ has diagram

$$\Big(\square\square, \square, \square \Big) \quad \text{or} \quad \Big(\square, \square, \square, \square \Big).$$

Then

$$G_{\boldsymbol{\nu}}(X;q)\big|_{q^2 s_\lambda} = 1 \quad \text{and} \quad \sum_{\mathbf{T}\in\mathbf{T}(\boldsymbol{\nu},\lambda)} q^{\mathrm{inv}_k(T)} s_\lambda\big|_{q^2 s_\lambda} = 0. \tag{2.4.3}$$

If $\lambda = (2,2)$ and $\mu = (4)$, then

$$\widetilde{H}_\mu(X;q,t)\big|_{q^2 s_\rho} = 1 \quad \text{and} \quad \sum_{T\in T(\mu,\lambda)} q^{\mathrm{inv}(T)} t^{\mathrm{maj}(T)} s_\lambda\big|_{q^2 s_\lambda} = 0. \tag{2.4.4}$$

If $\lambda = (4,2)$ and $\mu = (3,3)$, then

$$\widetilde{H}_\mu(X;q,t)\big|_{q^2 s_\rho} = 1 \quad \text{and} \quad \sum_{T\in T(\mu,\lambda)} q^{\mathrm{inv}(T)} t^{\mathrm{maj}(T)} s_\lambda\big|_{q^2 s_\lambda} = 0. \tag{2.4.5}$$

Corollary 2.4.6. *Let μ/ρ be a skew shape not containing $(2,2,2)$ or $(1,1,1,1)$ as a subdiagram. Then*

$$\widetilde{H}_{\mu/\rho}(X;q,t) = \sum_{\lambda\vdash|\mu/\rho|} \sum_{T\in T(\tilde{\mu}/\tilde{\rho},\lambda)} q^{\mathrm{maj}(T)} t^{\mathrm{inv}(T)} s_\lambda.$$

Proof: The corollary follows from (2.2.16) and Corollary 2.4.4. $\square$

In particular, Corollary 2.4.6 applies to all $\tilde{H}_\mu(X;q,t)$ where μ has at most three rows and $\mu_3 \leq 1$.

2.4.2 LLT Graphs

The goal of this section is to prove Theorem 2.4.3. We begin by following [Assaf, 2013] in defining an involution that will provide the edge sets of a signed colored graph. In this section, $\boldsymbol{\nu}$ will always denote a k-tuple of skew shapes whose sizes sum to $|\boldsymbol{\nu}| = n$. Also, w will always denote a permutation in S_n.

Let the involution $\tilde{d}_i : S_n \to S_n$ act by permuting the entries $i-1, i$, and $i+1$ as defined by,

$$\begin{aligned}
\tilde{d}_i(\ldots i-1 \ldots i \ldots i+1 \ldots) &= (\ldots i-1 \ldots i \ldots i+1 \ldots),\\
\tilde{d}_i(\ldots i+1 \ldots i \ldots i-1 \ldots) &= (\ldots i+1 \ldots i \ldots i-1 \ldots),\\
\tilde{d}_i(\ldots i \ldots i-1 \ldots i+1 \ldots) &= (\ldots i-1 \ldots i+1 \ldots i \ldots),\\
\tilde{d}_i(\ldots i \ldots i+1 \ldots i-1 \ldots) &= (\ldots i+1 \ldots i-1 \ldots i \ldots).
\end{aligned} \tag{2.4.6}$$

For instance, $\tilde{d}_3 \circ \tilde{d}_2(4123) = \tilde{d}_3(4123) = 3142$.

To decide when to apply d_i and when to use $\tilde{d}_i$, we appeal to the shifted content. Numbering the cells of a fixed $\boldsymbol{\nu}$ from 1 to n in shifted content reading order, let $\tilde{c}_i$ be the shifted content of the i^{th} cell. Define the weakly increasing word $\tau = \tau_1\tau_2 \ldots \tau_n$ by

$$\tau_i = \max\{j \in [n] : \tilde{c}_j - \tilde{c}_i \leq k\}. \tag{2.4.7}$$

See Figure 2.15 for an example. To emphasize the relationship between τ and $\boldsymbol{\nu}$, we will sometimes write $\tau = \tau(\boldsymbol{\nu})$. Notice that there are finitely many possible τ of any fixed length n. Specifically, τ will always satisfy $\tau_n = n$ and $i \leq \tau_i \leq \tau_{i+1}$ for all $i < n$. In fact, the number of possible τ is the n^{th} Catalan number (for details on the Catalan numbers and an extensive list on where they arise in mathematics, see [Stanley, 2012a]). Next, let $m(i)$ be the index of the the value in $\{i-1, i, i+1\}$ that occurs first in w, and let $M(i)$ be the index of the value in $\{i-1, i, i+1\}$ that occurs last in w. We now define the desired involution,

$$D_i^{(\tau)}(w) := \begin{cases} d_i(w) & \tau_{m(i)} < M(i) \\ \tilde{d}_i(w) & \tau_{m(i)} \geq M(i). \end{cases} \tag{2.4.8}$$

As an example, we may take $\tau = 456667899$ and $w = 534826179$, as in Figure 2.15. Then $D_3^{(\tau)}(w) = \tilde{d}_3(w) = 542836179$ and $D_5^{(\tau)}(w) = d_5(w) = 634825179$.

Figure 2.15: On the left, the shifted contents of a pair of skew diagrams with $\tau = 456667899$. On the right, a standard filling of the same tuple with shifted content word 534826179.

We may generalize $\mathcal{G}_n$ by defining $\mathcal{G}_n^{(\tau)}$ as the (n, n)-signed colored graph with vertex set indexed by S_n, signature function given by inverse descents, and each edge set E_i given by the nontrivial orbits of $D_i^{(\tau)}$. Direct inspection shows that if $\tau = \tau(\boldsymbol{\nu})$, then $D_i^{(\tau)}$ takes shifted content words of standard fillings of $\boldsymbol{\nu}$ to shifted content words of other standard fillings of $\boldsymbol{\nu}$. Thus, $D_i^{(\tau)}$ has a well-defined action on $\mathrm{SYT}(\boldsymbol{\nu})$ inherited from the action of $D_i^{(\tau)}$ on shifted content words. We may then define the following subgraph of $\mathcal{G}_n^{(\tau)}$.

Definition 2.4.7. Given some tuple of skew shapes $\boldsymbol{\nu}$, the *LLT graph* $\mathcal{L}_{\boldsymbol{\nu}} = (V, \sigma, E)$ is defined to be the (n, n)-signed colored graph with the following data:

1. $V = \{w \in S_n : w \text{ is the shifted content word of some } \mathbf{T} \in \mathrm{SYT}(\boldsymbol{\nu})\}$,
2. The signature function σ is given by the inverse descent sets of $w \in V$,
3. The edge sets E_i are defined by the nontrivial orbits of $D_i^{(\tau)}$ for all $1 < i < |\boldsymbol{\nu}|$, where $\tau = \tau(\boldsymbol{\nu})$.

Example 2.4.8. Consider $\boldsymbol{\nu} = ((2), (2), (1), (1))$. A portion of the LLT graph $\mathcal{L}_{\boldsymbol{\nu}}$ is presented in Figure 2.16. Here, $\mathcal{L}_{\boldsymbol{\nu}}$ is a subgraph of $\mathcal{G}_6^{(\tau)}$ with $\tau = 566666$. The entire connected component of the vertices in Figure 2.16 has 47 vertices. In the figure, the edge $\{312654, 412653\}$ is defined by the action of d_3 and d_4, while all other edges are defined by the action of $\tilde{d}_i$ for $1 < i < 6$.

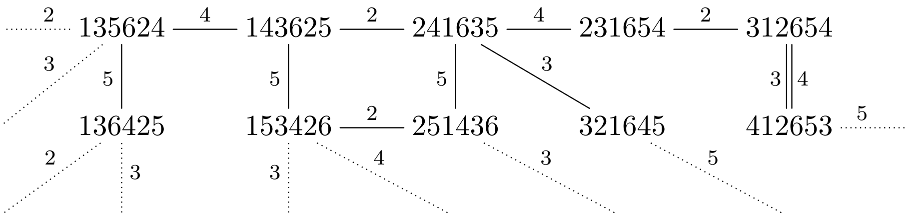

Figure 2.16: A portion of $\mathcal{L}_{\boldsymbol{\nu}}$ with signatures omitted. Here $\boldsymbol{\nu} = ((2), (2), (1), (1))$.

Remark 2.4.9.

1. If $\boldsymbol{\nu} = (\lambda/\rho)$ is a 1-tuple, then each $\mathbf{T} \in \mathrm{SYT}(\boldsymbol{\nu})$ corresponds to a unique $T \in \mathrm{SYT}(\lambda/\rho)$ in the obvious fashion. In particular, the shifted content word of $\mathbf{T}$ is equal to the content reading word of T. By Remark 2.3.3, $\mathcal{G}_{\lambda/\rho} \cong \mathcal{L}_{\boldsymbol{\nu}}$. In the case of straight shapes, $\boldsymbol{\nu} = (\lambda)$, Theorem 2.3.2 implies that $P\colon S_n \to \mathrm{SYT}(n)$ induces an isomorphism from $\mathcal{L}_{\boldsymbol{\nu}}$ to $\mathcal{G}_{\lambda}$.

2. Given some w in the vertex set of $\mathcal{L}_{\boldsymbol{\nu}}$, we may readily describe the component of w in a restriction of $\mathcal{L}_{\boldsymbol{\nu}}$. Let w be the shifted content word of $\mathbf{T} \in \mathrm{SYT}(\boldsymbol{\nu})$, and let $\mathcal{C}$ be the (m, m)-component of w in $\mathcal{L}_{\boldsymbol{\nu}}$. The isomorphism type of $\mathcal{C}$ can be found by removing the cells of $\mathbf{T}$ containing values in $\{m+1, \ldots, n\}$, creating some $\mathbf{T}' \in \mathrm{SYT}(\boldsymbol{\nu}')$ with shifted content word w'. Then $\mathcal{C}$ is isomorphic to the component of w' in $\mathcal{L}_{\boldsymbol{\nu}'}$.

   Similarly, if $\mathcal{C}$ is the component of w in the h-upward restriction of $\mathcal{L}_{\boldsymbol{\nu}}$, then the isomorphism type of $\mathcal{C}$ can be found by removing the cells of $\mathbf{T}$ containing values in $\{1, \ldots, h\}$ and subtracting h from each of the remaining cells, creating some $\mathbf{T}' \in \mathrm{SYT}(\boldsymbol{\nu}')$ with shifted content word w'. Then $\mathcal{C}$ is isomorphic to the component of w' in $\mathcal{L}_{\boldsymbol{\nu}'}$.

   In both cases above, $\boldsymbol{\nu}'$ is a tuple of skew shapes, so restrictions take components of LLT graphs to components of other LLT graphs. Furthermore, $\operatorname{diam}(\boldsymbol{\nu}') \leq \operatorname{diam}(\boldsymbol{\nu})$, since $\boldsymbol{\nu}'$ is obtained by removing cells from $\boldsymbol{\nu}$.

While LLT graphs do not necessarily satisfy Axiom 4 or Axiom 6, as can be seen in Example 2.4.8, they do satisfy a subset of the dual equivalence axioms. This is made precise in the following proposition and theorem.

Proposition 2.4.10 ([Assaf, 2013] Prop. 4.6)**.** *Any LLT graph $\mathcal{L}_{\boldsymbol{\nu}}$ obeys Axioms 1, 2, 3, and 5. Furthermore, the* $\operatorname{inv}_{\boldsymbol{\nu}}$ *statistic is constant on each connected component of $\mathcal{L}_{\boldsymbol{\nu}}$.*

With Proposition 2.4.10 in mind, it is natural to try to classify which LLT graphs satisfy all of the dual equivalence axioms. This classification is accomplished in the following theorem.

Theorem 2.4.11. *The LLT graph $\mathcal{L}_{\boldsymbol{\nu}}$ is a dual equivalence graph if and only if* $\operatorname{diam}(\boldsymbol{\nu}) \leq 3$.

Proof: First suppose $\operatorname{diam}(\boldsymbol{\nu}) \leq 3$. By Proposition 2.4.10, $\mathcal{L}_{\boldsymbol{\nu}}$ obeys Axioms 1, 2, 3, and 5. By Theorem 2.3.17, we need only show that $\mathcal{L}_{\boldsymbol{\nu}}$ satisfies Axiom 4^+ to prove that $\mathcal{L}_{\boldsymbol{\nu}}$ is a dual equivalence graph.

We begin by showing that $\mathcal{L}_{\boldsymbol{\nu}}$ obeys Axiom 4. Applying the same logic as Part 1 of Remark 2.3.18, we may use restrictions to only consider signed colored graphs of type $(5, 5)$. By Part 2 of Remark 2.4.9, each connected component of any such restriction is a component of $\mathcal{L}_{\boldsymbol{\mu}}$ for some $\boldsymbol{\mu}$ with $\operatorname{diam}(\boldsymbol{\mu}) \leq 3$. We may then reduce to the case where $|\boldsymbol{\nu}| = 5$ and $\operatorname{diam}(\boldsymbol{\nu}) \leq 3$. It thus suffices to consider $\mathcal{G}_5^{(\tau)}$ for all $\tau = \tau(\boldsymbol{\nu})$ with $\boldsymbol{\nu}$ satisfying these properties, in which case, it can be checked explicitly that all components of $\mathcal{G}_5^{(\tau)}$ that are contained in some $\mathcal{L}_{\boldsymbol{\nu}}$ with $\operatorname{diam}(\boldsymbol{\nu}) \leq 3$ are dual equivalence graphs. This fact is verified via computer at

<http://www.math.washington.edu/∼austinis/Proof_LLTandDEG.sws>,

cited as [Roberts, 2012].

To demonstrate Axiom 4^+, we may similarly inspect $\mathcal{G}_6^{(\tau)}$ for all τ that are derived from $\boldsymbol{\nu}$ with $|\boldsymbol{\nu}| = 6$ and $\operatorname{diam}(\boldsymbol{\nu}) \leq 3$. It is possible to explicitly check that $\mathcal{L}_{\boldsymbol{\nu}}$ is a dual equivalence graph by showing that the components of $\mathcal{G}_6^{(\tau)}$ that satisfy Axiom 4 also satisfy the hypotheses of Corollary 2.3.20. This fact is also verified via computer at [Roberts, 2012]. Thus, $\mathcal{L}_{\boldsymbol{\nu}}$ is a dual equivalence graph whenever $\operatorname{diam}(\boldsymbol{\nu}) \leq 3$.

Now suppose $\operatorname{diam}(\boldsymbol{\nu}) \geq 4$. By the definition of diameter, there exist four cells of $\boldsymbol{\nu}$ whose shifted contents are distinct and differ by at most k. If we label these four cells c_1, c_2, c_3, and c_4 in shifted content reading order, then c_1, c_2, and c_3 must occur in different skew tableaux of $\boldsymbol{\nu}$. Similarly, c_2, c_3, and c_4 occur in different skew tableaux of $\boldsymbol{\nu}$. There is then some value of i and some standard filling of $\boldsymbol{\nu}$ with values $i, i-1, i+2, i+1$ or $i+1, i+2, i-1, i$ in c_1, c_2, c_3, and c_4, respectively. Call the shifted content word of this standard filling w. Direct computation shows that w and $D_i^{(\tau)}(w) = \tilde{d}_i(w)$ are contained in distinct $(i+1)$-edges. That is, the restriction of $\mathcal{L}_{\boldsymbol{\nu}}$ to the edge colored graph $(V, E_i \cup E_{i+1})$ contains a component with at least three distinct edges. Therefore, $\mathcal{L}_{\boldsymbol{\nu}}$ violates Axiom 4 and is not a dual equivalence graph. □

Remark 2.4.12. By applying [Assaf, 2008/09, Prop. 5.3], it can be proven that $\mathcal{L}_{\boldsymbol{\nu}}$ is a dual equivalence graph whenever $\boldsymbol{\nu}$ is a pair of ribbons. Both the previous theorem and the following lemma can be viewed as extensions of this fact.

While the previous theorem describes which LLT graphs are dual equivalence graphs, there are other cases in which only specific components of an LLT graph are dual equivalence graphs. Sometimes, we may even give an explicit isomorphism from a component of an LLT graph to a standard dual equivalence graph. The next three results describe cases where such an explicit isomorphism exists.

Lemma 2.4.13. *Let $\mathcal{C} = (V, \sigma, E)$ be a component of $\mathcal{G}_n^{(\tau)}$ such that $D_i^{(\tau)}$ never acts on $w \in V$ nontrivially via $\tilde{d}_i$ unless $i-1, i$, and $i+1$ have adjacent indices in w. Then $\mathcal{C}$ is a dual equivalence graph, and $P\colon V \to \mathrm{SYT}(n)$ induces an isomorphism from $\mathcal{C}$ to some standard dual equivalence graph $\mathcal{G}_\lambda$.*

Proof : By Proposition 2.4.10, $\mathcal{C}$ satisfies Axioms 1, 2, 3, and 5. By Theorem 2.3.17, we need only demonstrate that $\mathcal{C}$ satisfies Axiom 4^+ to show that $\mathcal{C}$ is a dual equivalence graph. As in the proof of Theorem 2.4.11, we may reduce to the case where $\mathcal{C}$ has type (6,6). The requirement on the action of $D_i^{(\tau)}$ allows us to further reduce to the case where τ satisfies

$\tau_i \leq i+2$ for all $1 \leq i \leq n$. In order to show that $\mathcal{C}$ is a dual equivalence graph, it is thus sufficient to check that $\mathcal{G}_6^{(\tau)}$ is a dual equivalence graph for all τ of length 6 with $\tau_i \leq i+2$. This fact is verified via computer at [Roberts, 2012].

To demonstrate that P induces a morphism, we show that $\mathcal{C}$ satisfies the hypotheses of Theorem 2.3.13. That is, we need to show that $E_i(w)$ is in the Knuth class of $d_i(w)$ for all $w \in V$ and all $1 < i < n$. When $D_i^{(\tau)}$ acts via d_i, this edge requirement is clearly satisfied. Now notice that if $D_i^{(\tau)}$ acts via $\tilde{d}_i$, then the restriction on indices implies that $\tilde{d}_i = K_j \circ d_i$ for some j. Hence, $E_i(w)$ is in the same Knuth class as $d_i(w)$. Thus, P induces a morphism between connected dual equivalence graphs. This morphism must be an isomorphism by Corollary 2.3.12. □

Lemma 2.4.13 may be readily applied to LLT graphs because every LLT graph is a subgraph of some $\mathcal{G}_n^{(\tau)}$. While this lemma does give a particularly nice isomorphism, it does not apply in the more general case of LLT graphs indexed by tuples of shapes with diameter at most 3, as in Theorem 2.4.3. Instead, we will first need to identify specific components of LLT graphs satisfying the hypotheses of Lemma 2.4.13.

Lemma 2.4.14. *Let* $\boldsymbol{\nu}$ *be a tuple of skew shapes such that* $|\boldsymbol{\nu}| = n$ *and* $\mathrm{diam}(\boldsymbol{\nu}) \leq 3$. *Let* $\mathcal{C} = (V, \sigma, E)$ *be a connected component of* $\mathcal{L}_{\boldsymbol{\nu}}$, *let* $\tau = \tau(\boldsymbol{\nu})$, *and let* v *be any vertex in* $\mathcal{C}$. *If* $\mathrm{sh}(v)$ *has strictly less than three rows, then* $D_i^{(\tau)}$ *never acts on any vertex* $w \in V$ *nontrivially via* $\tilde{d}_i$ *unless* $i-1, i,$ *and* $i+1$ *have adjacent indices in* w.

Proof : Let v be the shifted content word of $\mathbf{T} = (T^{(0)}, \ldots, T^{(k-1)}) \in \mathrm{SYT}(\boldsymbol{\nu})$. We claim that it is enough to show that $D_i^{(\tau)}$ never acts on v nontrivially via $\tilde{d}_i$ unless $i-1, i$, and $i+1$ have adjacent indices in v. In this case, $D_i^{(\tau)}$ acts on v as the identity, d_i, or $K_j \circ d_i$ for some j, as was mentioned in the proof of Lemma 2.4.13. Thus $\mathrm{sh}(D_i^{(\tau)}(v)) = \mathrm{sh}(v)$. In particular, $D_i^{(\tau)}(v)$ satisfies the same hypotheses as v and recursively, so does every vertex of $\mathcal{C}$.

We may recharacterize the condition that $\mathrm{sh}(v)$ has strictly less than three rows in terms of the values in v. If $P(v)$ has strictly less than three rows, then v has no decreasing subword of length 3, as noted in Section 2.2.2. Suppose $D_i^{(\tau)}(v) = \tilde{d}_i(v) \neq v$ for some fixed choice of i, and let a, b, and c be the values $i-1, i$, and $i+1$ given in the order they appear in v.

Because $D_i^{(\tau)}$ acts nontrivially on v, either $a > b$ or $b > c$. Hence, there cannot be any value strictly greater than $i+1$ that occurs before a in v or any value strictly less than $i-1$ that occurs after c in v.

We can now show that a, b, and c occur consecutively in v by considering shifted contents. Let $\tilde{c}(x)$ denote the shifted content of the cell of $\mathbf{T}$ containing the value x. Suppose that there is a value x that appears after a in v with $\tilde{c}(a) = \tilde{c}(x)$. That is, x occurs above a on the same diagonal of some skew tableau $T^{(j)}$ of $\mathbf{T}$. Then there is a value y that occurs directly north of a and directly west of x in $T^{(j)}$. It follows that $y > i+1 \geq a$, and y occurs before a in v, a contradiction. Similarly, there cannot be a value x that occurs before c in v with $\tilde{c}(c) = \tilde{c}(x)$. In particular, $\tilde{c}(a) < \tilde{c}(b) < \tilde{c}(c)$.

Because $\operatorname{diam}(\boldsymbol{\nu}) \leq 3$, the cells of $\boldsymbol{\nu}$ can occupy at most one shifted content strictly between $\tilde{c}(a)$ and $\tilde{c}(c)$, so it suffices to show that there is exactly one cell in $\boldsymbol{\nu}$ with shifted content $\tilde{c}(b)$. Otherwise, there must be a cell on the same diagonal as b in some skew tableau of $\mathbf{T}$. There is then a value z in a cell either directly north of b or directly south of b with $\tilde{c}(z) = \tilde{c}(b) - k$ or $\tilde{c}(z) = \tilde{c}(b) + k$, respectively. Also, $z > b$ or $z < b$, respectively. Because $D_i^{(\tau)}(v) = \tilde{d}_i(v) \neq v$, $\tilde{c}(c) - \tilde{c}(a) \leq k$. Using the fact that $\tilde{c}(a) < \tilde{c}(b) < \tilde{c}(c)$, it follows that $\tilde{c}(b) - k < \tilde{c}(a)$ and $\tilde{c}(c) - \tilde{c}(b) < k$. Thus, z is greater than $i+1$ and occurs before a in v, or z is less than $i-1$ and occurs after c in v, forcing a contradiction. Therefore, a, b, and c occur consecutively in v. □

Corollary 2.4.15. *Let $\boldsymbol{\nu}$ be a tuple of skew shapes such that $|\boldsymbol{\nu}| = n$ and $\operatorname{diam}(\boldsymbol{\nu}) \leq 3$. Let $\mathcal{C} = (V, \sigma, E)$ be a connected component of $\mathcal{L}_{\boldsymbol{\nu}}$, and let w be any vertex in $\mathcal{C}$. If $\lambda = \operatorname{sh}(w)$ has strictly less than three rows, then $P\colon V \to SYT(n)$ induces an isomorphism from $\mathcal{C}$ to the standard dual equivalence graph $\mathcal{G}_\lambda$.*

Proof : The corollary follows immediately from Lemma 2.4.13 and Lemma 2.4.14.

We now change our focus from finding isomorphism types to finding a set of vertices to represent the components of an LLT graph. The following lemma is crucial to the proof of Theorem 2.4.3.

Lemma 2.4.16. *Let $\boldsymbol{\mu}$ and $\boldsymbol{\nu}$ be tuples of skew shapes such that $\operatorname{diam}(\boldsymbol{\mu}), \operatorname{diam}(\boldsymbol{\nu}) \leq 3$ and*

$|\boldsymbol{\mu}| = |\boldsymbol{\nu}| = n$. Let $\mathcal{C}$ and $\mathcal{D}$ be connected components of $\mathcal{L}_{\boldsymbol{\mu}}$ and $\mathcal{L}_{\boldsymbol{\nu}}$, respectively, and let $\phi\colon \mathcal{C} \to \mathcal{D}$ be an isomorphism. If w is a vertex in $\mathcal{C}$ and $\lambda \vdash n$, then $w \in \mathrm{SYam}(\lambda)$ if and only if $\phi(w) \in \mathrm{SYam}(\lambda)$.

Proof. Suppose, for the sake of contradiction, that $v = \phi(w) \in \mathrm{SYam}(\lambda)$ and $w \notin \mathrm{SYam}(\lambda)$. We begin by setting some definitions. Let $P(w) = T$. In particular, $P(\phi(w)) = U_\lambda \neq T$. Consider the lowest value that does not occur in the same cell of T and U_λ. By signature considerations, this value must be in a lower row of T than in U_λ. Let m be the smallest number such that some entry of the m^{th} row of U_λ occurs in a lower row of T (see Figure 2.17 for an example). Now define

$$p = \sum_{j=1}^{m-2} \lambda_j, \qquad q = \lambda_{m-1} + \lambda_m, \qquad S_i = p + i, \qquad \text{and} \qquad S = \{S_1, S_2, \ldots, S_q\}. \tag{2.4.9}$$

That is, S is the set of values in rows $m-1$ and m of U_λ. Notice that $m > 1$, so S is nonempty. Let w_S and v_S be the subwords of w and v, respectively, with values in S.

We now consider the longest increasing subwords of w_S and v_S. Because v is Knuth equivalent to the row reading word of U_λ, we may restrict the row reading word of U_λ to values in S in order to find the longest increasing subword of v_S, as described in Section 2.2.2. Specifically, the longest increasing subword of v_S has length λ_{m-1}. Similarly, we may restrict the row reading word of T to find the longest increasing subword of w_S. By the definitions of m and S, the values S_1 through $S_{\lambda_{m-1}}$ occur in row $m-1$ of T, while the value S_q occurs no higher than row $m-1$ of T. Hence, the row reading word of T has $S_1 S_2 \ldots S_{\lambda_{m-1}} S_q$ as an increasing subword (see Figure 2.17 for an example). In particular, if l is the length of the longest increasing subword of w_S, then $l > \lambda_{m-1}$. We will use these facts about w_S and v_S to create a contradiction to the assumption that $T \neq U_\lambda$.

$$T = \begin{array}{|c|c|c|c|c|} \cline{1-1} \mathbf{7} \\ \cline{1-3} \mathbf{4} & \mathbf{5} & \mathbf{6} \\ \hline 1 & 2 & 3 & \mathbf{8} & \mathbf{9} \\ \hline \end{array} \qquad\qquad U_{(3,3,3)} = \begin{array}{|c|c|c|} \hline \mathbf{7} & \mathbf{8} & \mathbf{9} \\ \hline \mathbf{4} & \mathbf{5} & \mathbf{6} \\ \hline 1 & 2 & 3 \\ \hline \end{array}$$

Figure 2.17: An example where T has the same signature as $U_{(3,3,3)}$. Here, $m = 3, p = 3$, $q = 6, S = \{4, 5, 6, 7, 8, 9\}$, $l = 5$, and the subword $S_1 S_2 \ldots S_{\lambda_{m-1}} S_q = 4569$.

Consider the (q, q)-component of w in the p-upward restriction of $\mathcal{C}$. Call this component $\mathcal{C}'$. We proceed by finding the isomorphism type of $\mathcal{C}'$ in two different ways. First, we may find the isomorphism type of $\mathcal{C}'$ in the manner described by Part 2 of Remark 2.4.9. Specifically, consider $\mathbf{T} \in \mathrm{SYT}(\boldsymbol{\mu})$ such that w is the shifted content word of $\mathbf{T}$. Remove the cells of $\mathbf{T}$ not containing values in S and then subtract p from each of the remaining cells, creating some $\mathbf{T}'$ of shape $\boldsymbol{\mu}'$ with shifted content word w'. Then $\mathcal{C}'$ is isomorphic to the component of $\mathcal{L}_{\boldsymbol{\mu}'}$ containing w'. Here, w' can also be found by subtracting p from each entry in w_S.

To find the isomorphism type of the component of w' in $\mathcal{L}_{\boldsymbol{\mu}'}$, we will apply Corollary 2.4.15. The signature of w' is equal to the restriction of the signature of w to the coordinates S_1 through S_{q-1}. The signature of w is equal to the signature of v, which is equal to the signature of U_λ. Thus, the signature of w' is equal to the restriction of the signature of U_λ to the coordinates S_1 through S_{q-1}. In particular, the signature of w' has exactly one -1, implying that $P(w')$ has exactly two rows. Applying Corollary 2.4.15, P induces an isomorphism from $\mathcal{C}'$ to the standard dual equivalence graph $\mathcal{G}_{\mathrm{sh}(w')}$. The length of the first row in $\mathrm{sh}(w')$ is the length of the longest increasing subword in w', as noted in Section 2.2.2. The length of this subword is equal to the length of the longest increasing subword of w_S, which we know to have length $l > \lambda_{m-1}$. Since $\mathrm{sh}(w')$ has exactly two rows, it follows that $\mathrm{sh}(w') = (l, q - l)$. Therefore, $\mathcal{C}' \cong \mathcal{G}_{(l,q-l)}$.

Alternatively, we may find the isomorphism type of $\mathcal{C}'$ by applying the same restrictions to $\mathcal{D}$. Let $\mathcal{D}'$ be the (q, q)-component of v in the p-upward restriction of $\mathcal{D}$. In particular, $\mathcal{C}' \cong \mathcal{D}'$. As before, $\mathcal{D}'$ is isomorphic to the component of some $\mathcal{L}_{\boldsymbol{\nu}'}$ containing the vertex v', where v' is obtained by subtracting p from each of the values in v_S. Also as before, $\mathrm{sh}(v')$ has two rows. Therefore Corollary 2.4.15 guarantees that $D' \cong \mathcal{G}_{\mathrm{sh}(v')}$. The length of the longest increasing subword of v' is equal to the length of the longest increasing subword of v_S, which we know to have length λ_{m-1}. Thus, $\mathrm{sh}(v') = (\lambda_{m-1}, \lambda_m)$. Since $l > \lambda_{m-1}$, $\mathcal{C}' \cong \mathcal{G}_{(l,q-l)} \ncong \mathcal{G}_{(\lambda_{m-1},\lambda_m)} \cong \mathcal{D}'$, a contradiction. Therefore, $w \in \mathrm{SYam}(\lambda)$ whenever $v = \phi(w) \in \mathrm{SYam}(\lambda)$.

We still need to prove that $\phi(w) \in \mathrm{SYam}(\lambda)$ whenever $w \in \mathrm{SYam}(\lambda)$. This follows via symmetry by considering w as the image of $\phi(w)$ under the isomorphism ϕ^{-1}. Thus,

$w \in \mathrm{SYam}(\lambda)$ if and only if $\phi(w) \in \mathrm{SYam}(\lambda)$. $\square$

Proof of Theorem 2.4.3:

We begin by reducing Theorem 2.4.3 to a statement about signed colored graphs. Let $V_1, V_2, \ldots, V_m$ be the vertex sets of the connected components of $\mathcal{L}_{\boldsymbol{\nu}} = (V, \sigma, E)$. Applying (2.2.8) and Definition 2.4.7,

$$G_{\boldsymbol{\nu}}(X;q) = \sum_{\mathbf{T} \in \mathrm{SYT}(\boldsymbol{\nu})} q^{\mathrm{inv}_k(\mathbf{T})} F_{\sigma(\mathbf{T})}(X) = \sum_{v \in V} q^{\mathrm{inv}_{\boldsymbol{\nu}}(v)} F_{\sigma(v)}(X) = \sum_{j=1}^{m} \sum_{v \in V_j} q^{\mathrm{inv}_{\boldsymbol{\nu}}(v)} F_{\sigma(v)}(X), \tag{2.4.10}$$

where it should be noted that the signature function σ changes between the first and second sums. Specifically, σ changes from the signature function on standard fillings of $\boldsymbol{\nu}$ to the signature function given by the inverse descents of the permutations in the vertex set of $\mathcal{L}$. The second equality then follows because the signature of a standard filling of $\boldsymbol{\nu}$ is defined via the signature of its row reading word.

To further simplify (2.4.10), we turn our attention to the individual components of $\mathcal{L}_{\boldsymbol{\nu}}$ when $\mathrm{diam}(\boldsymbol{\nu}) \leq 3$. Let $\mathcal{C} = (V_j, \sigma, E)$ be a connected component of $\mathcal{L}_{\boldsymbol{\nu}}$, and choose any fixed $v_j \in V_j$. By Proposition 2.4.10, the $\mathrm{inv}_{\boldsymbol{\nu}}$ statistic is constant on V_j. By Theorems 2.2.12 and 2.4.11, $\mathcal{C}$ is isomorphic to $\mathcal{G}_\lambda$ for some $\lambda \vdash |\boldsymbol{\nu}|$. From (2.2.21), it follows that

$$\sum_{v \in V_j} q^{\mathrm{inv}_{\boldsymbol{\nu}}(v)} F_{\sigma(v)}(X) = q^{\mathrm{inv}_{\boldsymbol{\nu}}(v_j)} s_\lambda. \tag{2.4.11}$$

To prove Theorem 2.4.3, we need only guarantee that there is a unique choice of $v_j \in V_j$ such that v_j is a standardized Yamanouchi word and, moreover, that this choice of v_j is in $\mathrm{SYam}(\lambda)$.

We will find such a v_j explicitly by applying Lemma 2.4.16. Consider the LLT graph $\mathcal{L}_{\boldsymbol{\lambda}}$, where $\boldsymbol{\lambda}$ is the 1-tuple (λ). It follows from Part 2 of Remark 2.4.9 that $\mathrm{diam}(\boldsymbol{\lambda}) \leq 2$. By Part 1 of Remark 2.4.9, P induces an isomorphism from $\mathcal{L}_{\boldsymbol{\lambda}}$ to $\mathcal{G}_\lambda$. In particular, the vertex set of $\mathcal{L}_{\boldsymbol{\lambda}}$ has exactly one vertex w such that $P(w) = U_\lambda$. Thus, $w \in \mathrm{SYam}(\lambda)$, and w

is the only vertex of $\mathcal{L}_{\boldsymbol{\lambda}}$ that is a standardized Yamanouchi word. Because $\mathcal{C}$ is isomorphic to $\mathcal{G}_\lambda$, and $\mathcal{G}_\lambda$ is isomorphic to $\mathcal{L}_{\boldsymbol{\lambda}}$, there exists an isomorphism $\phi\colon \mathcal{L}_{\boldsymbol{\lambda}} \to \mathcal{C}$. Applying Lemma 2.4.16, $v_j = \phi(w)$ is the unique standardized Yamanouchi word in V_j. $\square$

2.5 Conclusion

There are a number of persistent open questions involving dual equivalence graphs. We conclude by mentioning some of these questions as possibilities for further research.

Can dual equivalence graphs be used to give nice Schur expansions for symmetric functions other than LLT and Macdonald polynomials? Many functions have known expansions in terms of fundamental quasisymmetric functions but are lacking nice descriptions for their Schur expansions. As was the case with LLT polynomials, it may be possible to apply dual equivalence graphs to find Schur expansions for such functions. Examples include plethysms of Schur functions and k-Schur functions. The former is described in [Loehr and Warrington, 2012], while the latter has already seen some progress in [Assaf and Billey, 2012]. Though currently unproven, the shuffle conjecture provides another example. This conjecture describes the composition of the nabla operator with an elementary symmetric function in terms of two statistics and fundamental quasisymmetric functions. This sum, in turn, is equal to a sum of LLT polynomials with an additional statistic. It may then be possible to use the results of Section 2.4 as an aid to understanding the combinatorics of the shuffle conjecture. For the original statement of the shuffle conjecture, along with proofs of the facts mentioned above, see [Haglund et al., 2005b].

Is there a good description of when a component of an LLT graph is a dual equivalence graph? While Theorem 2.4.11 classifies when an LLT graph is a dual equivalence graph, it makes no claims about specific components of LLT graphs in the $\mathrm{diam}(\boldsymbol{\nu}) \geq 4$ case. Similarly, there is currently no good description for when a specific component of $\mathcal{G}_n^{(\tau)}$ is a dual equivalence graph.

Is there an axiomatic description for when a signed colored graph admits a morphism onto a standard dual equivalence graph? While Theorem 2.3.13 gives one criterion for when a signed colored graph admits a morphism onto a standard dual equivalence graph, a more axiomatic description would be desirable. Such a description would necessarily provide sufficient conditions for when a signed colored graph corresponds to a positive integer multiple of a Schur function.

Can the axiomatization of dual equivalence graphs be generalized to an axiomatization of the family of signed colored graphs that correspond to Schur positive functions? Such an axiomatization would necessarily be less strict than the one given for dual equivalence graphs and would thus be satisfied by a larger set of signed colored graphs. An axiomatization of this family of graphs could provide expanded methods for proving the Schur positivity of a variety of symmetric functions.

Chapter 3

ON THE SCHUR EXPANSION OF HALL-LITTLEWOOD AND RELATED POLYNOMIALS VIA YAMANOUCHI WORDS

3.1 Introduction

Adriano Garsia posed the question, when can the modified Hall-Littlewood polynomial $\widetilde{H}_\mu(X;0,t)$ be expanded in terms of Schur functions as a particular sum over Yamanouchi words, and is there a way to fix the expansion when it is not? At the time, he had already realized that the expansion he had in mind does not apply when $\widetilde{H}_\mu(X;0,t)$ is indexed by the partition shape $\mu = (3,3,3)$. The results of this paper are in direct response to Garsia's question. In fact, the results we found proved to be more general than the question as originally posed.

In this paper, we will concentrate on three main families of polynomials, the modified Macdonald polynomials $\widetilde{H}_\mu(X;q,t)$, the modified Hall-Littlewood polynomials $\widetilde{H}_\mu(X;0,t)$, and a refinement of the modified Hall-Littlewood polynomials, which we denote $R_{\gamma,\delta}(X)$. Firstly, the Macdonald polynomials were introduced in [Macdonald, 1988] and are often defined as the set of q,t-symmetric functions that satisfy certain orthogonality and triangularity conditions. The modified Macdonald polynomials were shown to be Schur positive by Mark Haiman via representation-theoretic and geometric means in [Haiman, 2001]. Adhering to recent combinatorial work in [Haglund et al., 2005a] and [Assaf, 2013], we choose to work with the modified Macdonald polynomials, though it is relatively straightforward to transition to other forms of Macdonald polynomials. For instance, the modified Macdonald polynomials may be transformed into the integral form of the Macdonald polynomials $H_\mu(X;q,t)$ by using $H_\mu(X;q,t) = t^{n(\mu)}\widetilde{H}_\mu(X;q,1/t)$, where $n(\mu)$ is a constant given by the highest exponent of t in $\widetilde{H}_\mu(X;q,t)$. Macdonald polynomials also specialize to several well known functions, including Hall-Littlewood polynomials and Jack polynomials. While combinatorial descriptions of the Schur expansions of specific families of Macdonald polynomials can be found in [Haglund et al., 2005a] (which drew on the earlier work in [Carré and Leclerc, 1995] and [van Leeuwen, 2001]), [Fishel, 1995], [Zabrocki, 1998], [Zabrocki, 1999], [Lapointe and Morse, 2003], [Assaf, 2008/09], and [Roberts, 2014], an explicit combinatorial description remains elusive outside of these special cases. For details on the combinatorics of Macdonald polynomials, see [Haglund, 2008].

As just noted, Macdonald polynomials specialize to Hall-Littlewood polynomials, specifically by letting $q = 0$. Hall-Littlewood polynomials, in turn, specialize to the Schur functions, Schur Q-functions, and the monomial symmetric functions. They were first studied by Philip Hall in relation to the Hall algebra in [Hall, 1957] and later by D.E. Littlewood in [Littlewood, 1961]. It should be noted that the earliest known work on Hall-Littlewood polynomials actually dates back to the lectures of Ernst Steinitz in [Steinitz, 1901]. Similar to the case with Macdonald polynomials there are several easily related forms of Hall-Littlewood polynomials. In particular, a transformed Hall-Littlewood polynomial, $H_\mu(X, 0, t)$, is related to a modified Hall-Littlewood polynomial via $H_\mu(X; 0, t) = t^{n(\mu)}\widetilde{H}_\mu(X; 0, 1/t)$. Hall-Littlewood polynomials have proven to be a rich mathematical topic, with recent combinatorial work including (but certainly not limited to) [Nakayashiki and Yamada, 1997], [Carbonara, 1998], [Dalal and Morse, 2012],and [Loehr et al., 2013]. Expanding modified Hall-Littlewood polynomials into Schur functions can be achieved via the charge statistic, as found in [Lascoux and Schützenberger, 1978], though we will present a new expansion in this paper as a sum over a subset of the Yamanouchi words. Equivalently, our result gives a new combinatorial rule for the coefficients of the Kostka-Foulkes polynomial in one variable t. For more details on the topic of Hall-Littlewood polynomials, see [Macdonald, 1995].

We use the statistics defined in [Haglund et al., 2005a] to generalize the definition of the modified Macdonald polynomials $\widetilde{H}_\mu(X; q, t)$ and the modified Hall-Littlewood polynomials $\widetilde{H}_\mu(X; 0, t)$ to any diagram $\delta \subset \mathbb{Z} \times \mathbb{Z}$, giving the functions $\widetilde{H}_\delta(X; q, t)$ and $\widetilde{H}_\delta(X; 0, t)$. We may then write $\widetilde{H}_\delta(X; 0, t)$ in terms of the refinement polynomials $R_{\gamma,\delta}(X)$, defined via row reading words of fillings of δ with a fixed descent set $\gamma \subset \delta$. We will discuss these polynomials in the general context of diagrams, though the reader with a refined taste for the specific is free to replace δ with a partition shape. We may then write the main theorem of this paper as follows.

Theorem 3.1.1. *If γ and δ are any diagrams such that $\gamma \subset \delta$, then*

$$\widetilde{H}_\delta(X;0,t) = \sum_{\lambda \vdash |\delta|} \sum_{\substack{w \in \mathrm{Yam}_\delta(\lambda) \\ \mathrm{inv}_\delta(w)=0}} t^{\mathrm{maj}_\delta(w)} s_\lambda, \qquad R_{\gamma,\delta}(X) = \sum_{\lambda \vdash |\delta|} \sum_{\substack{w \in \mathrm{Yam}_\delta(\lambda) \\ \mathrm{inv}_\delta(w)=0 \\ \mathrm{Des}_\delta(w)=\gamma}} s_\lambda. \tag{3.1.1}$$

Here, $\mathrm{Yam}_\delta(\lambda)$ is the subset of the Yamanouchi words with content λ whose elements, when thought of as row reading words of a filling of the diagram δ, never have the j^{th} from last i in the same pistol of δ as the $j+1^{th}$ from last $i+1$. The above definitions and notation will be given a more thorough treatment in Section 3.2.

The main tool used in the proof of Theorem 3.1.1 is the theory of dual equivalence graphs. Dual equivalence has its roots in the work of Schützenberger in [Schützenberger, 1977], Mark Haiman in [Haiman, 1992], and Donald Knuth in [Knuth, 1970]. Sami Assaf introduced the theory of dual equivalence graphs in her Ph.D. dissertation [Assaf, 2007] and subsequent paper [Assaf, 2013]. The theory was further advanced by the author in [Roberts, 2014], from which we will derive the definition of dual equivalence graph used in this paper. In these papers, a dual equivalence graph is associated to a symmetric function so that each component of the graph corresponds to a single Schur function. Thus, the Schur expansion of said symmetric function is described by a sum over the set of components of the graph. Variations of dual equivalence graphs have also been given for k-Schur functions in [Assaf and Billey, 2012], for the product of a Schubert polynomial with a Schur polynomial in [Assaf et al., 2012], and for shifted tableaux in relation to the type B Lie group in [Billey et al.] and [Assaf, 2014]. Dual equivalence graphs were also connected to Kazhdan-Lusztig polynomials and W-graphs by Michael Chmutov in [Chmutov, 2013].

This paper will focus on dual equivalence graphs that emerge as components of a larger family of graphs. The involution $D_i^\delta \colon S_n \to S_n$ was first introduced in [Assaf, 2007] and can be used to define the edge sets of a signed colored graph $\mathcal{H}_\delta$ with vertex set S_n and vertices labeled by the signature function σ, which is defined via the inverse descent sets of permutations. We may then associate $\widetilde{H}_\mu(X;0,t)$ and $R_{\gamma,\delta}(X)$ to subgraphs of $\mathcal{H}_\delta$. We show that these two subgraphs are dual equivalence graphs in Theorem 3.3.2. The main contribution of this paper to the theory of dual equivalence graphs can then be stated in the following theorem.

Theorem 3.1.2. *Let* δ *be a diagram of size* n*, and let* $\mathcal{G} = (V, \sigma, E)$ *be a dual equivalence graph such that* $\mathcal{G}$ *is a component of* $\mathcal{H}_\delta$ *and* $\mathcal{G} \cong \mathcal{G}_\lambda$*. Then there is a unique vertex of* V *in* $\mathrm{SYam}_\delta(\lambda)$*, and* $V \cap \mathrm{SYam}_\delta(\mu) = \emptyset$ *for all* $\mu \neq \lambda$*.*

Here, $\mathrm{SYam}_\delta(\lambda)$ is the set of permutations resulting from standardizing the words in $\mathrm{Yam}_\delta(\lambda)$.

This paper is organized as follows. We begin with the necessary material from the literature in Section 3.2, discussing tableaux, symmetric functions, and dual equivalence graphs. In Section 3.3, we give a classification of which connected components of $\mathcal{H}_\delta$ are dual equivalence graphs in Lemma 3.3.1, and use this to prove that the signed colored graphs associated to $R_{\gamma,\delta}(X)$ and $\widetilde{H}_\delta(X; 0, t)$ are dual equivalence graphs. We then prove Theorem 3.1.2, followed by Theorem 3.1.1, as well as some related results. Next, Conjecture 3.3.8 gives a possible direction towards the Schur expansion of Macdonald polynomials. Section 3.4 is dedicated to further analysis of $\widetilde{H}_\mu(X; q, t)$ and $\widetilde{H}_\mu(X; 0, t)$. After classifying when $\widetilde{H}_\mu(X; q, t)$ and $\widetilde{H}_\mu(X; 0, t)$ expand via Yamanouchi words in Corollary 3.4.1 and Proposition 3.4.2, we then end by classifying when $R_{\gamma,\delta}(X) = 0$ in Proposition 3.4.4 and giving a description of the leading term in the Schur expansion of $R_{\gamma,\delta}(X)$ in Proposition 3.4.6.

3.2 Preliminaries

3.2.1 Tableaux and Permutations

By a *diagram* δ, we mean a finite subset of $\mathbb{Z} \times \mathbb{Z}$. We let $|\delta|$ denote the size of this subset. By reflecting a diagram δ over the line $x = y$ in the Cartesian plane, we may obtain the *conjugate* diagram, denoted δ'. A *partition* λ is a weekly decreasing finite sequence of nonnegative integers $\lambda_1 \geq \ldots \geq \lambda_k \geq 0$. We write $|\lambda| = n$ or $\lambda \vdash n$ if $\sum \lambda_i = n$. We will give the diagram of a partition in French notation by drawing left justified rows of boxes, where λ_i is the number of boxes in the i^{th} row, from bottom to top, with bottom left cell at the origin, as in the left diagram of Figure 3.1. Any diagram that arrises from a partition in this fashion is said to have *partition shape.* Notice that conjugating a partition shape provides another partition shape.

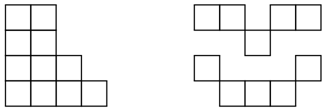

Figure 3.1: The diagrams for (4,3,2,2) and an arbitrary $\delta \subset \mathbb{Z} \times \mathbb{Z}$.

A *filling* is a function that takes each cell of a diagram δ to a positive integer. We express a filling visually by writing the value assigned to a cell inside of the cell. A *standard* filling uses each value in some $[n] = \{1, \ldots, n\}$ exactly once. Here, we say that T is a standard filling of δ, and we define $\mathrm{SF}(\delta)$ as the set of standard fillings of δ. A *standard Young tableau* is a standard filling in which all values are required to be increasing up columns and across rows from left to right. The set of all standard Young tableaux on diagrams of partition shape λ is denoted by $\mathrm{SYT}(\lambda)$, and the union of $\mathrm{SYT}(\lambda)$ over all $\lambda \vdash n$ is denoted $\mathrm{SYT}(n)$. For more information, see [Fulton, 1997, Part I], [Sagan, 1991, Ch. 3], or [Stanley, 2001, Ch. 7].

Define the *row reading word* of a filling T, denoted $\mathrm{rw}(T)$, by reading across rows from left to right, starting with the top row and working down, as in Figure 3.2. The row reading word of a standard filling is necessarily a permutation. For ease of reading, we will use π for permutations and use w in the more general context of words. Given a word w and diagram δ, both of size n, we let $T_\delta(w)$ denote the filling of diagram δ with reading word w.

Figure 3.2: On the left, a standard young tableau with row reading word 483691257. On the right, a standard filling with row reading word 3214.

By a *pistol* of a diagram δ, we mean a set of cells, in row reading order, between some cell c and the position one row below c, inclusive. By a pistol of a filling T, we mean a pistol of the diagram of T. If $|\delta| = n$, we may associate each cell to a number in $[n]$ in row reading order. In turn, this associates each pistol of δ with an interval $I \subset [n]$. In this case, we say that I is a pistol of δ and any collection of numbers in I is said to be *δ-pistoled.* In particular, given a word w of length $|\delta|$, any collection of indices corresponding to cells in a pistol of $T_\delta(w)$ is δ-pistoled.

Figure 3.3: Four pistols filled with bullets.

Given a permutation π in one-line notation, the signature of π is a string of 1's and -1's, or $+$'s and $-$'s for short, where there is a $+$ in the i^{th} position if and only if i comes before $i+1$ in π. Notice that a permutation is one entry longer than its signature. The signature of a standard filling T is defined as $\sigma(T) := \sigma(\mathrm{rw}(T))$. As an example, the signatures of the standard fillings in Figure 3.2 are $+--+-+-+$ and $--+$, respectively.

The Robinson-Schensted-Knuth (R-S-K) correspondence gives a bijection between permutations in S_n and pairs of standard Young tableaux (P, Q), where P and Q have the same shape $\lambda \vdash n$. Here, P is called the *insertion tableau* and Q is called the *recording tableau.* Let $P\colon S_n \to \mathrm{SYT}(n)$ be the function taking a permutation to its insertion tableau. Given any $T \in \mathrm{SYT}(n)$, the set of permutations π such that $P(\pi) = T$ is called a *Knuth class.*

For $I \subset [n]$ and $\pi \in S_n$, let $\pi|_I$ be the word given by reading the values of I in the order they appear in π. For any such I, $\pi|_I$ is referred to as a *subword* of π. Given any standard filling T of size n, we let $T|_I$ denote the filling that results from removing all cells of T with values not in I. Given a set of permutations $S \subset S_n$, we let $S|_I = \{\pi|_I \colon \pi \in S\}$.

We may *standardize* a word w of length n by replacing the values in w with the values in $[n]$ as follows. If there are k many 1's, replace the 1's in w with the values 1 through k from left to right. Then replace the 2's in w in a similar fashion, replacing the first 2 with the value $k+1$. Repeat this recursively until w has been replaced by a permutation, which we will denote $\mathrm{st}(w)$. Notice $\mathrm{st}(\pi) = \pi$ for all $\pi \in S_n$. We may *unstandardize* a permutation π by replacing each value i with 1 plus the number of -1's in $\sigma(\pi|_{[i-1]})$, resulting in a word that we will denote $\mathrm{unst}(\pi)$. That is, i and $i+1$ are taken to the same value if i occurs before $i+1$. Otherwise, $i+1$ is taken to the value that is one lager than that of i. We may unstandardize a word w by letting $\mathrm{unst}(w) = \mathrm{unst}(\mathrm{st}(w))$. Notice that $\mathrm{st}(\mathrm{unst}(\pi)) = \pi$ and $\mathrm{unst}(\mathrm{unst}(w)) = \mathrm{unst}(w)$. It can also be shown that for any word w, both $\mathrm{st}(w)$ and $\mathrm{unst}(w)$ have the same recording tableau as w.

For $\lambda \vdash n$, let U_λ be the standard Young tableau of shape λ given by placing the numbers in $[n]$ in order across the first row of λ, then across the second row, and so on. Now define $\mathrm{SYam}(\lambda) := \{\pi \in S_n \colon P(\pi) = U_\lambda\}$, and call this set of permutations the *standardized Yamanouchi words of shape* λ. Let $\mathrm{Yam}(\lambda)$ denote the set of all words w of length n such that there are never more $i+1$'s than i's while reading from right to left with the further requirement that i occurs λ_i times in w. Any such word is called a Yamanouchi word. It then follows from properties of the R-S-K correspondence that $\mathrm{SYam}(\lambda) = \{\mathrm{st}(w) \colon w \in \mathrm{Yam}(\lambda)\}$. Similarly, $\mathrm{Yam}(\lambda) = \{\mathrm{unst}(\pi) \colon \pi \in \mathrm{SYam}(\lambda)\}$.

The following definitions are referenced in the statement of Theorems 3.1.1 and 3.1.2.

Definition 3.2.1. Let δ be any diagram, and let w be any Yamanouchi word of length $|\delta|$. We say that *w jams δ* if there exists some i and some j such that indices of the j^{th} from last i in w and the $j+1^{th}$ from last $i+1$ in w are δ-pistoled. A standardized Yamanouchi word is said to jam δ if $\mathrm{unst}(\pi) = w$ jams δ. In the context of having such a w, we refer to the index of w containing the j^{th} from last i as jamming said pistol of δ.

We may then define,

$$\begin{aligned} \mathrm{Yam}_\delta(\lambda) &:= \{w \in \mathrm{Yam}(\lambda) \colon w \text{ does not jam } \delta\}, \\ \mathrm{SYam}_\delta(\lambda) &:= \{\pi \in \mathrm{SYam}(\lambda) \colon \pi \text{ does not jam } \delta\}, \end{aligned} \tag{3.2.1}$$

with examples of each set given in Figure 3.4.

Remark 3.2.2.

1. One method for listing Yamanouchi words is to begin with the number 1 and add numbers to the left of it, as in the description of Yamanouchi words in Section 3.2.1. The further condition that a word not jam δ simply means that upon adding the $j+1^{th}$

$$\mathrm{Yam}_{(3,3)}(2,2,2) = \left\{\mathrm{rw}\left(\begin{array}{|c|c|c|}\hline 3&2&1\\\hline 3&2&1\\\hline\end{array}\right), \mathrm{rw}\left(\begin{array}{|c|c|c|}\hline 3&2&3\\\hline 1&2&1\\\hline\end{array}\right)\right\}, \quad \mathrm{rw}\left(\begin{array}{|c|c|c|}\hline 3&3&\mathbf{2}\\\hline 1&2&\mathbf{1}\\\hline\end{array}\right), \mathrm{rw}\left(\begin{array}{|c|c|c|}\hline \mathbf{3}&2&3\\\hline \mathbf{2}&1&1\\\hline\end{array}\right) \notin \mathrm{Yam}_{(3,3)}(2,2,2)$$

$$\mathrm{SYam}_{(3,3)}(2,2,2) = \left\{\mathrm{rw}\left(\begin{array}{|c|c|c|}\hline 5&3&1\\\hline 6&4&2\\\hline\end{array}\right), \mathrm{rw}\left(\begin{array}{|c|c|c|}\hline 5&3&6\\\hline 1&4&2\\\hline\end{array}\right)\right\}, \quad \mathrm{rw}\left(\begin{array}{|c|c|c|}\hline 5&6&\mathbf{3}\\\hline 1&4&\mathbf{2}\\\hline\end{array}\right), \mathrm{rw}\left(\begin{array}{|c|c|c|}\hline \mathbf{5}&3&6\\\hline \mathbf{4}&1&2\\\hline\end{array}\right) \notin \mathrm{SYam}_{(3,3)}(2,2,2)$$

Figure 3.4: At left, the sets $\mathrm{Yam}_{(3,3)}(2,2,2)$ and $\mathrm{SYam}_{(3,3)}(2,2,2)$. At right, examples of words in $\mathrm{Yam}(2,2,2)$ and $\mathrm{SYam}(2,2,2)$ that jam $(3,3)$. The bottom row is achieved by standardizing the top row.

$i+1$, it needs to be checked that this $i+1$ is not in a pistol with the j^{th} i. That is, the process is readily integrated into the procedure for finding Yamanouchi words.

2. For the reader that prefers permutations, we may describe $\mathrm{SYam}_\delta(\lambda)$ as follows. Consider the result of right justifying U_λ, and let A_λ be the set of pairs of values in cells that are touching on a southeasterly diagonal. Now treat $\pi \in \mathrm{SYam}(\lambda)$ as a row reading word of δ. Then $\pi \in \mathrm{SYam}_\delta(\lambda)$ if and only if no pairs in A_λ are in a pistol of $T_\delta(\pi)$. See Figure 3.5 for an example.

$$\begin{array}{|c|c|c|c|c|}\cline{1-2} 8&9&\multicolumn{3}{c}{}\\\cline{1-2} 6&7&\multicolumn{3}{c}{}\\\hline 1&2&3&4&5\\\hline\end{array} \qquad \begin{array}{|c|c|c|c|c|}\cline{4-5}\multicolumn{3}{c|}{}&8&9\\\cline{4-5}\multicolumn{3}{c|}{}&6&7\\\hline 1&2&3&4&5\\\hline\end{array} \qquad A_{(5,2,2)} = \{(6,5),(8,7)\} \qquad \begin{array}{|c|c|c|}\hline \mathbf{8}&6&9\\\hline \mathbf{7}&1&2\\\hline 3&4&5\\\hline\end{array}$$

Figure 3.5: From the left, $U_{(5,2,2)}$, followed by the result of right justifying, followed by $A_{(5,2,2)}$, followed by a standard filling T of partition shape $\mu = (3,3,3)$ such that $\mathrm{rw}(T) \in \mathrm{SYam}((5,2,2))$ and the pair $(8,7) \in A_{(5,2,2)}$ is in a pistol of μ. Thus $\mathrm{rw}(T) \notin \mathrm{SYam}_\mu((5,2,2))$.

3. The set of permutations in $\mathrm{SYam}(\lambda)$ is a Knuth class. The set $\mathrm{SYam}_\delta(\lambda)$ is necessarily a subset of this class, and so can be expressed via some set of recording tableaux of a given partition shape. Finding a more explicit way of generating all such recording tableaux remains an interesting open problem.

3.2.2 Symmetric Functions

We briefly recall the interplay between standard fillings, quasisymmetric functions, Schur functions, Macdonald polynomials, and Hall-Littlewood polynomials. All three have rich connections to the theory of symmetric functions. The curious reader may also refer to [Stanley, 2001, Ch. 7], [Fulton, 1997, Part I], or [Sagan, 1991, Ch. 4].

Definition 3.2.3. Given any signature $\sigma \in \{\pm 1\}^{n-1}$, define the *fundamental quasisymmetric function* $F_\sigma(X) \in \mathbb{Z}[x_1, x_2, \ldots]$ by

$$F_\sigma(X) := \sum_{\substack{i_1 \leq \ldots \leq i_n \\ i_j = i_{j+1} \Rightarrow \sigma_j = +1}} x_{i_1} \cdots x_{i_n}.$$

We may now use the previous definition to define the Schur functions, relying on a result of Ira Gessel. While it is not the standard definition, it is the most functional for our purposes.

Definition 3.2.4. [Gessel, 1984] Given any partition λ, define

$$s_\lambda := \sum_{T \in \mathrm{SYT}(\lambda)} F_{\sigma(T)}(X), \tag{3.2.2}$$

where s_λ *is a Schur function of shape* λ*.*

In order to define the modified Macdonald polynomials and the modified Hall-Littlewood polynomials, we will first need to define some statistics, relying on the results in [Haglund et al., 2005a] for our definitions. Let T be a filling of a diagram δ. Given a cell $c \in \delta$, let $T(c)$ denote the value of T in cell c. A descent of T is a cell c in δ such that $T(c) > T(d)$, where d is the cell directly below and adjacent to c. We denote the set of descents of T as $\mathrm{Des}(T)$. As an example, the filling on the left of Figure 3.2 has descents in the cells containing 3, 4, 6, 8, and 9, while the filling on the right has no descents. Given a cell c of δ, define $\mathrm{leg}(c)$ as the number of cells in δ strictly above and in the same column as c. Letting $w = \mathrm{rw}(T)$, we may then define

$$\mathrm{maj}_\delta(w) := \mathrm{maj}(T) := \mathrm{maj}_\delta(\mathrm{Des}(T)) := \sum_{c \in \mathrm{Des}(T)} 1 + \mathrm{leg}(c). \tag{3.2.3}$$

Let c, d, and e be cells of δ in row reading order. Then c, d, and e form a *triple* if c and d are in the same row and e is the cell immediately below c, as in Figure 3.6. If, in addition, T is a filling of the diagram δ, then c, d, and e form an *inversion triple* of T if $T(e) < T(d) < T(c)$, $T(c) \leq T(e) < T(d)$, or $T(d) < T(c) \leq T(e)$. As a mnemonic, in each set of inequalities, the three cells are presented in a counterclockwise order. If either c or e is not in δ, then the remaining two cells form an *inversion pair* of T if either $T(c) > T(d)$ or $T(d) > T(e)$. See Figure 3.6 for an example of each of these types of inversions. By letting $w = \mathrm{rw}(T)$, we may now define the final statistic as

$$\mathrm{inv}_\delta(w) := \mathrm{inv}(T) := |\{\text{inversion triples of } T\}| + |\{\text{inversion pairs of } T\}|. \tag{3.2.4}$$

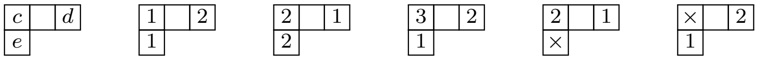

Figure 3.6: From left, a generic triple, three inversion triples, and then two inversion pairs, where × denotes the lack of a cell.

We are now able to define a combinatorial generalization of the modified Macdonald polynomials,

$$\widetilde{H}_\delta(X;q,t) := \sum_{T \in \mathrm{SF}(\delta)} q^{\mathrm{inv}(T)} t^{\mathrm{maj}(T)} F_{\sigma(T)}. \tag{3.2.5}$$

We similarly define the modified Hall-Littlewood polynomials as,

$$\widetilde{H}_\delta(X;0,t) := \widetilde{H}_\delta(X;0,t) = \sum_{\substack{T \in \mathrm{SF}(\delta) \\ \mathrm{inv}(T)=0}} t^{\mathrm{maj}(T)} F_{\sigma(T)}. \tag{3.2.6}$$

Further, if we let $\gamma \subset \delta$, we may define

$$R_{\gamma,\delta}(X) := \sum_{\substack{T \in \mathrm{SF}(\delta) \\ \mathrm{inv}(T)=0 \\ \mathrm{Des}(T)=\gamma}} F_{\sigma(T)}. \tag{3.2.7}$$

It follows immediately from these definitions that

$$\widetilde{H}_\delta(X;0,t) = \sum_{\gamma \subset \delta} t^{\mathrm{maj}_\delta(\gamma)} R_{\gamma,\delta}(X). \tag{3.2.8}$$

Using [Haglund et al., 2005a], it is possible to write $\widetilde{H}_\delta(X;q,t)$ as a sum of Lascoux-Leclerc-Thibon polynomials, which are shown to be Schur positive in [Grojnowski and Haiman, 2007]. Hence, $\widetilde{H}_\delta(X;q,t)$ is both symmetric and Schur positive. While the coefficients of the Schur expansion of $\widetilde{H}_\delta(X;0,t)$ and $R_{\gamma,\delta}(X)$ are given by Theorem 3.1.1, finding a combinatorial description of this expansion for $\widetilde{H}_\delta(X;q,t)$ remains an open problem.

We should emphasize that our definition of $\widetilde{H}_\delta(X;q,t)$ and $\widetilde{H}_\delta(X;0,t)$ are combinatorial generalizations, chosen to agree with definitions in [Haglund et al., 2005a] and related definitions of LLT polynomials. Hence, they need not agree with any algebraic generalizations of Macdonald polynomials. Specifically, Garsia and Haiman conjectured a generalization of Macdonald polynomials to diagram indices in [Garsia and Haiman, 1995], with further results contributed by Jason Bandlow in his Ph.D. dissertation [Bandlow, 2007]. Their conjecture would require that $\widetilde{H}_\delta(X;q,t) = \widetilde{H}_{\delta'}(X;t,q)$. While this is the case when δ is a partition, it fails for the diagram $\{(0,0),(1,1)\}$. Finding a way of recovering this symmetry, perhaps by modifying the maj statistic, is an important open problem.

Remark 3.2.5. In order to use Theorem 3.1.1 to expand $\widetilde{H}_\delta(X;0,t)$ in terms of Schur functions, it is necessary to generate $\{w \in \cup_{\lambda \vdash n} \mathrm{Yam}_\delta(\lambda) \colon \mathrm{inv}_\delta(w) = 0\}$. We may do this by making a tree: proceeding as mentioned in Part 1 of Remark 3.2.2 by filling μ in reverse row reading order and checking that there are no inversions, that we still have a Yamanouchi word, and that no pistol is jammed with the addition of each new entry. In fact, it is readily shown that upon filling the bottom three rows, such a filling must satisfy one of the three cases in Figure 3.7. That is, the bottom row must be all 1's, the second row starts with k

many 2's followed by all 1's, and then the third row has $j \leq k$ many 3's followed by one of three options. Either the rest of the third row is 1's, or there are $k - j$ 1's followed by all 2's, or the rest may be all 2's if the result is still a Yamanouchi word. It is, in theory, possible to precompute more rows in this fashion at the expense of more complicated rules.

$$\begin{array}{|c|c|c|c|c|c|c|}\hline 3&3&1&1&1&1&1\\\hline 2&2&2&2&1&1&1\\\hline 1&1&1&1&1&1&1\\\hline\end{array} \qquad \begin{array}{|c|c|c|c|c|c|c|}\hline 3&3&1&1&2&2&2\\\hline 2&2&2&2&1&1&1\\\hline 1&1&1&1&1&1&1\\\hline\end{array} \qquad \begin{array}{|c|c|c|c|c|c|c|}\hline 3&3&2&2&2&2&2\\\hline 2&2&2&2&1&1&1\\\hline 1&1&1&1&1&1&1\\\hline\end{array}$$

Figure 3.7: The three ways to fill the first three rows of μ when expanding $\widetilde{H}_\mu(X;0,t)$ into Schur functions.

It should be noted that the tree described above may still have dead ends. In that respect, a key open problem is to find an algorithm that avoids any dead ends in order to maximize efficiency. Such an algorithm was provided for the Littlewood-Richardson coefficients in [Remmel and Whitney, 1984], suggesting that it may be possible in this case as well.

3.2.3 Dual Equivalence Graphs

The key tool used in this paper is the theory of dual equivalence graphs. We quickly lay out the necessary background on the subject in this section.

Definition 3.2.6 (Haiman [1992])**.** Given a permutation in S_n expressed in one-line notation, define an *elementary dual equivalence* as an involution d_i that interchanges the values $i-1, i$, and $i+1$ as

$$\begin{aligned} d_i(\ldots i \ldots i-1 \ldots i+1 \ldots) &= (\ldots i+1 \ldots i-1 \ldots i \ldots), \\ d_i(\ldots i-1 \ldots i+1 \ldots i \ldots) &= (\ldots i \ldots i+1 \ldots i-1 \ldots), \end{aligned} \tag{3.2.9}$$

and that acts as the identity if i occurs between $i-1$ and $i+1$. Two permutations are *dual equivalent* if one may be transformed into the other by successive elementary dual equivalences.

For example, 21345 is dual equivalent to 41235 because $d_3(d_2(21345)) = d_3(31245) = 41235$.

We may also let d_i act on the entries of a standard Young tableau by applying them to the row reading word. It is not hard to check that the result of applying this action to a standard Young tableau is again a standard Young tableau. The transitivity of this action is described in the following theorem.

Theorem 3.2.7 ([Haiman, 1992, Prop. 2.4]). *Two standard Young tableaux on partition shapes are dual equivalent if and only if they have the same partition shape.*

It follows from [Haiman, 1992, Lem. 2.3] that the action of d_i is further described by

$$d_i(P(\pi)) = P(d_i(\pi)). \tag{3.2.10}$$

By definition, d_i is an involution, and so we define a graph on standard Young tableaux by letting each nontrivial orbit of d_i define an edge colored by i. By Theorem 3.2.7, the graph on $\mathrm{SYT}(n)$ with edges colored by $1 < i < n$ has connected components with vertices in $\mathrm{SYT}(\lambda)$ for each $\lambda \vdash n$. We may further label each vertex with its signature to create a *standard dual equivalence graph* that we will denote $\mathcal{G}_\lambda$. Refer to Figure 3.8 for examples of $\mathcal{G}_\lambda$ with $\lambda \vdash 5$.

Definition 3.2.4 and Theorem 3.2.7 determine the connection between Schur functions and dual equivalence graphs as highlighted in [Assaf, 2013, Cor. 3.10]. Given any standard dual equivalence graph $\mathcal{G}_\lambda = (V, \sigma, E)$,

$$\sum_{v \in V} F_{\sigma(v)} = s_\lambda. \tag{3.2.11}$$

Here, $\mathcal{G}_\lambda$ is an example of the following broader class of graphs.

Definition 3.2.8. A *signed colored graph* consists of the following data:

1. *a finite vertex set* V*,*
2. *a signature function* $\sigma : V \to \{\pm 1\}^{n-1}$ *for some fixed positive integer* n*,*
3. *a collection* E_i *of unordered pairs of distinct vertices in* V *for each* $i \in \{2, \ldots, n-1\}$ *and the same positive integer* n*.*

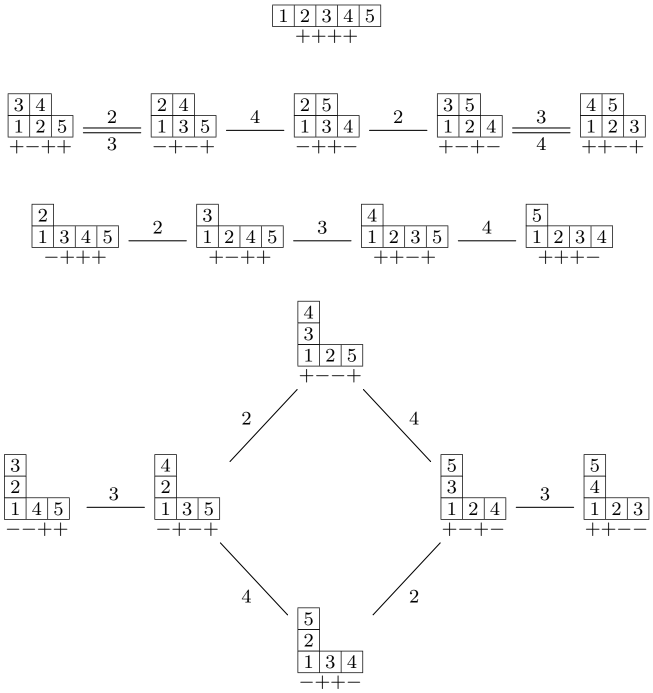

Figure 3.8: Some dual equivalence graphs with standard Young tableaux as vertices.

We denote a signed colored graph by $\mathcal{G} = (V, \sigma, E_2 \cup \cdots \cup E_{n-1})$ or simply $\mathcal{G} = (V, \sigma, E)$.

In order to give an abstract definition of dual equivalence graphs, we will need definitions for isomorphisms and restrictions. Given two signed colored graphs $\mathcal{G}(V, \sigma, E)$ and $\mathcal{H}(V', \sigma', E')$, an *isomorphism* $\phi\colon \mathcal{G} \to \mathcal{H}$ is a bijective map from V to V' such that both ϕ and ϕ^{-1} preserve colored edges and signatures. The definition of a restriction is a bit more technical.

Definition 3.2.9. Given a signed colored graph $\mathcal{G} = (V, \sigma, E)$ and an interval of nonnegative integers I, define the restriction of $\mathcal{G}$ to I, denoted $\mathcal{G}|_I$, as the signed colored graph $\mathcal{H} = (V, \sigma', E')$, where

1. $\sigma'(v)_i = \sigma(v)_{\min(I)+i-1}$ when $i \in \{1, \ldots, |I| - 1\}$ and $\sigma_{\min(I)+i-1}$ is defined.

2. $E'_i = E_{\min(I)+i-1}$ when $i \in \{2, 3, \ldots, |I| - 1\}$ and $E_{\min(I)+i-1}$ is defined.

We now proceed to the definition of a dual equivalence graph. Here, we use a result of [Roberts, 2014] as our definition. For more general definitions, see [Assaf, 2013] and [Roberts, 2014].

Definition 3.2.10. A signed colored graph $\mathcal{G} = (V, \sigma, E)$ is a *dual equivalence graph* if the following two properties hold.

Locally Standard Property: If I is any interval of integers with $|I| = 6$, then each component of $\mathcal{G}|_I$ is isomorphic to some $\mathcal{G}_\lambda$.

Commuting Property: If $\{v, w\} \in E_i$ and $\{w, x\} \in E_j$ for some $|i-j| > 2$, then there exists $y \in V$ such that $\{v, y\} \in E_j$ and $\{x, y\} \in E_i$.

Theorem 3.2.11 ([Assaf, 2013, Thm 3.9])**.** *A connected component of a signed colored graph is a dual equivalence graph if and only if it is isomorphic to a unique $\mathcal{G}_\lambda$.*

Next, we will associate to every Macdonald polynomial and Hall-Littlewood polynomial a signed colored graph. To do this, we need to define an involution D_i^δ to provide the edge sets of a signed colored graph, as defined originally in [Assaf, 2013]. First let $\tilde{d}_i : S_n \to S_n$ be the involution that cyclically permutes the values $i-1, i$, and $i+1$ as

$$\begin{aligned} \tilde{d}_i(\ldots i \ldots i-1 \ldots i+1 \ldots) &= (\ldots i-1 \ldots i+1 \ldots i \ldots), \\ \tilde{d}_i(\ldots i \ldots i+1 \ldots i-1 \ldots) &= (\ldots i+1 \ldots i-1 \ldots i \ldots), \end{aligned} \tag{3.2.12}$$

and that acts as the identity if i occurs between $i-1$ and $i+1$. For example, $\tilde{d}_3 \circ \tilde{d}_2(4123) = \tilde{d}_3(4123) = 3142$.

We now define the desired involution. Given $\pi \in S_n$ and a diagram δ of size n,

$$D_i^\delta(\pi) := \begin{cases} \tilde{d}_i(\pi) & \text{if the indices of } i-1,\ i, \text{ and } i+1 \text{ in } \pi \text{ are } \delta\text{-pistoled}, \\ d_i(\pi) & \text{otherwise.} \end{cases} \tag{3.2.13}$$

As an example, we may take $\pi = 53482617$ and δ as in Figure 3.9. Then $D_3^\delta(\pi) = \tilde{d}_3(\pi) = 54283617$ and $D_5^\delta(\pi) = d_5(\pi) = 63482517$.

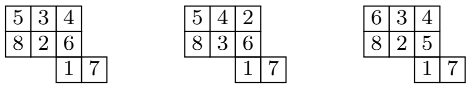

Figure 3.9: Three standard fillings of a diagram δ. At left, a standard filling with row word $\pi = 53482617$ followed by $T_\delta(D_3^\delta(\pi))$ and then $T_\delta(D_5^\delta(\pi))$.

Given some δ of size n, we may then define an *Assaf Graph* as the signed colored graph $\mathcal{H}_\delta = (V, \sigma, E)$ with vertex set $V = S_n$, signature function σ given via inverse descents, and edge sets E_i defined via the nontrivial orbits of D_i^δ. It is readily shown that the action of D_i^δ on π preserves $\mathrm{inv}_\delta(\pi)$, $\mathrm{Des}_\delta(\pi)$, and $\mathrm{maj}_\delta(\pi)$. Thus, these functions are all constant on components of $\mathcal{H}_\delta$. We may study $\widetilde{H}_\delta(X; 0, t)$ by restricting our attention to components of $\mathcal{H}_\delta$ where inv_δ is zero, as in the following definition.

Definition 3.2.12. Let γ and δ be diagrams such that $\gamma \subset \delta$ and $|\delta| = n$. We define two subgraphs of $\mathcal{H}_\delta = (V, \sigma, E)$ induced by restricting the vertex set V as follows.

1. $\mathcal{P}_\delta := (V', \sigma, E')$ is the subgraph with $V' = \{\pi \in S_n \colon \mathrm{inv}_\delta(\pi) = 0\}$.
2. $\mathcal{R}_{\gamma,\delta} := (V'', \sigma, E'')$ is the subgraph with $V'' = \{\pi \in S_n \colon \mathrm{inv}_\delta(\pi) = 0, \mathrm{Des}_\delta(\pi) = \gamma\}$.

Notice that each subgraph is a union of connected components of $\mathcal{H}_\delta$. Furthermore,

$$\begin{aligned}
\widetilde{H}_\delta(X; q, t) &= \sum_{v \in V} q^{\mathrm{inv}_\delta(v)} t^{\mathrm{maj}_\delta(v)} F_{\sigma(v)} \\
\widetilde{H}_\delta(X; 0, t) &= \sum_{v \in V'} t^{\mathrm{maj}_\delta(v)} F_{\sigma(v)}, \qquad (3.2.14) \\
R_{\gamma,\delta}(X) &= \sum_{v \in V''} F_{\sigma(v)}.
\end{aligned}$$

The Assaf graph $\mathcal{H}_\delta$ is the primary object of interest in the proof of the following lemma, which was originally stated in terms of Lascoux-Leclerc-Thibon polynomials but is easily translated using results in [Haglund et al., 2005a].

Lemma 3.2.13 ([Roberts, 2014, Lem. 4.13]). *Let δ be a diagram such that no pistol of δ contains more than three cells, then*

$$\widetilde{H}_\delta(X;q,t) = \sum_{\lambda \vdash |\delta|} \sum_{w \in \mathrm{Yam}(\lambda)} q^{\mathrm{inv}_\delta(w)} t^{\mathrm{maj}_\delta(w)} s_\lambda.$$

3.3 Dual Equivalence Graphs in $\mathcal{H}_\delta$

3.3.1 The Graphs $\mathcal{P}_\delta$ and $\mathcal{R}_{\gamma,\delta}$

Next we give a classification for when a component of $\mathcal{H}_\delta$ is a dual equivalence graph. To do this, we will need the following definition. Given permutations $p \in S_m$ and $\pi \in S_n$ with $m \leq n$, we say that p is a *strict pattern* of π if there exists some sequence $i_1 < i_2 < \ldots < i_m$ such that $\pi_{i_j} = p_j + k$ for some fixed integer k and all $1 \leq j \leq m$. Furthermore, we say that p is a δ-strict pattern of π if the indices $i_1, i_2, \ldots, i_m$ of π are δ-pistoled. As an example, on the left side of Figure 3.9, $p = 231$ is a strict pattern of $\pi = 534826179$ occurring in π at indices 2, 3, and 5. These indices correspond to the second, third, and fifth cell in row reading order of the given diagram δ. These cells are contained in a pistol δ. Thus, $p = 231$ is a δ-strict pattern of π.

Lemma 3.3.1. *Let $\mathcal{G} = (V, \sigma, E)$ be a component of $\mathcal{H}_\delta$. Then $\mathcal{G}$ is a dual equivalence graph if and only if for every vertex $\pi \in V$, the following two conditions holds.*

1. *The permutations 1342 and 2431 are not δ-strict patterns of π.*
2. *If the permutation 12543 or 34521 is a strict pattern of π occurring in indices $i_1, \ldots, i_5$, then $\{i_1, \ldots, i_5\}$ is δ-pistoled, $\{i_1, \ldots, i_4\}$ is not δ-pistoled, or $\{i_2, \ldots, i_5\}$ is not δ-pistoled.*

Proof. If $\mathcal{G}$ contains a vertex with one of the patterns described above, then we may consider the smallest interval I such that restricting to $\pi|_I$ gives a word that still contains one of the strict patterns. Direct inspection shows that the components containing those patterns do not satisfy the Locally Standard Property in Definition 3.2.10.

If, on the other hand, $\mathcal{G}$ does not contain a vertex with one of the above patterns, we need only show that $\mathcal{G}$ satisfies Definition 3.2.10. The proof of the Commuting Property follows from the fact that D_i^δ fixes all values except $i-1, i$, and $i+1$. To demonstrate the Locally Standard Property, it suffices to check $\mathcal{H}_\delta$ with $|\delta| = 6$. By considering the finitely many ways that the indices of π can be put into pistols, this property can be verified via computer. More specifically, there are 6! permutations in S_6, and it can be shown that the number of possible ways to give subsets of [6] corresponding to pistols of a diagram of size 6 is equal to the 6^{th} Catalan number, 132, providing an upper bound of $6! \cdot 132 = 95040$ cases to check. The code for said verification can be found at <http://www.math.washington.edu/~austinis/Proof_DEG_by_strict_patterns.pdf>, cited as [Roberts, 2013b]. □

Theorem 3.3.2. *If γ and δ are any diagrams such that $\gamma \subset \delta$, then $\mathcal{R}_{\gamma,\delta}$ and $\mathcal{P}_\delta$ are dual equivalence graphs.*

Proof. Let π be some arbitrary vertex of $\mathcal{R}_{\gamma,\delta}$ or $\mathcal{P}_\delta$. It suffices to show that if the hypotheses of Lemma 3.3.1 are not satisfied by a permutation π and diagram δ, then $\mathrm{inv}_\delta(\pi) \neq 0$. Let $T = T_\delta(\pi)$ such that π contains one of the strict patterns mentioned in Lemma 3.3.1. Let $\pi|_I$ be this pattern. We will use the location of $\pi|_I$ in T to show that $\mathrm{inv}_\delta(\pi) \geq 1$.

First notice that as a word, $\pi|_I$ ends in a descent. If the last two values of $\pi|_I$ occur in the same row of T, they must be part of an inversion triple or an inversion pair, since the cell completing said triple cannot have a value in between the last two values in $\pi|_I$, by the definition of a strict pattern. We may thus restrict our attention to the case where the last value of $\pi|_I$ does not share its row in T with any other value of I.

We will now demonstrate an inversion triple or inversion pair by ignoring any rows and columns that do not contain the last four values of $\pi|_I$. By the assumption that the last four values of $\pi|_I$ are contained in a pistol, we have restricted to a diagram with exactly two rows. We may then demonstrate an inversion triple or a inversion pair in all possible cases, as shown in Figure 3.10. Thus $\mathrm{inv}_\mu(\pi) \geq 1$, completing our proof. □

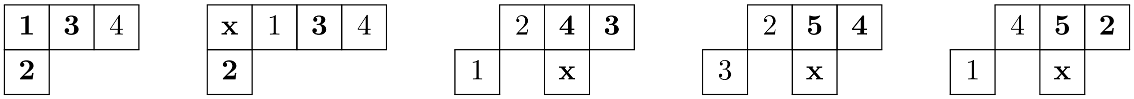

Figure 3.10: Configurations for the last four values of $\pi|_I$ with inversion triples in bold. Here, numbers may be shifted by some nonnegative integer k, and the variable x may represent an omitted cell or any value not in I.

3.3.2 The Proofs of Theorems 3.1.1 and 3.1.2

We begin by giving several lemmas necessary for the proof of Theorem 3.1.2. The results driven reader may safely skip these lemmas without hindering their understanding of later sections. In the proof of Lemmas 3.3.3–3.3.5, we will use the variables s, t, u, and v to denote vertices of graphs. Moreover, the vertices of graphs in the proofs of these lemmas will always be permutations.

Lemma 3.3.3. *Let γ be a diagram such that $|\gamma| = n$, and let $\lambda \vdash n$ have at most two rows. Let $\pi \in \mathrm{SYam}_\gamma(\lambda)$ and $u \in S_{n-1}$ such that $u = \pi|_{[n-1]}$. Let δ be the diagram of $T_\gamma(\pi)\big|_{[n-1]}$. Then u is connected by a path p in $\mathcal{H}_\delta$ to some vertex v such that:*

1. *each edge of p is defined via the action of d_i,*

2. *$v_{n-1} = n - 1$,*

3. *$u_i = v_i$ whenever $u_i < \lambda_1$.*

Furthermore, the sequence of edge colors in p is not dependent on the choice of $\pi \in \mathrm{SYam}_\delta(\lambda)$.

Proof. We proceed by induction on the size of λ_2. If $\lambda_2 \leq 1$, then u must be the identity permutation $123 \cdots n-1$, since $\pi \in \mathrm{SYam}(\lambda)$. We may then let $u = v$ in order to satisfy the result. For the inductive step, suppose that $\lambda_2 > 1$ and that the result holds for all two row partitions whose second row is smaller than λ_2.

We will apply a sequence of edges to find v, as portrayed in Figure 3.11. Here, it may be helpful to notice that $P(\pi) = U_\lambda$, by the definition of $\mathrm{SYam}(\lambda)$, and that $P(u)$ is the result of removing n from U_λ. It may also be helpful to recall that the action of d_i on $P(u)$ can be understood via (3.2.10). It follows from the fact that $\pi \in \mathrm{SYam}(\lambda)$ and $\lambda_2 \geq 2$ that $u_{n-2} = \lambda_1 - 1$ and $u_{n-1} = \lambda_1$. To apply the inductive hypotheses, notice that restricting u to values in $[n-2]$ yields a permutation equal to some $\pi'|_{[n-2]}$ where $\pi' \in \mathrm{SYam}(\mu)$ and $\mu = (\lambda_1, \lambda_2 - 1)$. Specifically, $\pi' = u$. We may then apply induction to move $n-2$ into position $n-1$ of u. Similarly restricting to values in $[n-3]$, we may then move $n-3$ into position $n-2$ via some path q.

We have not changed the index of $n-1$, and so $n-1$ must now occur before $n-3$, which occurs before the $n-2$ in the last index. We may thus apply an $n-2$-edge to move $n-1$ into the last index. Since the last index of π does not jam a pistol of γ, the last index of u and the index of $n-1$ in u cannot be δ-pistoled. In particular, the $n-2$-edge must be defined via d_i. Finally, we may consider the restriction to values in $[n-3]$ again to apply the edges of q in reverse order, ensuring $u_i = v_i$ whenever $u_i < \lambda_1$. The result is the desired v, as given by applying a sequence of edges that was not dependent on the choice of $\pi \in \mathrm{SYam}(\lambda)$. ☐

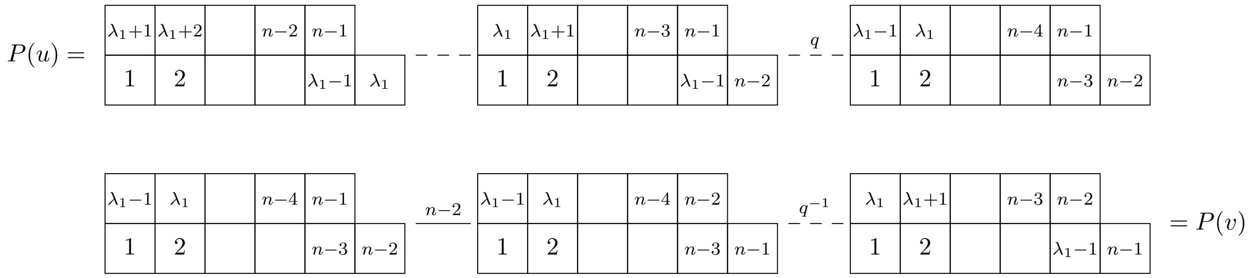

Figure 3.11: The sequence of edges connecting $P(u)$ and $P(v)$.

Lemma 3.3.4. *Let $\mathcal{G}$ and $\mathcal{H}$ be connected components of $\mathcal{H}_\gamma$ and $\mathcal{H}_\delta$, respectively. Further suppose that there exists an isomorphism $\phi\colon \mathcal{G} \to \mathcal{H}$. Let s be a vertex of $\mathcal{G}$ such that $s \in \mathrm{SYam}_\gamma(\lambda)$, $\lambda \vdash n$, and $\phi(s)|_{[n-1]} \in \mathrm{SYam}(\lambda)|_{[n-1]}$. Then $\phi(s)$ does not jam δ.*

Proof. Let s be as in the statement of the lemma. We begin with the case where λ has at most two rows, and proceed by induction on the size of λ_2. If $\lambda_2 \leq 1$, then it follows from the hypothesis that $s \in \mathrm{SYam}(\lambda)$ and the fact that the isomorphism ϕ preserves signatures that $\mathrm{unst}(\phi(s))$ has at most one 2. Hence, $\phi(s)$ cannot jam δ, by Definition 3.2.1.

Next assume that λ has exactly two rows and that $\lambda_2 \geq 2$. Applying induction, we further assume the result for all two row partitions whose second row has fewer than λ_2 cells. We wish to show that $\phi(s)$ does not jam δ. Because λ has two rows, all values of s weakly less than λ_1 are taken to 1 in $\mathrm{unst}(\phi(s))$ and all values greater than λ_1 are take to 2 in $\mathrm{unst}(\phi(s))$. Applying Definition 3.2.1, it thus suffices to show that the indices of $\phi(s)$ with values weakly less than λ_1 do not jam δ.

It follows from (3.2.13) that a permutation π is only contained in an $i/i+1$-double edge of $\mathcal{H}_\delta$ if the indices of i and $i+1$ in π are not δ-pistoled. We will use this fact to show that the index of λ_1 in $\phi(s)$ does not jam δ by demonstrating a particular double edge in $\mathcal{G}$, as shown in Figure 3.12. By considering the restriction to values in $[n-2]$, we may apply Lemma 3.3.3 to move $n-2$ into the index of λ_1 in s. Call this vertex t. Next we may consider the restriction to values in $[n-3]$ to similarly move $n-3$ into the index of $\lambda_1 - 1$ in s. Call the resulting vertex u. Notice that s is connected to u by a sequence of edges defined via the action of d_i and whose labels are less than $n-2$. In particular, $n-1$ and n are not moved. Thus the indices of $n-2$ and $n-1$ in u are not γ-pistoled, since the index of λ_1 of s does not jam γ. Furthermore, n and $n-3$ are between $n-2$ and $n-1$ in u, so u must admit an $n-2/n-1$-double edge.

Now consider the effect of the same sequence of edges on $\phi(s)$. Notice that $\phi(s)|_{[n-2]}$ also satisfies the hypotheses of Lemma 3.3.3. Thus, $\phi(s)$ is connected to $\phi(u)$ by a sequence of edges that are defined via d_i and do not move $n-1$ or n. Because isomorphisms preserve edges, $\phi(u)$ must also admit an $n-2/n-1$-double edge. In particular, the indices of $n-2$ and $n-1$ in $\phi(u)$ are not δ-pistoled. Thus, the indices of λ_1 and $n-1$ in s are not δ-pistoled.

We still need to show that no index of $\phi(s)$ with value less than λ_1 can jam δ. Let v be the vertex connected to t by an $n-1$-edge, as in Figure 3.12. Applying the above analysis, s is connected to v by a sequence of edges that are all defined via d_i, and similarly for $\phi(s)$ and $\phi(v)$. Notice that $v|_{[n-2]} \in \mathrm{SYam}((\lambda_1 - 1, \lambda_2 - 1))$ and that $v|_{[n-2]}$ does not jam γ', where γ' is the diagram of $T_\gamma(v)|_{[n-2]}$. By our inductive hypothesis, $\phi(v)|_{[n-2]}$ does not jam δ', where δ' is the diagram of $T_\delta(\phi(v))|_{[n-2]}$. Comparing $\phi(s)$ to $\phi(v)$, it follows that no

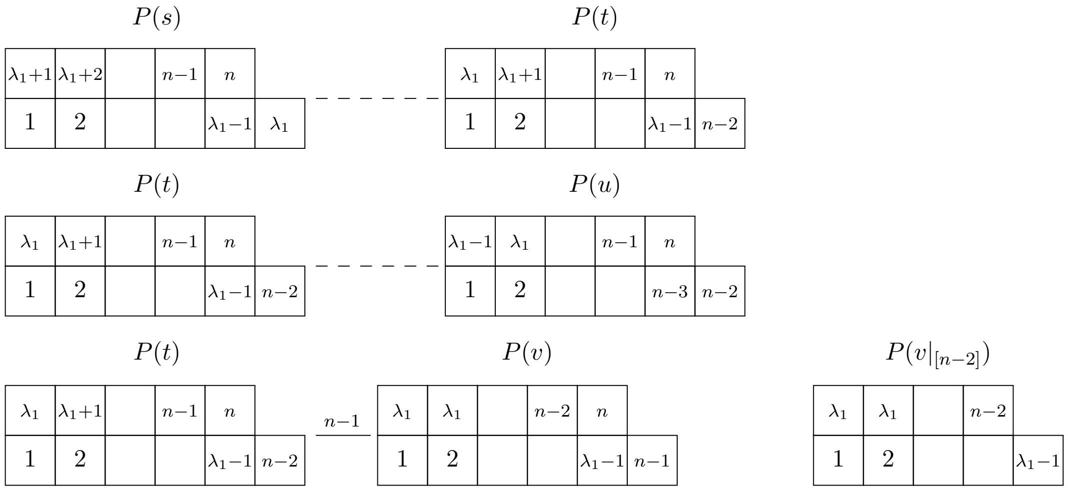

Figure 3.12: The relationships between the standard Young tableaux $P(s)$, $P(t)$, $P(u)$, $P(v)$, and $P(v|_I)$.

index of $\phi(s)$ with value less than λ_1 can jam δ. Hence, $\phi(s)$ does not jam δ. We have thus completed our inductive argument for λ with at most two rows.

Finally, consider the case where λ has more than two rows. Suppose, for the sake of contradiction, that $\phi(s)$ jams δ. In particular, there are some values i and $i+1$ in $\operatorname{unst}(s)$ that satisfy the definition of jamming. We wish to restrict our attention to those values in s that are sent to i and $i+1$ in $\operatorname{unst}(s)$ and then force a contradiction by applying the two row case above. Specifically, there must exist some interval of integers I such that $\operatorname{st}(s|_I) \in \operatorname{SYam}(\mu)$, where μ has exactly two rows, and $\operatorname{st}(\phi(s)|_I)$ jams δ', where δ' is the diagram of $T_\delta(\phi(s))|_I$. If we let γ' be the diagram of $T_\gamma(s)|_I$, then the component of $\operatorname{st}(s|_I)$ in $\mathcal{H}_{\gamma'}$ is isomorphic to the component of s in $\mathcal{H}_\gamma|_I$. Similarly, the component of $\operatorname{st}(\phi(s)|_I)$ in $\mathcal{H}_{\delta'}$ is isomorphic to the component of $\phi(s)$ in $\mathcal{H}_\delta|_I$. Hence, there exists an isomorphism from the component of $\operatorname{st}(s|_I)$ in $\mathcal{H}_{\gamma'}$ to the component of $\operatorname{st}(\phi(s)|_I)$ in $\mathcal{H}_{\delta'}$. Appealing to the two row case provides the desired contradiction. Thus $\phi(s)$ does not jam δ. □

Lemma 3.3.5. *Let $\mathcal{G}$ and $\mathcal{H}$ be connected components of $\mathcal{H}_\gamma$ and $\mathcal{H}_\delta$, respectively. Further suppose that there exists an isomorphism $\phi\colon \mathcal{G} \to \mathcal{H}$. If s is a vertex of $\mathcal{G}$ such that*

$s \in \mathrm{SYam}_\gamma(\lambda)$, then $\phi(s) \in \mathrm{SYam}_\delta(\lambda)$.

Proof. Let s be as in the statement of the lemma. We begin with the case where λ has at most two rows and proceed by induction on the size of λ_2. If $\lambda_2 \leq 1$, then the fact that isomorphisms preserve signatures guarantees that $\mathrm{unst}(\phi(s))$ has at most one 2. Hence, $\phi(s)$ cannot jam δ by Definition 3.2.1. Also via signature considerations, it follows that $\phi(s) \in \mathrm{SYam}(\lambda)$.

Now assume that λ has exactly two rows and $\lambda_2 \geq 2$. Applying our induction, we further assume the result for all two row partitions whose second row has fewer than λ_2 cells. Applying Lemma 3.3.4, $\phi(s)$ does not jam δ. To complete the argument, we need to show that $\phi(s) \in \mathrm{SYam}(s)$. By induction, $\phi(s)|_{[n-1]} \in \mathrm{SYam}(\lambda)|_{[n-1]}$, so it suffices to show that n is not in the first row of $P(\phi(s))$.

To show that n is not in the first row of $P(\phi(s))$, consider the vertex v as defined in the proof of Lemma 3.3.4 and depicted in Figure 3.12. Recall that s is connected to v by a sequence of edges defined via d_i. Similarly, $\phi(s)$ is connected to $\phi(v)$ by a sequence of edges defined via d_i. Also recall that $v|_{[n-2]} \in \mathrm{SYam}((\lambda_1 - 1, \lambda_2 - 1))$. Thus, $\phi(v)|_{[n-2]} \in \mathrm{SYam}((\lambda_1 - 1, \lambda_2 - 1))$, by our inductive hypothesis. Because $\phi(s)$ is connected to $\phi(v)$ by a sequence of edges defined via d_i, we may use (3.2.10) to consider the relationship between $P(\phi(s))$ and $P(\phi(v))$. It can then be shown that if n is in the bottom row of $P(\phi(s))$, then $n-2$ is in the bottom row of $\phi(v)$, as shown in Figure 3.13. In particular, $\phi(v)|_{[n-2]} \notin \mathrm{SYam}((\lambda_1 - 1, \lambda_2 - 1))$. This contradiction completes our inductive argument for λ with at most two rows.

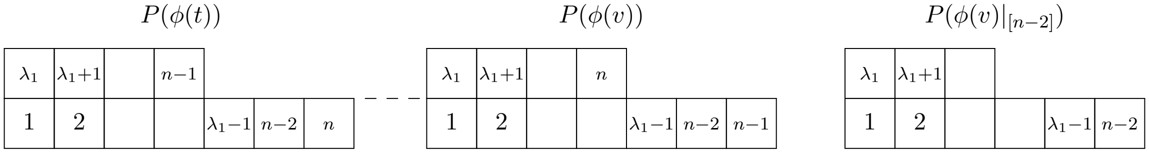

Figure 3.13: The standard Young tableaux $P(\phi(t)), P(\phi(v))$, and $P(\phi(v)|_{[n-2]})$ if n were in the bottom row of $P(s)$.

Finally, consider the case where λ has more than two rows. By inducting on $|\lambda|$, we may apply Lemma 3.3.4 again to show that $\phi(s)$ does not jam γ. Now suppose, for the sake of contradiction, that $\phi(s)$ is not a standardized Yamanouchi word. Then there must be some minimal integer r such that there is a value in row r of $P(s)$ occurring in a lower row of $P(\phi(s))$. We may let I be the set of values in row $r-1$ and r of $P(s)$. Thus, $\mathrm{st}(s|_I) \in \mathrm{SYam}(\mu)$ for some partition μ with exactly two rows, and $\mathrm{st}(\phi(s)|_I)$ is not a standardized Yamanouchi word. As in the end of the proof of Lemma 3.3.4, there must exist an isomorphism from the component of $\mathrm{st}(s|_I)$ in some $\mathcal{H}_{\gamma'}$ to the component of $\mathrm{st}(\phi(s)|_I)$ in some $\mathcal{H}_{\delta'}$. However, this contradicts the conclusion of the two row case above. Thus $\phi(s) \in \mathrm{SYam}_\delta(\lambda)$. □

Proof of Theorem 3.1.2. Recall that δ is a diagram and that $\mathcal{G} = (V, \sigma, E)$ is a dual equivalence graph such that $\mathcal{G}$ is a component of $\mathcal{H}_\delta$. By Theorem 3.2.11, we may further suppose that $\mathcal{G} \cong \mathcal{G}_\lambda$, where $\lambda \vdash n$. We need to show that $|V \cap \mathrm{SYam}_\delta(\lambda)| = 1$ and that $|V \cap \mathrm{SYam}_\delta(\mu)| = 0$ if $\lambda \neq \mu$.

In order to apply Lemma 3.3.5, we first need to show that $\mathcal{G}_\lambda$ is isomorphic to a component of a well behaved $\mathcal{H}_\gamma$. By conflating standard Young tableaux in $\mathcal{G}_\lambda$ with their row reading words, we may consider $\mathcal{G}_\lambda$ as a component $\mathcal{C}$ of $\mathcal{H}_\gamma$, where γ is the subset of the vertical-axis $\{(0,i) \in \mathbb{Z} \times \mathbb{Z} \colon 0 \leq i < n\}$. Notice that the unique standardized Yamanouchi word in $\mathcal{C}$ is $\mathrm{rw}(U_\lambda)$ and that $\mathrm{SYam}_\gamma(\lambda) = \mathrm{SYam}(\lambda)$. By composing isomorphisms, we may find an isomorphism $\phi \colon \mathcal{G} \to \mathcal{C}$. By Lemma 3.3.5, ϕ acts as a bijection between the set of standardized Yamanouchi words in $\mathcal{C}$ that do not jam γ and the set of standardized Yamanouchi words in $\mathcal{G}$ that do not jam δ. Thus, there is a unique permutation in $V \cap \mathrm{SYam}_\delta(\lambda)$ corresponding to $\phi^{-1}(\mathrm{rw}(U_\lambda))$. Hence, $|V \cap \mathrm{SYam}_\delta(\lambda)| = 1$. Since $\mathrm{SYam}(\lambda)$ and $\mathrm{SYam}(\mu)$ are disjoint when $\lambda \neq \mu$, it follows that $V \cap \mathrm{SYam}_\delta(\mu) = \emptyset$ when $\lambda \neq \mu$. □

Corollary 3.3.6. *If $\mathcal{G} = (V, \sigma, E)$ is a dual equivalence graph contained in $\mathcal{H}_\delta$, then*

$$\sum_{v \in V} F_{\sigma(v)}(X) = \sum_{\lambda \vdash n} |V \cap \mathrm{SYam}_\delta(\lambda)| \cdot s_\lambda \tag{3.3.1}$$

Proof. This is an immediate consequence of Theorem 3.1.2 and (3.2.11). □

Proof of Theorem 3.1.1. Expressing $R_{\gamma,\delta}(X)$ and $\widetilde{H}_\delta(X;0,t)$ as in (3.2.14), the result follows from Theorems 3.3.2, Corollary 3.3.6, and the fact that the maj_δ statistic is constant on components of $\mathcal{P}_\delta$. □

Remark 3.3.7.

1. For each partition λ and diagram δ, there exists a set $Y_\delta(\lambda)$ defined as the intersection of $\mathrm{SYam}_\delta(\lambda)$ with the set of permutations of length $|\delta|$ whose component in $\mathcal{H}_\delta$ is a dual equivalence graph. That is, if $\mathcal{G} = (V, \sigma, E)$ is a component of $\mathcal{H}_\delta$, then $|V \cap Y_\delta(\lambda)| = 1$ if $\mathcal{G} \cong \mathcal{G}_\lambda$, and $|V \cap S_\delta(\lambda)| = 0$ otherwise. Finding a more direct way to generate Y_δ, however, is an open problem.

2. We may use Theorem 3.3.2 to find two related families of dual equivalence graphs in $\mathcal{H}_\delta$. Consider the graphs obtained by inverting the vertex sets of $\mathcal{P}_\delta$ and $\mathcal{R}_{\gamma,\delta}$, sending i to $|\delta| + 1 - i$. If we similarly invert the edge labels and multiply the signatures by -1, this necessarily creates dual equivalence graphs in $\mathcal{H}_\delta$, and so we may apply Corollary 3.3.6. Combinatorially, we are restricting our attention to permutations π that achieve the maximal $\mathrm{inv}_\delta(\pi)$, denoted $m(\delta)$. The associated symmetric function may also be computed by applying the function ω, which sends s_λ to $s_{\lambda'}$, and replacing maj_δ with comaj_δ, which is the result of subtracting maj from the maximal value of $\mathrm{maj}_\delta(\pi)$. That is,

$$\widetilde{H}_\delta(X;q,t)\big|_{q^{m(\delta)}} = \sum_{\lambda \vdash |\delta|} \sum_{\substack{w \in \mathrm{Yam}_\delta(\lambda) \\ \mathrm{inv}_\delta(w) = m(\delta)}} t^{\mathrm{maj}_\delta(w)} s_\lambda = \sum_{\lambda \vdash |\delta|} \sum_{\substack{w \in \mathrm{Yam}_\delta(\lambda) \\ \mathrm{inv}_\delta(w) = 0}} t^{\mathrm{comaj}_\delta(w)} s_{\lambda'}.$$

3. At this point it is appropriate to briefly discuss a connection to the Lascoux-Leclerc-Thibon (LLT) polynomials. The involution D_i^δ was originally defined in [Assaf, 2013] in order to assign a signed colored graph to the LLT polynomial $G_{\boldsymbol{\nu}}(X;q)$, which is the set of symmetric functions generated as sums over $\mathrm{SYT}(\boldsymbol{\nu})$, the set of standard Young fillings of the tuples of skew tableaux $\boldsymbol{\nu} = (\nu^{(0)}, \ldots, \nu^{(k-1)})$. In this case, we use the involution $D_i^{\boldsymbol{\nu}}$, and the resulting signed colored graphs are called *LLT graphs*. In $G_{\boldsymbol{\nu}}(X;q)$, the exponent of q is defined by the inv statistic on $\mathbf{T} \in \mathrm{SYT}(\boldsymbol{\nu})$, which

is closely related to our earlier definition of inv_δ. The action of $D_i^{\boldsymbol{\nu}}$ on $\mathbf{T} \in \mathrm{SYT}(\boldsymbol{\nu})$ preserves $\mathrm{inv}(\mathbf{T})$, where this action is defined via a reading word $\mathrm{rw}(\mathbf{T})$. For the moment, assume that $\boldsymbol{\nu}$ is a tuple of ribbon tableaux. Using the relationship between Macdonald polynomials and LLT polynomials described in [Haglund et al., 2005a], we may send each $\nu^{(j)}$ to a column of some δ and then treat the graph associated to $G_{\boldsymbol{\nu}}(X;0)$ as a collection of components of some $\mathcal{R}_{\gamma,\delta}$. In particular, the graph associated to $G_{\boldsymbol{\nu}}(X;0)$ is a dual equivalence graph, by Theorem 3.3.2. As in Theorem 3.1.1, we may further conclude that

$$G_{\boldsymbol{\nu}}(X;0) = \sum_{\lambda \vdash |\boldsymbol{\nu}|} \sum_{\substack{\mathbf{T} \in \mathrm{SYT}(\boldsymbol{\nu}) \\ \mathrm{rw}(\mathbf{T}) \in \mathrm{Yam}_\delta(\lambda) \\ \mathrm{inv}(\mathbf{T}) = 0}} s_\lambda$$

For the more general case where $\boldsymbol{\nu}$ is a tuple of skew shapes, we need only associate each skew tableau in $\boldsymbol{\nu}$ with the union of separate ribbons given by its columns, providing some $\widetilde{\boldsymbol{\nu}}$ that is a tuple of ribbons through which to gain the analogous result.

3.3.3 Related Conjectures

The above analysis lends itself to an interesting conjecture about the Schur expansion of the quasisymmetric function associated to any graph comprised of components of $\mathcal{H}_\delta$.

Conjecture 3.3.8. *Given any diagram δ such that $|\delta| = n$, there exists an injective function $f_\delta\colon \mathrm{SYam}(n) \hookrightarrow S_n$ fixing $\mathrm{SYam}_\delta(n)$ and preserving σ such that for any component $\mathcal{C} = (V, \sigma, E)$ of $\mathcal{H}_\delta$,*

$$\sum_{v \in V} F_{\sigma(v)}(X) = \sum_{\lambda \vdash n} |\{\pi \in \mathrm{SYam}(\lambda)\colon f_\delta(\pi) \in V\}| \cdot s_\lambda \tag{3.3.2}$$

Conjecture 3.3.9 (Corollary of Conjecture 3.3.8)**.** *Given any diagram δ and function f_δ as in Conjecture 3.3.8,*

$$\widetilde{H}_\delta(X;q,t) = \sum_{\lambda \vdash |\delta|} \sum_{\pi \in \mathrm{SYam}(\lambda)} q^{\mathrm{inv}_\delta(f_\delta(\pi))} t^{\mathrm{maj}_\delta(f_\delta(\pi))} s_\lambda.$$

Conjecture 3.3.9 has been explicitly checked when δ is a partition shape of size at most seven. It should be mentioned, however, that f_δ was defined in an ad hoc fashion for each new δ.

3.4 Further Applications to Symmetric Functions

3.4.1 More Analysis of $\widetilde{H}_\mu(X;0,t)$ and $\widetilde{H}_\mu(X;q,t)$

We can now explicitly answer the question of Garsia mentioned in Section 3.1. We also provide the analogous result for Macdonald polynomials.

Corollary 3.4.1. *Given a partition μ, the following equality holds if and only if μ does not contain $(3,3,3)$ as a subdiagram.*

$$\widetilde{H}_\mu(X;0,t) = \sum_{\lambda \vdash |\mu|} \sum_{\substack{w \in \mathrm{Yam}(\lambda) \\ \mathrm{inv}_\mu(w)=0}} t^{\mathrm{maj}_\mu(w)} s_\lambda. \tag{3.4.1}$$

Proof. First assume that $(3,3,3)$ is a subdiagram of μ. In light of Theorem 3.1.1, it suffices to show that there exists $w \in \mathrm{Yam}(\lambda)$ for some $\lambda \vdash n$ such that $\mathrm{inv}_\mu(w) = 0$ and w jams μ. We may explicitly demonstrate the desired w by placing 1's in all cells of μ except the first cell of the second row and the first three cells of the third row. Now fill the four remaining cells with 3232 in row reading order, as in Figure 3.14, and then define w as the row reading word of this filling. Thus, $\mathrm{inv}(w) = 0$, and the index of the last 2 of w jams the pistol starting at the first cell of the third row from the bottom.

Now suppose that $(3,3,3)$ is not a subdiagram of μ. We need to show that $\mathrm{Yam}_\mu(\lambda) = \mathrm{Yam}(\lambda)$ to apply Theorem 3.1.1. Let $w \in \mathrm{Yam}(\lambda)$ be the row reading word of a filling T of μ such that $\mathrm{inv}(T) = 0$. Since a Yamanouchi word must end in a 1, and the bottom row of T must be weakly increasing to avoid inversion pairs, the bottom row must be all 1's. Similarly focusing on the construction of Yamanouchi words and the description of inversion triples in Figure 3.6, it is readily shown that the second row starts with some number of 2's followed by all 1's. Thus, w cannot jam any of the pistols contained in the bottom two rows. Now consider any pistol that ends before the bottom row. Since μ does not contain $(3,3,3)$, any such pistol has at most three cells. In a Yamanouchi word, the index of the j^{th} i differs by at least three from the index of the $j+1^{th}$ $i+1$, and so cannot share a pistol containing less than four cells. Hence, no pistol that ends before the bottom row can be jammed by an index of w. That is, $\mathrm{Yam}_\mu(\lambda) = \mathrm{Yam}(\lambda)$. $\square$

$$\begin{array}{llll} 1 & 1 & & \\ 3 & 2 & 3 & \\ 2 & 1 & 1 & \\ 1 & 1 & 1 & 1 \end{array}$$

Figure 3.14: A filling with row reading word $w \in \mathrm{SYam}(\lambda)$ such that $\mathrm{inv}_\mu(w) = 0$ and w jams μ.

Proposition 3.4.2. *Given a partition μ, the following equality holds if and only if μ does not contain (4) or $(3,3)$ as a subdiagram.*

$$\widetilde{H}_\mu(X; q, t) = \sum_{\lambda \vdash |\mu|} \sum_{w \in \mathrm{Yam}(\lambda)} q^{\mathrm{inv}_\mu(w)} t^{\mathrm{maj}_\mu(w)} s_\lambda. \tag{3.4.2}$$

Proof. If we assume that μ does not contain (4) or $(3,3)$ as a subdiagram, the result is given by Proposition 3.2.13. It then suffices to assume that μ contains (4) or $(3, 3)$ and show that Equation (3.4.2) does not hold.

We proceed by considering the coefficients of q^2t^0. Focusing on the right hand side of Equation (3.4.2), consider a Yamanouchi word w such that $\mathrm{inv}_\mu(w) = 2$ and $\mathrm{maj}_\mu(w) = 0$. In particular, $\mathrm{maj}_\mu(w) = 0$ forces the columns of $T_\mu(w)$ to be weakly increasing when read downward. The bottom row of $T_\mu(w)$ must contribute at most two inversions, so w must end in $1 \ldots 111$, $1 \ldots 121$, or $1 \ldots 1211$. That is, the bottom row can have at most one 2. Applying the fact that all columns are weakly increasing, all but one column must be all 1's.

We now consider each of these possibilities for the bottom row separately. If the bottom row is all 1's, then every column must be all 1's, and so there cannot be any inversions. If the bottom row is $1 \ldots 121$ or $1 \ldots 1211$ and there is an $x > 1$ above the 2 in the bottom row, then there are at least two inversion triples containing x in the first case and at least one in the second case. Both situations force $\mathrm{inv}_\mu(w) > 2$. Hence, any values greater than one must occur in the bottom row. Because $\mathrm{inv}_\mu(w) = 2$, we are left with only the filling containing all 1's except for a bottom row filled by $1 \ldots 1211$. The conclusion of this analysis is that

$$\sum_{\lambda \vdash |\mu|} \sum_{w \in \mathrm{Yam}(\lambda)} q^{\mathrm{inv}_\mu(w)} t^{\mathrm{maj}_\mu(w)} s_\lambda \Big|_{q^2 t^0} = s_{(n-1,1)}. \tag{3.4.3}$$

Now consider the coefficient of q^2t^0 in $\widetilde{H}_\mu(X;q,t)$, as described in (3.2.6), when μ contains (4) or $(3,3)$. By letting π be the standardization of $w = 1 \ldots 1\,1\,2\,1\,3\,2$, we have $\mathrm{maj}_\mu(\pi) = 0$ and $\mathrm{inv}_\mu(\pi) = 2$ (see Figure 3.15). Thus, $q^2t^0F_{\sigma(\pi)}$ has a positive coefficient in the expansion of $\widetilde{H}_\mu(X;q,t)$ into fundamental quasisymmetric functions. However, $F_{\sigma(\pi)}$ has coefficient 0 in the expansion of $s_{(n-1,1)}$. This is clear because $\sigma(\pi) = + \cdots + + - + -$, but all fundamental quasisymmetric functions that contribute to $s_{(n-1,1)}$ have exactly one minus sign. Hence, there must be some term of $\widetilde{H}_\mu(X;q,t)$ of the form $q^2t^0s_\lambda$, where $\lambda \neq (n-1,1)$. Therefore, (3.4.2) cannot hold if μ contains (4) or (3,3) as a subdiagram, completing our proof. $\square$

Figure 3.15: Fillings by $w = 1 \ldots 112132$ with $\mathrm{inv}_\mu(w) = 2$ and $\mathrm{maj}_\mu(w) = 0$.

3.4.2 Further Analysis of $R_{\gamma,\delta}(X)$

In this section we give a method for quickly finding the leading term of the expansion of $R_{\gamma,\delta}(X)$. We begin by using the relationship between γ and δ to give a characterization of when $R_{\gamma,\delta}(X) = 0$. To do so, we will need the following definition.

Definition 3.4.3. Given any diagrams γ and δ such that $\gamma \subset \delta$, we refer to γ as a *realizable descent set* of δ if the following hold.

1. If $(x,y) \in \gamma$, then $(x, y-1) \in \delta$.

2. If $x_1 < x_2$ are any integers and I is any integer interval such that $(x_1, \max(I)), (x_1, \min(I)) \notin \gamma$, and $(x_1, I \setminus \{\min(I), \max(I)\}) \subset \gamma$, then $(x_2, I \setminus \{\min(I)\}) \not\subset \gamma$.

Figure 3.16: Diagrams with bullets representing γ and boxes representing δ. From the left, three examples where γ is not a realizable descent set of δ, then three examples where γ is a realizable descent set of δ.

Proposition 3.4.4. *Given any two diagrams* γ *and* δ *such that* $\gamma \subset \delta$*, then* $R_{\gamma,\delta}(X) \neq 0$ *if and only if* γ *is a realizable descent set of* δ*.*

The proof of Proposition 3.4.4 is postponed until the end of this section.

Definition 3.4.5. Given diagrams γ and δ such that γ is a realizable descent set of δ, define the *leading Yamanouchi word* of $R_{\gamma,\delta}(X)$, denoted $w_{\gamma,\delta}$, as the row reading word of the filling of δ achieved by placing a 1 in every cell of $\delta \setminus \gamma$ and then placing values in the rows of γ from bottom to top, filling each cell with one plus the value in the cell immediately below it in δ.

Notice that $w_{\gamma,\delta}$ is indeed a Yamanouchi word. We can then use $w_{\gamma,\delta}$ to provide the leading term in the expansion of $R_{\gamma,\delta}(X)$ into Schur functions.

Proposition 3.4.6. *Given diagrams* γ *and* δ *such that* γ *is a realizable descent set of* δ*, let* $R_{\gamma,\delta}(X) = \sum c_\lambda s_\lambda$ *for some nonzero integers* c_λ*, and let* $w_{\gamma,\delta} \in \mathrm{Yam}(\mu)$*, then*

1. $c_\lambda = 1$ *if* $\lambda = \mu$*,*
2. $c_\lambda = 0$ *if* $\lambda > \mu$ *in lexicographic order.*

Proof. Let $w_{\gamma,\delta}$ be as described in the proposition. We begin by showing that $T = T_\delta(w_{\gamma,\delta})$ has $\mathrm{Des}(T) = \gamma$ and $\mathrm{inv}(T) = 0$. The fact that $\mathrm{Des}(T) = \gamma$ is a result of the definition of $w_{\gamma,\delta}$. We thus need to show that $\mathrm{inv}(T) = 0$.

First consider a triple of T consisting of cells c, d, and e, in row reading order. Notice that if c is a descent of T, then the value in c is one greater than the value in e. It then follows that cells c, d, and e cannot form an inversion triple, regardless of the value in d.

Now suppose that c is not a descent of T. Then the value in c is a 1, so it suffices to show that the value in d is weakly less than the value in e. If the value in d is 1, we are done. We may then assume that d is a descent. By the definition of $w_{\gamma,\delta}$, it follows that the value in d is equal to the number of consecutive descents of T occurring in the cells weakly below d. Appealing to the definition of realizable descent set, there must be at least as many consecutive descents of T starting at e and continuing down. Thus, the value in d is weakly less than the value in e. That is, the cells c, d, and e do not form an inversion triple.

Next consider two cells c and d, in row reading order, that could form an inversion pair. If d is in the row below c, then the same reasoning as above shows that d must be weakly greater than c. If c and d are in the same row, then the cell below c cannot be in δ. Thus c is not a descent of T and so has value 1 in T. In particular, c and d do not form an inversion pair. Therefore, $\operatorname{inv}(T) = 0$.

Notice that the values in each cell of T are as small as possible while respecting $\operatorname{Des}(T) = \gamma$. Thus, there cannot be any other filling with row reading word $w \in \operatorname{SYam}(\mu)$, $\operatorname{inv}_\delta(w) = 0$, and $\operatorname{Des}_\delta(w) = \gamma$, where $\lambda \leq \mu$ in lexicographic order. Appealing to Theorem 3.1.1 completes the proof. □

Proof of Proposition 3.4.4. We wish to show that γ is a realizable descent set of δ if an only if $R_{\gamma,\delta}(X) \neq 0$. First assume that γ is not a realizable descent set of δ. Notice that if γ does not satisfy Part 1 of Definition 3.4.3, then there are no fillings of δ with descent set γ. Thus, $R_{\gamma,\delta}(X) = 0$. Next, suppose that T is a filling of δ with $\operatorname{Des}(T) = \gamma$ such that γ satisfies Part 1 but not Part 2 of Definition 3.4.3. It suffices to show that $\operatorname{inv}(T) \neq 0$. We suppose, for the sake of contradiction, that $\operatorname{inv}(T) = 0$.

Choose $\{x_1, x_2\} \times I$ violating Part 2 of Definition 3.4.3. Label the values of T in this rectangle by a, b, c, and so on, in row reading order. Here, we will call the upper left value a, whether or not the value exists in T. Regardless of existence in δ, the cell containing a is not in γ. See Figure 3.17 for an illustration.

Because a is not in a descent of T, and T has no inversion triples, it follows that $a < b < c$. Also, $b > d$ because b is in a descent of T, and so $c > d$. Assuming that c is a descent of T, it then follows that $c > e > d > f$, since T has no inversions and d is in a descent of T. In particular, notice that $e > f$. Continuing in this fashion recursively, it follows that the value in the bottom left corner is greater than the value in the bottom right corner. However, the bottom left corner is not a descent of T, and so its value must be greater than the value to its right if $\mathrm{inv}(T) = 0$, as desired. Therefore, $R_{\gamma,\delta}(X) = 0$ whenever γ is not a realizable descent set of δ.

Figure 3.17: At left, a rectangle violating Part 2 of Definition 3.4.3 with γ represented by bullets and possibly the open circle. At right, the values of T in this rectangle.

We still need to consider the case were γ is a realizable descent set of δ and show that $R_{\gamma,\delta}(X) \neq 0$ in this case. By Proposition 3.4.6, the Schur expansion of $R_{\gamma,\delta}(X)$ has a nonzero term, and so $R_{\gamma,\delta}(X) \neq 0$, completing the proof. $\square$

Chapter 4

COXETER-KNUTH GRAPHS AND A SIGNED LITTLE MAP FOR TYPE B REDUCED WORDS

4.1 Introduction

Stanley symmetric functions F_w appear in the study of reduced words of permutations [Stanley, 1984], the representation theory of generalized Specht modules [Kraśkiewicz, 1995], and the geometry of positroid varieties [Knutson et al., 2013]. The F_w are known to have a Schur positive expansion with coefficients determined by the Edelman-Greene correspondence. This correspondence associates to each reduced word a pair of tableaux (P, Q) of the same shape where the second tableau is standard. These symmetric functions F_w can be defined as the sum of certain fundamental quasisymmetric functions where the sum is over all reduced words for $w \in S_n$, denoted $R(w)$. In particular, the coefficient of $x_1 x_2 \cdots x_{\ell(w)}$ in F_w equals $|R(w)|$. There is a recurrence relation for F_w derived from Lascoux and Schützenbergers' transition equation for Schubert polynomials [Lascoux and Schützenberger, 1982] of the form

$$F_w = \sum_{w' \in T(w)} F_{w'},$$

along with the base cases that F_w is a single Schur function if w has at most one descent; in this case we say w is *Grassmannian*. By taking the coefficient on both sides of the recurrence, we see that the sets $R(w)$ and $\cup_{w' \in T(w)} R(w')$ are equinumerous.

David Little gave a remarkable bijection between $R(w)$ and $\cup_{w' \in T(w)} R(w')$ [Little, 2003] inspired by the lectures of Adriano Garsia, which are published as a book [Garsia, 2002] . This algorithm is a finite sequence of steps, each of which decrements one letter in the word. If ever a 1 is decremented to a 0, then instead the whole reduced word is lifted up by one to make space for one extra generator. This bijection is an instance of a more general phenomenon known as *Little bumps*.

Recently, Hamaker and Young [Hamaker and Young, 2013] have shown that Little bumps preserve the recording tableaux under the Edelman-Greene correspondence. This proved a conjecture of Thomas Lam [Lam, 2010, Conj. 2.5]. They further show that all reduced words with a given recording tableau Q under the Edelman-Greene correspondence are connected via Little bumps. Edelman and Greene gave a refinement on the Coxeter relations in type A, which they call *Coxeter-Knuth* relations. These relations preserve the insertion tableaux under the Edelman-Greene correspondence, and the set of reduced words which have a fixed insertion tableau P is connected by elementary Coxeter-Knuth relations. Hamaker and Young further showed that two reduced words that differ by an elementary Coxeter-Knuth relation give rise to Q tableaux that differ in exactly two positions. This can be made more precise. Consider the graph $CK_A(w)$ on all reduced words for w with an edge labeled i between two reduced words $a = a_1a_2 \dots a_p$ and $b = b_1b_2 \dots b_p$ whenever a and b differ by an elementary Coxeter-Knuth relation in positions $i, i+1, i+2$. Call $CK_A(w)$ a *Coxeter-Knuth* graph. Using the theory of dual equivalence graphs due to Assaf [Assaf, 2013] and the equivalent axioms given by Roberts [Roberts, 2014], one can easily show that $CK_A(w)$ is a dual equivalence graph and the Q tableaux for two reduced words differing by an elementary Coxeter-Knuth move differ by one of Haiman's dual equivalence moves [Haiman, 1992].

In this paper, we define the analog of Little bumps $B^{\delta}_{(i,j)}$ on reduced words for the signed permutations B_n, and show that these maps satisfy many of the same properties as in the original case. In particular, there is a close connection to the Stanley symmetric functions for types B and C defined in [Billey and Haiman, 1995], see also [Fomin and Kirillov, 1996; Lam, 1995]. These Stanley symmetric functions again satisfy a transition equation [Billey, 1998], which proves that $R(w)$ is equinumerous with a certain union of $R(w')$'s.

To concretely state our first main result, we need to establish some notation. A signed permutation $w \in B_n$ is a bijection from $\{-n, \dots -1, 1, 2, \dots, n\}$ to itself such that $w(i) = -w(-i)$. One could represent w in one-line notation either by listing $[w(-n), w(-n-1), \dots, w(-1), w(1), \dots, w(n)]$ in long form or simply $[w(1), \dots, w(n)]$ in short form. For example, $[1, \bar{2}, \bar{4}, 3, \bar{3}, 4, 2, \bar{1}]$ and $[\bar{3}, 4, 2, \bar{1}]$ represent the same element in B_4 where $-i$ is denoted $\bar{i}$. For our purposes, we identify $v \in B_n$ with the element $w \in B_{n+1}$ such that $v(i) = w(i)$ for $1 \leq i \leq n$ and $w(n+1) = n+1$. Set $B_\infty = \cup B_n$ in this identification. For

$i < j \in \mathbb{Z} \setminus \{0\}$, let t_{ij} be the (signed) transposition such that $wt_{ij}(i) = w(j)$, $wt_{ij}(j) = w(i)$, $wt_{ij}(-i) = -w(j)$, $wt_{ij}(-j) = -w(i)$ and for every integer $k \notin \{\pm i, \pm j, 0\}$ we have $wt_{ij}(k) = w(k)$. If $w \in B_\infty$ has $w(1) < w(2) < \ldots$, we say w is *increasing.* If w is not increasing, let $(r < s)$ be the lexicographically largest pair of positive integers such that $w_r > w_s$. Set $v = wt_{rs}$. Let $T(w)$ be the set of all signed permutations $w' = vt_{ir}$ for $i < r, i \neq 0$ such that $\ell(w') = \ell(w)$.

Theorem 4.1.1. *Using the notation above, if $w \in B_\infty$ is not increasing, then the particular Little bump $B^-_{(r,s)} : R(w) \longrightarrow \bigcup_{w' \in T(w)} R(w')$ is the bijection predicted by the transition equation for type C Stanley symmetric functions.*

The analog of Edelman-Greene insertion and elementary Coxeter-Knuth relations for signed permutations were given by Kraśkiewicz [Kraśkiewicz, 1989]. Kraśkiewicz insertion inputs a reduced word a and outputs two shifted tableaux $(P'(a), Q'(a))$ of the same shifted shape where the recording tableau $Q'(a)$ is standard.

Theorem 4.1.2. *Say w and wt_{ij} are signed permutations such that $\ell(w) = \ell(wt_{ij}) + 1$.*

1. *The Little bump $B^\delta_{(i,j)}$ maps $R(w)$ to reduced words for some signed permutation $w' = wt_{ij}t_{kl}$ with $\ell(w) = \ell(w')$.*

2. *Two reduced words a and b are connected via Little bumps if and only if $Q'(a) = Q'(b)$ under Kraśkiewicz insertion.*

3. *For each standard shifted tableau Q', there exists a unique reduced word a for an increasing signed permutation such that $Q'(a) = Q'$.*

The Coxeter-Knuth relations given by Kraśkiewicz lead to a type B Coxeter-Knuth graph $CK_B(w)$ for each $w \in B_\infty$. An important step in proving Theorem 4.1.2 is showing that two reduced words for signed permutations that differ by an elementary Coxeter-Knuth relation give rise to two Q' tableaux that differ by one of Haiman's shifted dual equivalence moves [Haiman, 1992]. In fact, shifted dual equivalence completely determines the graph structure for a type B Coxeter-Knuth graphs and vice versa. Thus, we define shifted dual equivalence graphs in analogy with the work of Assaf and Roberts on dual equivalence graphs.

Theorem 4.1.3. *Every type B Coxeter-Knuth graph $CK_B(w)$ is a shifted dual equivalence graph with signature function given via peak sets of reduced words. The isomorphism is given by Q' in Kraśkiewicz insertion. Conversely, every connected shifted dual equivalence graph is isomorphic to the Coxeter-Knuth graph for some increasing signed permutation.*

Putting Theorem 4.1.2 and Theorem 4.1.3 together, one can see that Little bumps in both type A and type B play a similar role for Stanley symmetric functions as jeu de taquin plays in the study of Littlewood-Richardson coefficients for skew-Schur functions. We will show that there is exactly one reduced word for a unique increasing signed permutation in each communication class under Little bumps.

We give local axioms characterizing graphs isomorphic to shifted dual equivalence graphs or equivalently Coxeter-Knuth graphs of type B. We state the theorem here using some terminology that is developed in Section 4.5.

Theorem 4.1.4. *A signed colored graph $\mathcal{G} = (V, \sigma, E)$ of shifted degree $[n]$ is a shifted dual equivalence graph if and only if the following local properties hold.*

1. *If I is any interval of integers with $|I| \leq 9$, then each component of $\mathcal{G}|_I$ is isomorphic to the standard shifted dual equivalence graph of a shifted shape of size up to $|I|$.*

2. *If $i, j \in \mathbb{N}$ with $|i - j| > 3$, $(u, v) \in E_i$ and $(u, w) \in E_j$, then there exists a vertex $y \in V(\mathcal{G})$ such that $(v, y) \in E_j$ and $(w, y) \in E_i$.*

We propose that the study of Coxeter-Knuth graphs initiated in this paper is an interesting way to generalize dual equivalence graphs to other Coxeter group types. For example, in type A, dual equivalence graphs have been shown to be related to crystal graphs [Assaf, 2008]. Furthermore, the transition equation due to Lascoux and Schützenberger follows from Monk's formula for multiplying a special Schubert class of codimension 1 with an arbitrary Schubert class in the flag manifold of type A. The elementary Coxeter-Knuth relations could have been derived from the Little bijection provided one understood the Coxeter-Knuth relations for the base case of the transition equations in terms of Grassmannian permutations. The transition equations for the other classical groups follow from Chevellay's generalization for Monk's formula on Schubert classes [Chevalley, 1994]. In fact, there is a very general Chevellay Formula for all Kac-Moody groups [Lenart and Shimozono, 2012].

We comment on one generalization which did not work as hoped. In type A, Chmutov showed that the molecules defined by Stembridge's axioms can be given edge labels in such a way that the graphs are dual equivalence graphs [Chmutov, 2013]. Alas, in type B, this does not appear to be possible. The Kazhdan-Lusztig graph for B_3 has a connected component with an isomorphism type that does not occur for dual equivalence graphs or shifted dual equivalence graphs. Namely, the component of $[2, 1, \bar{3}]$ is a tree with 4 vertices and 3 leaves $[1, \bar{2}, \bar{3}], [\bar{2}, \bar{3}, 1], [2, 1, \bar{3}]$.

The paper proceeds as follows. In Section 4.2, we review the necessary background on permutations and signed permutations as Coxeter groups. In Section 4.3, we formally define the signed Little bumps and pushes. The key tool we use to visualize the algorithms is the wiring diagram of a reduced word. The relation between Little bumps, the recording tableaux under the Kraśkiewicz recording tableaux, Coxeter-Knuth moves of type B and shifted dual equivalence moves is discussed in Section 4.4. In Section 4.5, the shifted dual equivalence graphs are equivalently defined in terms of either shifted dual equivalence moves or Coxeter-Knuth moves proving Theorem 4.1.3. We go on to prove many lemmas leading up to the axiomitization of shifted dual equivalence graphs proving Theorem 4.1.4. We conclude with some interesting open problems in Section 4.6.

We recently learned that Assaf has independently considered shifted dual equivalence graphs in connection to a new Schur positive expansion of the Schur P-polynomials [Assaf, 2014].

4.2 Background

Let W be a Coxeter group with generators $S = \{s_1, \ldots, s_n\}$ and elementary relations $(s_i s_j)^{m(i,j)} = 1$. For $w \in W$, let $\ell(w)$ be the minimal length of any expression $s_{a_1} \cdots s_{a_q} = w$. If $\ell(w) = p$, we say $s_{a_1} \cdots s_{a_p}$ is a *reduced expression* and the list of subscripts $a_1, \ldots, a_p$ is a *reduced word* for $w = s_{a_1} \cdots s_{a_p}$. Let $R(w)$ be the set of reduced words for w.

For $w \in W$, one can define a graph $G(w)$ with vertices given by the reduced words of w using the Coxeter relations. In this graph, any two reduced words are connected by an edge if they differ only by an elementary relation of the form $s_is_js_i\cdots = s_js_is_j\cdots$ where each side is a product of $m(i,j)$ generators. It is a well known theorem, sometimes attributed to Tits, that this graph is connected [Björner and Brenti, 2005, Thm. 3.3.1].

4.2.1 Type A

The symmetric group S_n is the Coxeter groups of type A_{n-1}. For our purposes, we can think of $w \in S_n$ in one-line notation as $w = [w_1, w_2, \ldots, w_n]$ or as $w = [w_1, w_2, \ldots] \in S_\infty$ with $w_i = i$ for all $i > n$. Let t_{ij} be the transposition interchanging i and j and fixing all other values. Then right multiplication by t_{ij} interchanges the values in positions i and j in w.

The group S_n is minimally generated by the *adjacent transpositions* $s_1, \ldots s_{n-1}$ with $s_i = t_{i,i+1}$ with elementary Coxeter relations

1. **Commutation**: $s_is_j = s_js_i$ provided $|i-j| > 1$,

2. **Braid**: $s_is_{i+1}s_i = s_{i+1}s_is_{i+1}$.

For example, if $w = [2,1,5,4,3]$, then $R(w) = \{1343, 3143, 3413, 3431, 4341, 4314, 4134, 1434\}$ and $G(w)$ is a cycle on these 8 vertices.

In [Edelman and Greene, 1987], Edelman-Greene (EG) gave an insertion algorithm much like the famous Robinson-Schensted-Knuth (RSK) algorithm for inserting a reduced word into a tableau for some partition $\lambda = (\lambda_1 \geq \lambda_2 \geq \ldots \geq \lambda_k > 0)$. The one difference in EG insertion is that when inserting an i into a row that already contains an i and $i+1$, we skip that row and insert $i+1$ into the next row. If one keeps track of the recording tableau of the insertion, then the process is invertible. Let $P(a)$ be the EG insertion tableau for $a = a_1 \ldots a_p$ and let $Q(a)$ be the recording tableau. Define $EG(w) = \{P : P = P(a) \text{ for some } a \in R(w)\}$.

For example, using EG insertion, the reduced word 1343 inserts to give

$$\begin{ytableau} 1 \end{ytableau} \to \begin{ytableau} 1 & 3 \end{ytableau} \to \begin{ytableau} 1 & 3 & 4 \end{ytableau} \to \begin{ytableau} 4 \\ 1 & 3 & 4 \end{ytableau}.$$

So

$$P(1343) = \begin{ytableau} 4 \\ 1 & 3 & 4 \end{ytableau} \qquad \text{and} \qquad Q(1343) = \begin{ytableau} 4 \\ 1 & 2 & 3 \end{ytableau}$$

in French notation.

The EG recording tableau $Q(a)$ is a *standard Young tableau* of partition shape λ, denoted $SYT(\lambda)$. These are bijective fillings of the Ferrers diagram for the partition λ with rows and columns increasing. The *row reading word* of a standard tableau T is the permutation in one-line notation obtained by reading along the rows of T in the French way, left to right and top to bottom. The *ascent set* of T is the set of all i such that i preceeds $i+1$ in the row reading word of T. Similarly, define the *ascent set* of a reduced word $a = a_1 \dots a_p$ to be $\{j : a_j < a_{j+1}\}$.

Theorem 4.2.1. *[Edelman and Greene, 1987] Fix $w \in S_\infty$ and $P \in EG(w)$. Then, the recording tableau for EG insertion gives a bijection between $\{a \in R(w) : P(a) = P\}$ and the set of standard Young tableaux of the same shape as P. Furthermore, this bijection preserves ascent sets.*

Definition 4.2.2. Let $a_{\lambda,w}$ be the number of distinct tableaux $P \in EG(w)$ such that P has shape λ. We call these numbers the *Edelman-Greene* coefficients.

Definition 4.2.3. [Stanley, 1984] For $w \in S_\infty$ and $a = a_1 \cdots a_p \in R(w)$, let $I(a)$ be the set of all increasing integer sequences $1 \le i_1 \le i_2 \le \cdots \le i_p$ such that $i_j < i_{j+1}$ whenever $a_j < a_{j+1}$. The *Stanley symmetric function* $F_w = F_w^A$ is defined by

$$F_w^A = \sum_{a \in R(w)} \sum_{i_1 i_2 \cdots i_p \in I(a)} x_{i_1} x_{i_2} \cdots x_{i_p}.$$

Here, $\sum_{i_1 i_2 \cdots i_p \in I(a)} x_{i_1} x_{i_2} \cdots x_{i_p}$ is the fundamental quasisymmetric function indexed by the ascent set of a [Stanley, 2001, Ch. 7.19]. Edelman-Greene showed that the ascent set of $a \in R(w)$ agrees with the ascent set of $Q(a)$. Furthermore, Ira Gessel [Gessel, 1984] showed that the Schur function s_λ is the sum over all standard tableaux T of shape λ of the fundamental quasisymmetric function by the ascent set of T. Putting this together gives the following theorem.

Theorem 4.2.4. *[Edelman and Greene, 1987] Fix $w \in S_\infty$. Then*

$$F_w^A = \sum_\lambda a_{\lambda,w} s_\lambda,$$

where s_λ *is the Schur function indexed by the partition* λ *and each* $a_{\lambda,w}$ *is a nonnegative integer.*

Edelman-Greene also characterized when two reduced expressions give rise to the same P tableau by restricting the elementary Coxeter relations. For this characterization, they define the *elementary Coxeter-Knuth relations* to be either a braid move or a witnessed commutation move:

1. $ikj \leftrightarrow kij$ for all $i < j < k$,
2. $jik \leftrightarrow jki$ for all $i < j < k$,
3. $i(i+1)i \leftrightarrow (i+1)i(i+1)$,

Two words which are connected via a sequence of Coxeter-Knuth relations are said to be in the same *Coxeter-Knuth class.*

Theorem 4.2.5. *[Edelman and Greene, 1987] Let* $a, b \in R(w)$*. Then* $P(a) = P(b)$ *if and only if* a *and* b *are in the same Coxeter-Knuth class.*

In the example $w = [2, 1, 5, 4, 3]$, there are three Coxeter-Knuth classes $\{3143, 3413\}$, $\{3431, 4341, 4314\}$, and $\{1343, 4134, 1434\}$, which respectively insert to the three P tableaux:

$$\begin{array}{|c|c|}\hline 3 & 4 \\ \hline 1 & 3 \\ \hline \end{array} \qquad \begin{array}{|c|c|}\cline{1-1} 4 & \multicolumn{1}{c}{} \\ \cline{1-1} 3 & \multicolumn{1}{c}{} \\ \hline 1 & 4 \\ \hline \end{array} \qquad \begin{array}{|c|c|c|}\cline{1-1} 4 & \multicolumn{2}{c}{} \\ \hline 1 & 3 & 4 \\ \hline \end{array}.$$

Let $CK_A(w)$ be the *Coxeter-Knuth graph* for $w \in S_\infty$ with vertices $R(w)$ and colored (labeled) edges constructed using Coxeter-Knuth relations. An edge between a and b is labeled i if $a_j = b_j$ for all $j \notin \{i, i+1, i+2\}$ and $a_i a_{i+1} a_{i+2}$ and $b_i b_{i+1} b_{i+2}$ differ by an elementary Coxeter-Knuth relation. To each vertex $a \in R(w)$ associate a signature determined by its ascent set, $\sigma(a) = \{j : a_j < a_{j+1}\}$. If $\ell(w) = p$, we denote a subset $S \subset \{1, \ldots, p-1\}$ by a sequence in $\{+, -\}^{p-1}$ where $+$ in the j^{th} position means $j \in S$. Here $\sigma_j(a) = +$ if $a_j < a_{j+1}$ and $\sigma_j(a) = -$ if $a_j > a_{j+1}$. See Figure 4.1 for example and compare to $G(21543)$, which is a cycle with eight vertices, as mentioned above.

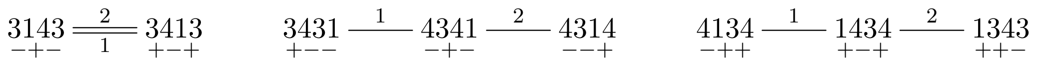

Figure 4.1: The Coxeter-Knuth graph of $w = [2, 1, 5, 4, 3]$

Let α_i be the involution defined by the edges of $CK_A(w)$, namely $\alpha_i(a) = b$ provided a and b are connected by an edge colored i, or equivalently an elementary Coxeter-Knuth relation on positions $i, i+1, i+2$. If a is not contained in an i-edge, then define $\alpha_i(a) = a$.

The type A Coxeter-Knuth graphs are closely related to dual equivalence graphs on standard tableaux as defined by Assaf [Assaf, 2013]. For a partition λ, one defines a *standard dual equivalence graph* $\mathcal{G}_\lambda$ to be the graph with vertex set given by $SYT(\lambda)$, and an edge colored i between any two tableaux that differ by an elementary dual equivalence defined as follows.

Definition 4.2.6. *[Haiman, 1992]* Given a permutation $\pi \in S_n$, define the *elementary dual equivalence operator* d_i for all $1 \leq i \leq n-2$ as follows. Say $\{i, i+1, i+2\}$ occur in positions $a < b < c$ in π, then $d_i(\pi) = \pi t_{ac}$ provided $\pi(b) \neq i+1$ and $d_i(\pi) = \pi$ otherwise. Dual equivalence operators also act on standard tableaux by acting on their row reading word.

Theorem 4.2.7. *The graph $CK_A(w)$ is isomorphic to a disjoint union of standard dual equivalence graphs for each $w \in S_n$. The isomorphism preserves ascent sets on vertices. On each connected component, the Edelman-Green Q function provides the necessary isomorphism. Furthermore, ascent sets are preserved.*

4.2.2 Type B/C

The hyperoctahedral group, or signed permutation group B_n is also a finite Coxeter group. This group is the Weyl group of both the root systems of types B and C of rank n. This group is generated as a Coxeter group by the adjacent transpositions $s_1, \ldots, s_{n-1}$ with $s_i = t_{i,i+1}$ plus an additional generator $s_0 = t_{-1,1}$. Thus, if $w = [w_1, \ldots, w_n] \in B_n$, then $ws_0 = [-w_1, w_2, \ldots, w_n]$. For example, $[\bar{3}, 2, 1] = s_1 s_2 s_1 s_0 \in B_3$. Again, let $R(w)$ denote the set of all reduced words for w. Note that if $w \in S_n$, then it can also be considered as an element in B_n with the same reduced words. The *elementary relations* on the generators are given by

1. **Commutation**: $s_i s_j = s_j s_i$ provided $|i - j| > 1$,

2. **Short Braid**: $s_i s_{i+1} s_i = s_{i+1} s_i s_{i+1}$ for all $i > 0$,

3. **Long Braid**: $s_0 s_1 s_0 s_1 = s_1 s_0 s_1 s_0$.

Recall from Section 4.1 that we have defined the *(signed) transposition* t_{ij} to be the signed permutation interchanging i with j and $-i$ with $-j$ for all $i, j \neq 0$. Thus, $s_0 = t_{-1,1}$ in this notation.

Shifted tableaux play the same role in types B/C as the usual tableaux play in type A. Given a strict partition $\lambda = (\lambda_1 > \lambda_2 > \cdots > \lambda_k > 0)$, the *shifted shape* λ is the set of squares in positions $\{(i, j) \ : \ 1 \leq i \leq k, i \leq j \leq \lambda_i + i - 1\}$. A *standard shifted tableau* T is a bijective filling of a shifted shape with positive integers with rows and columns increasing. For example, see T in Figure 4.2. Let $SST(\lambda)$ be the set of standard shifted tableaux of shifted shape λ.

$$S = \begin{ytableau} 7 \\ 5 & 6 & 8 \\ 1 & 2 & 3 & 4 & 9 \end{ytableau}, \qquad T = \begin{ytableau} \none & \none & 7 \\ \none & 3 & 6 & 8 \\ 1 & 2 & 4 & 5 & 9 \end{ytableau}, \qquad U = \begin{ytableau} \none & \none & 0 \\ \none & 2 & 1 & 0 \\ 3 & 2 & 1 & 0 & 1 \end{ytableau}$$

Figure 4.2: A standard tableau S, standard shifted tableau T and unimodal tableau U respectively of shape $\lambda = (5, 3, 1)$.

We will also need to consider another type of tableaux on shifted shapes. We say a list $r = r_1 \ldots r_l$ is *unimodal* if there exists an index j, referred to as the *middle*, such that $r_1 \ldots r_j$ is decreasing and $r_j \ldots r_l$ is increasing. A *unimodal tableau* T is a filling of a shifted shape with positive integers such that the reading word along each row is unimodal.

In 1989, Kraśkiewicz[Kraśkiewicz, 1989] gave an analog of Edelman-Greene insertion for reduced words of signed permutations. Kraśkiewicz insertion is a variant of the mixed shifted insertion of [Haiman, 1989] that maps a reduced word b of a signed permutation to the pair of shifted tableaux $(P'(b), Q'(b))$ where $Q'(a)$ is a standard shifted tableau and $P'(a)$ is a unimodal tableau of the same shape such that the reading word given by reading rows left to right from row 1 to row k is a reduced word for w. Once again, there is an analog of the Coxeter-Knuth relations. We will need the details of this insertion map and relations for our main theorems. Our description of this map is based on an equivalent algorithm in Tao Kai Lam's Ph.D. thesis [Lam, 1995].

First, there is an algorithm to insert a positive integer into a unimodal sequence. Given a number k and a (potentially empty) unimodal sequence $r = r_1 \dots r_l$ with middle index j, we insert k into r as follows:

1. If $k \neq 0$ or $r_j \neq 0$, perform Edelman-Greene insertion of k into $r_{j+1} \dots r_l$. Call the bumped entry k^-, if it exists. Call the resulting string after insertion $v_1 \dots v_{l-j}$.

2. If $k = 0$ and $r_j = 0$, set $k^- = 1$. Set $v_1 \dots v_{l-j} = r_{j+1} \dots r_l$.

3. If k^- exists, perform Edelman-Greene insertion of $-k^-$ into $-r_1 \cdots - r_j$. This time a bumped entry will exist, as $k^- > r_j$. Call it k'. Set $u_1 \dots u_p$ to be the result of negating every entry in the resulting string after insertion and reversing it.

4. If k^- does not exist, set $u_1 \dots u_p = r_1 \dots r_j$.

The *Kraśkiewicz insertion* of a positive integer k into a shifted unimodal tableau P' starts by inserting k into the first row of P' using the algorithm above. Replace the first row of P by $u_1 \dots u_p v_1 \dots v_q$. If k^- exists and k' is the output, then insert k' in the second row of P, etc. Continue until no output exists or no further rows of P' exist. In that case, add k' in a new final row along the diagonal so the result is again a shifted unimodal tableau. Call the final tableau $P' \leftarrow k$. For $w \in B_\infty$ and $a = a_1 \cdots a_p \in R(w)$, let $P'(w)$ be the result of inserting $\emptyset \leftarrow a_1 \cdots a_p$ consecutively into the empty shifted unimodal tableau denoted $\emptyset$.

For example, using Kraśkiewicz insertion, on the same reduced word 1343 as before inserts to give

$$\boxed{1} \to \boxed{1}\boxed{3} \to \boxed{1}\boxed{3}\boxed{4} \to \begin{array}{ccc} & \boxed{1} & \\ \boxed{4} & \boxed{3} & \boxed{4} \end{array}.$$

So

$$P(1343) = \begin{array}{ccc} & \boxed{1} & \\ \boxed{4} & \boxed{3} & \boxed{4} \end{array} \qquad \text{and} \qquad Q(1343) = \begin{array}{ccc} & \boxed{4} & \\ \boxed{1} & \boxed{2} & \boxed{3} \end{array}.$$

Also, $021032101 \in R([\bar{3}, \bar{4}, \bar{1}, 2])$ gives $P = U, Q = T$ from Figure 4.2.

Kraśkiewicz insertion behaves well with respect to the peaks of a reduced word. Given any word $a_1 \ldots a_p$, we say a has an *ascent* in position $0 < i < p$ if $a_i < a_{i+1}$ and a *descent* if $a_i > a_{i+1}$. Similarly, we say a has a *peak* in position $1 < i < p$ if $a_{i-1} < a_i > a_{i+1}$. Define the *peak set* of $a \in R(w)$ to be $peaks(a) = \{1 < i < p \ : \ a_{i-1} < a_i > a_{i+1}\}$. Given a standard (shifted) tableau T, we say j is a peak of T provided j appears after $j-1$ and $j+1$ in the row reading word of T, so there is an ascent from $j-1$ to j and a descent from j to $j+1$. The peak set of T, denoted again $peaks(T)$, is defined similarly.

Theorem 4.2.8. *[Lam, 1995, Theorem 2.10] Given a signed permutation w and a reduced word $a \in R(w)$, $peaks(a) = peaks(Q'(a))$.*

One important tool for studying Kraśkiewicz insertion is a family of local transformations on words known as the *type B Coxeter-Knuth moves.* These moves are based on certain type B elementary Coxeter relations that depend on exactly four adjacent entries of a word.

Definition 4.2.9. [Kraśkiewicz, 1989] The *elementary Coxeter-Knuth moves of type B* are given by the following rules on any reduced word $i_1i_2i_3i_4$. If $i_1i_2i_3i_4$ has no peak then $\beta(i_1i_2i_3i_4) = i_1i_2i_3i_4$. If $i_1i_2i_3i_4$ has a peak in position 3, $\beta(i_1i_2i_3i_4)$ is given by reversing $\beta(i_4i_3i_2i_1)$. If $i_1i_2i_3i_4$ has a peak in position 2, then we have three cases:

1. **Long braid**: If $i_1i_2i_3i_4 = 0101$, then define $\beta(0101) = 1010$. Note 1010 is another reduced word for the same signed permutation, and it has a peak in position 3.

2. **Short braid witnessed by smaller value**: If there are 3 distinct letters among $i_1i_2i_3i_4$ and there is a corresponding short braid relation possible, say $i_1i_2i_3i_4 =$

$a\ b+1\ b\ b+1$ or $b\ b+1\ b\ a$ for some $a<b$. Define

$$\beta(a\ b+1\ b\ b+1) = a\ b\ b+1\ b \text{ and} \tag{4.2.1}$$

$$\beta(b\ b+1\ b\ a) = b+1\ b\ b+1\ a. \tag{4.2.2}$$

Again the sequence $\beta(i_1 i_2 i_3 i_4)$ has a peak in position 3. Also, this word is another reduced word for the same signed permutation which differ by a short braid move.

3. **Peak moving commutation**: In all other cases,

$$\beta(i_1 i_2 i_3 i_4) = (i_1 i_2 i_3 i_4)s_j$$

for the smallest j such that $(i_1 i_2 i_3 i_4)s_j$ is related to $i_1 i_2 i_3 i_4$ by a commuting move and has peak in position 3.

Observe that $i_1 i_2 i_3 i_4$ is fixed by β if and only if $i_1 i_2 i_3 i_4$ has no peak. Furthermore, the map β is an involution on $R(w)$ for w a signed permutation with $\ell(w)=4$. Define a family of involutions β_i acting on reduced words $a_1 a_2 \cdots a_p$ by replacing $a_i a_{i+1} a_{i+2} a_{i+3}$ by $\beta(a_i a_{i+1} a_{i+2} a_{i+3})$,provided $0<i<p-3$.

Theorem 4.2.10. *[Kraśkiewicz, 1989] Let a and b be reduced words of signed permutations. Then $P'(a) = P'(b)$ if and only if there exist Coxeter-Knuth moves of type B relating a to b. Furthermore, for each standard shifted tableau Q of the same shape as $P'(a)$, there exists a reduced word c for the same signed permutation such that $P'(c) = P'(a)$ and $Q'(c) = Q$.*

Using the Coxeter-Knuth moves of type B, we can define an analogous graph $CK_B(w)$ on the reduced words for $w \in B_n$ with edges defined by the involutions β_i. Each connected component of $CK_B(w)$ has vertex set given by a Coxeter-Knuth equivalence class $\{a \in R(w) : P'(a) = P'\}$, and assuming this set is nonempty, Q' gives a bijection between this set and the standard shifted tableaux of the same shape as P'. In Section 4.4, we will show that every connected component of $CK_B(w)$ is isomorphic to some $CK_B(v)$ where v is increasing. In Section 4.5, we will show that Q' gives an isomorphism of signed colored graphs with a graph on standard shifted tableaux of the same shape with edges given by shifted dual equivalence.

4.2.3 Stanley symmetric functions revisited

For signed permutations, there are two forms of Stanley symmetric functions and their related Schubert polynomials, see [Billey and Haiman, 1995; Fomin and Kirillov, 1996; Lam, 1995]. The distinct forms correspond to the root systems of type B and C, which both have signed permutations as their Weyl group. The definition we will give is the type C version, from which the type B version can be readily obtained. First, we introduce an auxiliary family of quasisymmetric functions.

In type A, ascent sets of reduced words can be used to define the Stanley symmetric functions. In type B/C, the peak set of a reduced word plays a similar role.

Definition 4.2.11. [Billey and Haiman, 1995, Eq. (3.2)] Let $X = \{x_1, x_2, \dots\}$ be an alphabet of variables. The *peak fundamental quasisymmetric function* of degree d on a possible peak set P is defined by

$$\Theta_P^d(X) = \sum_{(i_1 \leq \cdots \leq i_d) \in A_d(P)} 2^{|\{i_1, i_2, \dots, i_d\}|} x_{i_1} x_{i_2} \cdots x_{i_d}$$

and $A_d(P)$ is the set of all *admissible sequences* $(1 \leq i_1 \leq \cdots \leq i_d)$ such that $i_{k-1} = i_k = i_{k+1}$ only occurs if $k \notin P$.

The peak fundamental quasisymmetric functions also arise in Stembridge's enumeration of P-partitions [Stembridge, 1997] and are a basis for the peak subalgebra of the quasisymmetric functions as studied by [Billera et al., 2003; Schocker, 2005] and many others. They are also related to the Schur Q-functions $Q_\mu(X)$ which are specializations of Hall-Littlewood polynomials $Q_\mu(X;t)$ with $t = -1$, see [Macdonald, 1995, III]. By [Billey and Haiman, 1995, Prop. 3.2], the following is an equivalent definition of Schur Q-functions.

Definition 4.2.12. *For a shifted shape μ, the Schur Q-function $Q_\mu(X)$ is*

$$Q_\mu(X) = \sum_T \Theta^{|\mu|}_{peaks(T)}(X)$$

where the sum is over all standard shifted tableaux T of shape μ.

Remark 4.2.13. In this way, the peak fundamental quasisymmetric functions play the role of the original fundamental quasisymmetric functions in Gessel's expansion of Schur functions [Gessel, 1984].

Let g^{μ}_{w} be the number of distinct shifted tableaux of shape μ that occur as $P'(a)$ for some $a \in R(w)$ under Kraśkiewicz insertion. The numbers g^{μ}_{w} can equivalently be defined as the number of reduced words in $R(w)$ mapping to any fixed standard tableaux of shape μ by Haiman's promotion operator [Haiman, 1992, Prop. 6.1 and Thm. 6.3]. Haiman's promotion operator on $a \in R(w)$ in type B is equivalent to Kraśkiewicz's $Q'(a)$. Recall from Theorem 4.2.8 that $Q'(a)$ and a have the same peak set which implies the equivalence in the following definition.

Definition 4.2.14. [Billey and Haiman, 1995, Prop. 3.4] For $w \in B_{\infty}$ with $d = \ell(w)$, define the type C *Stanley symmetric function* to be

$$\begin{aligned} F^{C}_{w}(X) &= \sum_{\mu} g^{\mu}_{w} Q_{\mu}(X) \\ &= \sum_{a \in R(w)} \Theta^{\ell(w)}_{peaks(a)}(X) \\ &= \sum_{a \in R(w)} \sum_{(i_1 \leq \cdots \leq i_d) \in A(P)} 2^{|\{i_1, i_2, \ldots, i_p\}|} x_{i_1} x_{i_2} \cdots x_{i_d}. \end{aligned}$$

Every Schur Q-function is itself a type C Stanley symmetric function. In particular, for the shifted partition $\mu = (\mu_1 > \mu_2 > \cdots > \mu_k > 0)$, we can construct an increasing signed permutation $w(\mu)$ in one-line notation starting with the negative values $\bar{\mu_1}, \bar{\mu_2}, \ldots, \bar{\mu_k}$ and ending with the positive integers in the complement of the set $\{\mu_1, \mu_2, \ldots, \mu_k\}$. For example, if $\mu = (5, 3, 1)$ then $w(\mu) = (\bar{5}, \bar{3}, \bar{1}, 2, 4)$. Then by [Billey and Haiman, 1995, Thm.3],

$$F^{C}_{w(\mu)}(X) = Q_{\mu}(X). \tag{4.2.3}$$

Conversely, every increasing signed permutation w gives rise to an F^{C}_{w} which is a single Schur Q-function defined by the negative numbers in $[w(1), \ldots, w(n)]$.

Theorem 4.2.15. *[Billey, 1998, Cor. 9] Let w be a signed permutation which is not increasing. Then we have the following transition equation*

$$F_w^C(X) = \sum_{w' \in T(w)} F_{w'}^C(X). \tag{4.2.4}$$

This expansion terminates in a finite number of steps as a sum with all terms indexed by increasing signed permutations.

Corollary 4.2.16. *Let w be a signed permutation that is not increasing. Then*

$$|R(w)| = \sum_{w' \in T(w)} |R(w')|.$$

Proof. Consider the coefficient of $x_1 x_2 \cdots x_{\ell(w)}$ in $F_w^C(x)$ and the right hand side of (4.2.4). $\square$

4.3 Pushes, Bumps, and the signed Little Bijection

In this section, we define the signed Little map on reduced words via two other algorithms called *push* and *bump*. A key tool is the wiring diagrams for reduced words of signed permutations. The main theorem proved in this section is Theorem 4.1.1, which says that the Little bumps determine a bijection on reduced words that realizes the transition equation for type C Stanley symmetric functions.

The *wiring diagram* of $a = a_1 a_2 \ldots a_p$ is the array $[p] \times [-n, n]$ in Cartesian coordinates with a cross in entry (j, a_j), denoted $\times$, and a cross in entry $(j, -a_j)$ for all $1 \leq i \leq p$. All other entries contain a horizontal line, denoted _, along the bottom of the entry. If $a_j = 0$, there will be just one cross in column j. The line segments connect to form "wires" with labels $[-n] = \{-1, -2, \ldots, -n\}$ and $[n] = \{1, 2, \ldots, n\}$ starting on the right hand edge of the diagram. So, if $a_p > 0$ then wires a_p and $a_p + 1$ cross in column p, and also wires $-a_p$ and $-a_p - 1$ cross. If $a_p = 0$, then wires 1 and -1 cross in column p. For each $0 \leq k < p$, define $w^k(a)$ to be $w^k(a) = s_{a_p} \ldots s_{a_{k+1}} s_{a_k}$ where $w^0(a) = w^{-1}$ and $w^p(a)$ is the identity. The sequence of wire labels along the left edge from bottom to top gives the long form of the signed permutation w^{-1}. More generally, the sequence of labels on the wires of the wiring diagram just before column k is $w^k(a)$. Every wiring diagram should be considered as a subdiagram of the diagram with wires labeled by all of $\mathbb{Z} - \{0\}$ where all constant trajectories are suppressed. See Figure 4.3 for several examples.

The *inversion set* $\mathrm{Inv}_B(w)$ of a signed permutation w is

$$\mathrm{Inv}_B(w) = \{(i,j) \in ([-n] \cup [n]) \times [n] : \ |i| \leq j \text{ and } w(i) > w(j)\}.$$

We have defined the wiring diagrams so that the inversion (i,j) corresponds with the crossing of wires i and j in any wiring diagram of a reduced word for w. Note, the wires $-j$ and $-i$ also cross in the same column in such a diagram. Thus, it is equivalent to refer to the inversion (i,j) by $(-j,-i)$. If $a \in R(w)$, then the wiring diagram for a is *reduced* and every crossing corresponds to an inversion for w.

For a word $a = a_1 \dots a_p$ and $\delta \in \{-1,1\}$, we define a *push* P_i^δ at index i to be the map that adds δ to a_i while fixing the rest of the word. We will write P_k^- and P_k^+ for $\delta = -1$ and $\delta = 1$ respectively. The effect pushes have on wiring diagrams can be observed in Figure 4.3.

If $a = a_1 \dots a_p$ is a word that is not reduced, we say a *defect* is caused by a_i and a_j with $i \neq j$ if the removal of either leaves a reduced word. The following lemma can be deduced for signed permutations from the wiring diagrams, but it holds more generally for Coxeter groups.

Lemma 4.3.1 (Lemma 21, [Lam and Shimozono, 2005/07]). *For W a Coxeter group and $w \in W$, let $a = a_1 \dots \hat{a_i} \dots a_p \in R(w)$ such that $a_1 \dots a_p$ is not reduced. Then there exists a unique $j \neq i$ such that $a_1 \dots \hat{a_j} \dots a_p$ is reduced. Moreover, $a_1 \dots \hat{a_j} \dots a_p \in R(w)$.*

Definition 4.3.2 (**Little Bump Algorithm**). Let $a = a_1 \dots a_p$ be a reduced word of the signed permutation w and $(i,j) \in \mathrm{Inv}_B(w)$ such that $\ell(wt_{ij}) = \ell(w) - 1$. Fix $\delta \in \{-1,1\}$. We define the *Little bump for w* at the inversion (i,j) in the direction δ, denoted $B_{(i,j)}^\delta$, as follows

Step 1: Identify the column k and row r containing the wire crossing (i,j) with $i < j$. Either $w^{k-1}(a_k) = i$ and $w^{k-1}(a_k+1) = j$ or $w^{k-1}(a_k) = -j$ and $w^{k-1}(a_k+1) = -i$. If $w^{k-1}(a_k) = i$, set $b = P_k^\delta(a)$. Otherwise $w^{k-1}(a_k) = -j$, and we set $\delta := -\delta$, and

$b = P_k^\delta(a)$. In either case, set $r := r + \delta$. Note that the order in which the variables are updated matters. Let $(x < y)$ be the new wires crossing in column k and row r.

Step 2: If b is reduced, return b. Otherwise, by Lemma 4.3.1 there is a unique defect caused by b_k and some b_l with $l \neq k$. If $b_l > 0$, then either $w^{l-1}(b_l) \in \{x, y\}$ or $w^{l-1}(b_l) \in \{-x, -y\}$. If $b_l = 0$, then $x = -y$ and $w^{l-1}(1) \in \{-x, x\}$.

- If $b_l > 0$ and $w^{l-1}(b_l) \in \{x, y\}$, set $r := r + \delta$, $k := l$, and $b := P_k^\delta(b)$. After updating the variables, let $(x < y)$ to be the wires crossing in the diagram for b in column k and row r. Repeat Step 2.
- Otherwise, $w^{l-1}(b_l) \in \{-x, -y\}$. Set $\delta := -\delta$, $r = r + \delta$, $k := l$ and $b := P^{\delta(b)}$. Again, the order matters. After updating the variables, let $(x < y)$ to be the wires crossing in column k and row r. Repeat Step 2.

Figure 4.3 shows each step of a Little bump in terms of wiring diagrams. The corresponding effect on reduced words can be read off the diagrams by noting the row numbers of the wire crossings in the upper half plane including the x-axis.

Remark 4.3.3. The Little bump algorithm is best thought of as acting on wiring diagrams. At every step, the pushes move the (x, y) crossings consistently in the initial direction of δ. In the first step, we move the (i, j)-crossing in the wiring diagram up if $\delta = +1$ and down if $\delta = -1$. If wires i and j cross in the upper half plane then the swap a_k is replaced with $a_k + \delta$. However, if wires i, j cross in the lower half plane then the swap a_k is replaced with $a_k - \delta$ and the sign of δ is switched. If a new defect crossing is later found on the other side of the x-axis from the last crossing, then the sign of δ will switch again so that the crossing continues to move in the same direction. Thus, if the initial push moved (i, j) down, each subsequent iteration will continue to move a crossing down, but the effect on the word from the corresponding push can vary.

Remark 4.3.4. Observe that in each iteration of Step 2, the word b has the property that its subword $b_1 b_2 \ldots \widehat{b_k} \ldots b_p$ is reduced.

When analyzing Little bumps and pushes, we will need to track where the next defect can occur. Given the wiring diagram for a word $b = b_1 \dots b_p$, not necessarily reduced, and a crossing (x, y) in the diagram, define the *(lower) boundary* of b for crossing (x, y), denoted $\partial^k_{(x,y)}(b)$, to be the union of the trajectory of x from columns 0 to k and the trajectory of y from columns k to p. Note using the notation of Step 2 above, any defect found in this iteration will occur along $\partial^k_{(x,y)}(b)$. In Figure 4.3, the boundary of each crossing that will be pushed is dashed. A similar concept of an upper boundary could be defined if the initial step pushes the (i, j) cross up.

Lemma 4.3.5. *Let $a \in R(w)$ and $B^{\delta}_{(i,j)}$ be a Little bump for w consisting of the sequence of pushes $P^{\delta_1}_{t_1}, P^{\delta_2}_{t_2}, \dots$ acting on a. Then for all k and $\delta \in \{-1, 1\}$ the push P^{δ}_k appears at most once in this sequence. Hence, the Little bump algorithm terminates in at most $2\ell(w)$ pushes.*

This proof is a slight extension of the proof of Lemma 5 in [Little, 2003].

Proof. Let $a = a_1 a_2 \dots a_p \in R(w)$ and a_k denote the swap introducing the inversion (i, j), with $i < j$. Since $B^-_{(i,j)} = B^+_{(-j,-i)}$, we need only demonstrate the result when the algorithm starts with a push that moves the (i, j)-crossing down.

In Step 1 of the Little bump algorithm, either $b = P^{\delta}_k(a)$ is reduced or there is some $l \neq k$ such that b_k and b_l cause a defect. Suppose the latter case holds. Then b_k and b_l swap the wire i with some wire $h \neq i, j$. By considering the reverse of the word if necessary, we may assume $k < l$. The defect in column l must occur on the boundary of $\partial^k_{(x,y)}(b)$. Observe that $\partial^k_{(i,j)}(a)$ and $\partial^l_{(i,h)}(b)$ coincide from 1 to k and from l to p. Moreover, between k and l, the boundary $\partial^l_{(i,h)}(b)$ is strictly lower in the wiring diagram than $\partial^k_{(i,j)}(a)$. This can be seen in the first two diagrams shown in Figure 4.3. Observe that the trajectories of $-i$ and $-j$ will not interact with $\partial^l_{(i,h)}(b)$ unless $j = -i$. Therefore the boundary $\partial^l_{(i,h)}(b)$ is weakly below the boundary $\partial^k_{(i,j)}(a)$. Similar reasoning show that on each iteration of Step 2 in the algorithm, the boundary always moves weakly down provided the initial push moves a crossing down.

Now we can verify that no push occurs twice in the Little bump algorithm. In particular, we claim that both P^-_k and P^+_k can occur, but they are never repeated. For example, P^-_5

Figure 4.3: The sequence of pushes corresponding to the Little bump $B^-_{(-2,1)}$ as applied to $a = 1021201 \in R([\bar{1}, 3, 2])$. The boundary of each crossing about to be moved is dashed. Here, red is negative and blue is positive.

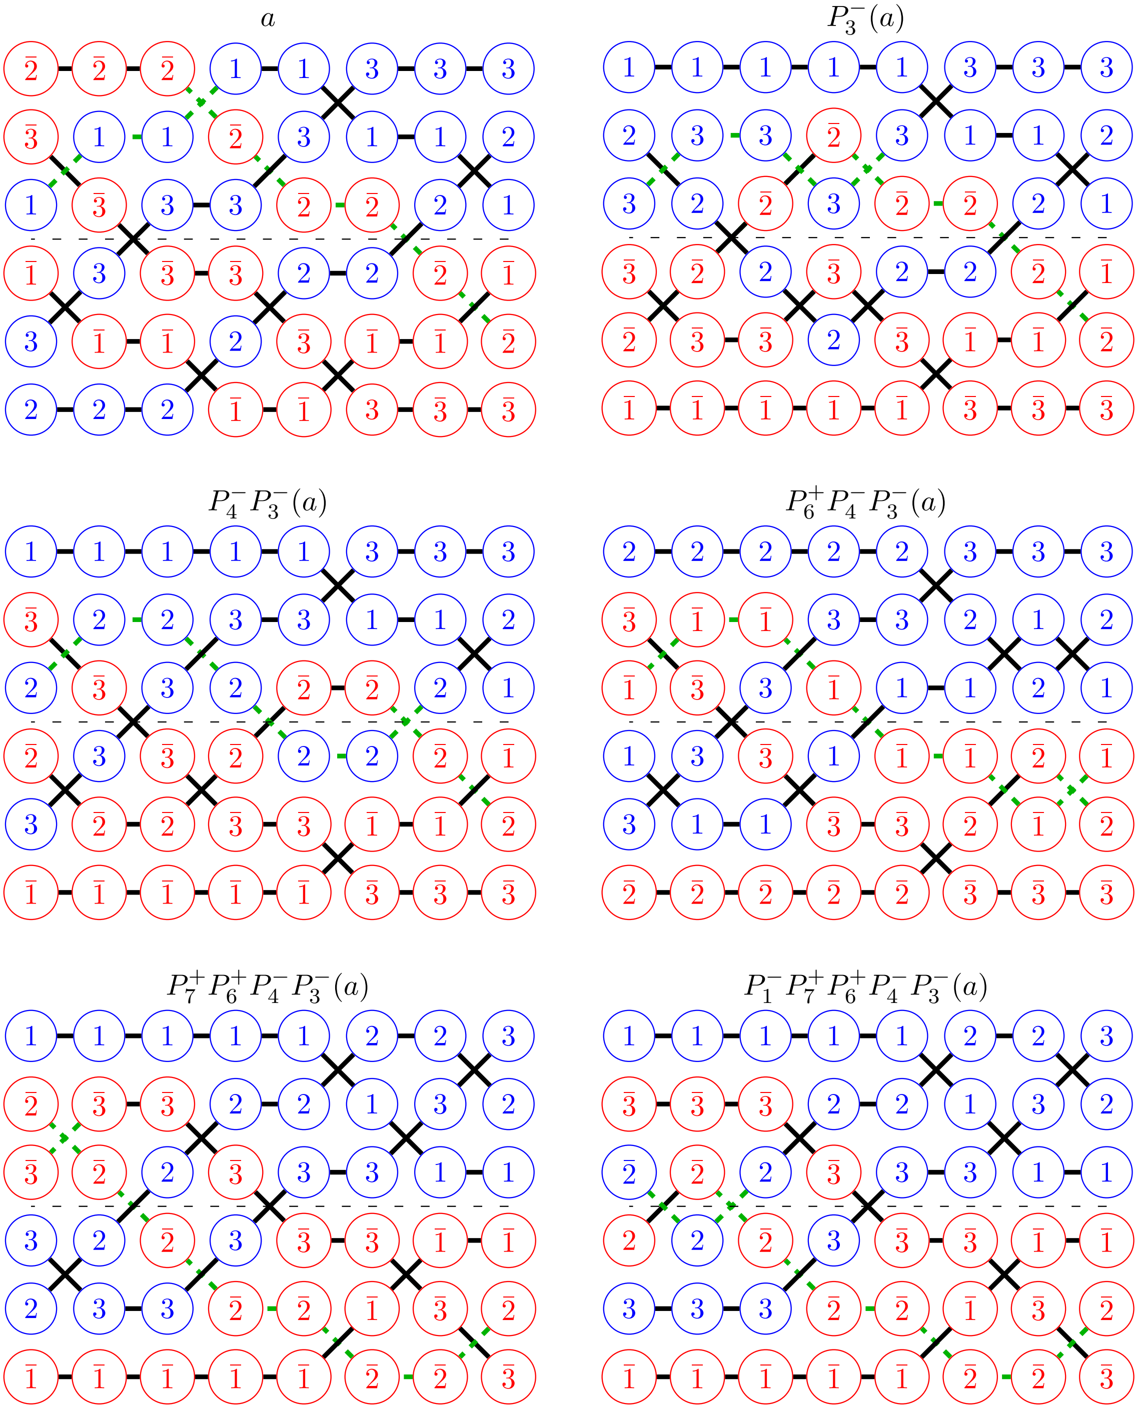

a

$P_3^-(a)$

$P_4^- P_3^-(a)$

$P_6^+ P_4^- P_3^-(a)$

$P_7^+ P_6^+ P_4^- P_3^-(a)$

$P_1^- P_7^+ P_6^+ P_4^- P_3^-(a)$

$P_2^+ P_1^- P_7^+ P_6^+ P_4^- P_3^-(a)$

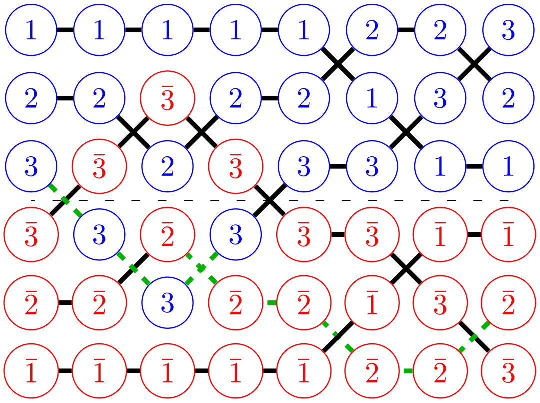

$P_3^+ P_2^+ P_1^- P_7^+ P_6^+ P_4^- P_3^-(a)$

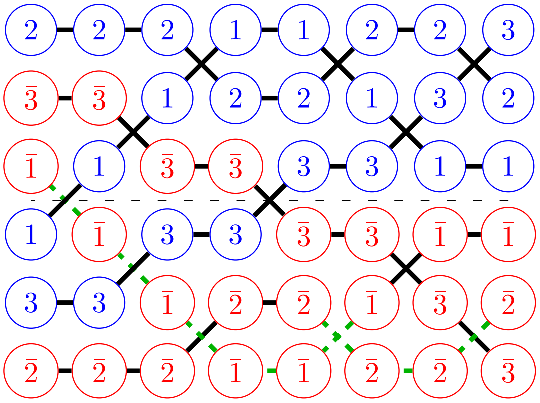

$B^-_{(-2,1)}(a) = P_5^+ P_3^+ P_2^+ P_1^- P_7^+ P_6^+ P_4^- P_3^-(a)$

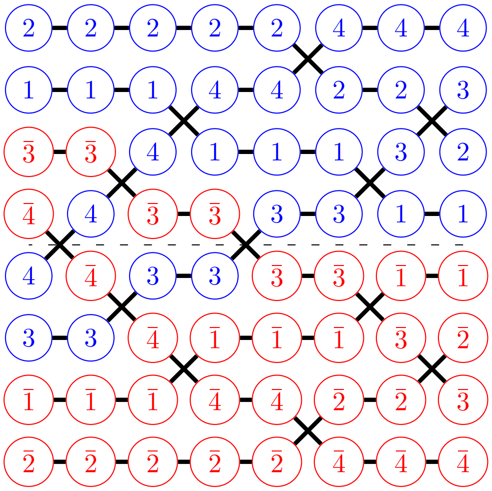

and P_5^+ both occur in Figure 4.3. To do this, we need to examine the argument above more closely. Assume that Step 2 starts with a push P_k^δ moving a crossing into position (k, r) in the wiring diagram of b. Assume this wiring diagram has a defect in columns k and l and $k < l$. Then the boundary before and after the push P_k^δ agree weakly to the left of column k. If successive pushes occur strictly to the right of column k, then none of these pushes will repeat P_k^δ. Furthermore, the boundary to the left of column k will be constant. The first time another iteration of Step 2 finds a defect weakly to the left of column k, we claim it must occur in a column strictly to the left of column k, thus the boundary moves strictly below (k, r). The reason column k cannot be part of the defect this time is that the boundary has negative slope just to the right of the crossing in row r column k, but to create a defect with a string passing below (k, r) the boundary must have positive slope where the two strings meet to the right of column k. Furthermore, if another push in column k occurs later in the algorithm, it must be on the other side of the x-axis so it must be $P_k^{-\delta}$ since the boundary moves monotonically. Once the boundary has moved beyond both crossings in column k, neither crossing will be pushed again, so P_k^δ occurs at most once in the Little Bump algorithm. $\square$

Lemma 4.3.6. *Let w be a signed permutation, $a = a_1 \dots a_p \in R(w)$, and $B_{(i,j)}^\delta$ be a Little bump for w. Then*

$$peaks(a) = peaks(B_{(i,j)}^\delta(a)).$$

Proof. Say $i \in peaks(a)$, so $a_{i-1} < a_i > a_{i+1}$. The statement holds unless a push applied to one of these entries during the algorithm leaves an equality in the resulting word b, say $b_{i-1} = b_i$. In this case, there is a defect caused by b_{i-1} and b_i, so we push the other next. The direction of the new push for defects caused by adjacent entries will be the same unless $b_{i-1} = b_i = 0$. This cannot occur since $a_i > a_{i-1} \geq 0$ and if $a_i = 1$, then $a_{i-1} = 0 = a_{i+1}$. Hence, $a_1 \dots \hat{a_i} \dots a_p$ would not be reduced, which is not possible by Remark 4.3.4. $\square$

Theorem 4.3.7 (Theorem 4.1.2(1))**.** *Let $x \in B_\infty$, and let $a \in R(x)$. Say $j < k$.*

1. *If $B_{(j,k)}^-$ is a Little bump for x, then $B_{(j,k)}^-(a) \in R(xt_{jk}t_{ij})$ for some $i < j$.*

2. If $B^+_{(j,k)}$ is a Little bump for x, then $B^-_{(j,k)}(a) \in R(xt_{jk}t_{kl})$ for some $k < l$.

Proof. When applying $B^-_{(j,k)}$ on a, the initial push is some $P^{\delta_1}_{l_1}$ with $a_1 \dots \widehat{a_{l_1}} \dots a_p \in R(v)$. Let $b = P^{\delta_1}_{l_1}(a)$. Then, one can observe from the wiring diagrams that $s_{b_1} \cdots s_{b_p} = vt_{ji_1}$ for some $i_1 \neq j, k$.

If b is reduced, the bump is done. Otherwise, by the Little bump algorithm, there is some unique defect between b_{l_1} and b_{l_2} so we push next in column l_2. We know $b_1 \dots \widehat{b_{l_2}} \dots b_p \in R(v)$ by construction and Remark 4.3.4. So when the next push occurs in column l_2 the new crossing will be between j and another string i_2. Continuing the algorithm, we see recursively that $B^-_{(j,k)}(a) \in R(vt_{ji})$ for some $i \neq j, k$.

Assume for the sake of contradiction that $i > j$, and say the (i,j)-crossing in the wiring diagram of $c = B^-_{(j,k)}(a)$ occurs in column l. By removing the l^{th} swap from c we get a wiring diagram for v that does not have (j,i) as an inversion. Thus, the i-wire must stay entirely above the j-wire. Hence, the i wire is above the boundary of the last push. Thus, it cannot be a part of the last push since the boundary moves monotonically according to the proof of Lemma 4.3.5. We can then conclude that $c \in R(vt_{ij})$ for some $i < j$.

A similar proof holds for the second statement. □

Recall the notation of transition equations from Section 4.1. If w is not increasing, let $(r < s)$ be the lexicographically largest pair of positive integers such that $w_r > w_s$. Set $v = wt_{rs}$. Let $T(w)$ be the set of all signed permutations $w' = vt_{ir}$ for $i < r$, $i \neq 0$ such that $\ell(w') = \ell(w)$.

Next, we show that the *canonical* Little bump $B^-_{(r,s)}$ for w respects the transition equations in Theorem 4.2.15. This is best done by describing the domain and range of Little bumps in greater generality. For $v \in B_\infty$ and $j \in \mathbb{Z} - \{0\}$, we define

$$D(v,j) = \{vt_{ij} : \ i < j, \ i \neq 0 \text{ and } \ell(vt_{ij}) = \ell(v) + 1\}$$
$$U(v,j) = \{vt_{jk} : \ j < k, \ i \neq 0 \text{ and } \ell(vt_{jk}) = \ell(v) + 1\}.$$

Observe that we have $D(v,-j) = U(v,j)$. We now prove the analogue of [Little, 2003, Theorem 3], from which we can deduce Theorem 4.1.1.

Lemma 4.3.8. *Let $v \in B_\infty$ and $j \neq 0$. Then*

$$\sum_{y \in D(v,j)} |R(y)| = \sum_{x \in U(v,j)} |R(x)|.$$

Proof. We will prove the equality bijectively by using a collection of Little bumps. Define a map $M_{w,j}$ on $\cup_{x \in U(v,j)} R(x)$ as follows. Say $a = a_1 \dots a_p \in R(x)$ for some $x \in U(v,j)$. Then $x = vt_{jk}$ for some unique $k > j$. Furthermore, $B^-_{(j,k)}$ is a Little bump for x. By Theorem 4.3.7, we know that $B^-_{(j,k)}(a) \in R(vt_{ij})$ for some $i < j$ and $\ell(vt_{ij}) = \ell(vt_{jk})$. Thus, $vt_{ij} \in D(v,j)$. Set $M_{w,j}(a) := B^-_{(j,k)}(a)$ for all $a \in R(x)$. In this way, we construct a map

$$M_{w,j} : \cup_{x \in U(v,j)} R(x) \longrightarrow \cup_{y \in D(y,j)} R(y).$$

Since the Little bump algorithm is reversible with $B^+_{(i,j)}(c) = a$ in the notation above, we know $M_{w,j}$ is injective.

The bijective proof is completed by observing that $D(v,j) = U(v,-j)$, $U(v,j) = D(v,-j)$, and that $B^-_{(-j,-i)} = B^+_{(i,j)}$ is a Little bump for $vt_{ij} \in D(v,j)$ whose image, by the above argument, is a reduced word of some $x \in U(v,j)$. □

Corollary 4.3.9 (Theorem 4.1.1)**.** *Let $a \in R(w)$ and $B^-_{(r,s)}$ be the canonical Little bump for w. Recall that (r,s) is the lexicographically last inversion in w. Then $B^-_{(r,s)}(a)$ is a reduced word for w' where*

$$w' \in T(w) = \{wt_{rs}t_{lr} \;\mid\; l < r \text{ and } \ell(w) = \ell(wt_{rs}t_{rl})\}.$$

Proof. Observe $U(wt_{rs}, r) = \{w\}$ and

$$D(wt_{rs}, r) = \{wt_{rs}t_{lr} \;\mid\; l < r,\ l \neq 0, \text{ and } \ell(w) = \ell(wt_{rs}t_{rl})\} = T(w).$$

The result now follows from Lemma 4.3.8. □

4.4 Kraśkiewicz insertion and the signed Little Bijection

In this section, we show that Coxeter-Knuth moves act on $Q'(a)$ by shifted dual equivalence, as defined in [Haiman, 1992]. We then prove Theorem 4.1.2 by applying properties of shifted dual equivalence and showing that Little bumps and Coxeter-Knuth moves commute on reduced words of signed permutations.

For a permutation $\pi \in S_n$, let $\pi|_I$ be the subword consisting of values in the interval I. Let $\mathrm{fl}(\pi|_I) \in S_{|I|}$ be the permutation with the same relative order as $\pi|_I$. Here fl is the *flattening operator*. Similarly, for Q a standard shifted tableau $Q|_I$ denotes the shifted skew tableau obtained by restricting the tableau to the cells with values in the interval I.

Definition 4.4.1. [Haiman, 1992] Given a permutation $\pi \in S_n$, define the *elementary shifted dual equivalence* h_i for all $1 \leq i \leq n-3$ as follows. If $n \leq 3$, then $h_1(\pi) = \pi$. If $n = 4$, then $h_1(\pi)$ acts by swapping x and y in the cases below,

$$1x2y \quad x12y \quad 1x4y \quad x14y \quad 4x1y \quad x41y \quad 4x3y \quad x43y, \tag{4.4.1}$$

and $h_1(\pi) = \pi$ otherwise. If $n > 4$, then h_i is the involution that fixes values not in $I = \{i, i+1, i+2, i+3\}$ and permutes the values in I via $\mathrm{fl}(h_i(\pi)|_I) = h_1(\mathrm{fl}(\pi|_I))$.

As an example, $h_1(24531) = 14532$, $h_2(25134) = 24135$, and $h_3(314526) = 314526$.

Recall from Definition 4.2.9 that a type B Coxeter-Knuth move starting at position i is denoted by β_i. One can verify that this definition is equivalent to defining h_i as

$$h_i(\pi) = (\beta_i(\pi^{-1}))^{-1}.$$

Given a standard shifted tableau T, we define $h_i(T)$ as the result of letting h_i act on the row reading word of T. Observe $h_i(T)$ is also a standard tableau. We can define an equivalence relation on standard shifted tableaux by saying T and $h_i(T)$ are *shifted dual equivalent* for all i.

Theorem 4.4.2 ([Haiman, 1992, Prop. 2.4])**.** *Two standard shifted tableaux are shifted dual equivalent if and only if they have the same shape.*

Recall the notion of jeu de taquin is an algorithm for sliding one cell at a time in a standard tableau on a skew shape in such a way that the result is still a standard tableau [Sagan, 1991]. The analogous notion for shifted tableaux was introduced independently in [Sagan, 1987] and [Worley, 1984].

Lemma 4.4.3 (Lemma 2.3, [Haiman, 1992]). *Given two standard shifted tableaux T and U with $T = h_i(U)$, let T' and U' be the result of applying any fixed sequence of jeu de taquin slides to T and U, respectively. Then $T' = h_i(U')$.*

Definition 4.4.4. Given a standard shifted tableau Q', define $\Delta(Q')$ as the result of removing the cell containing 1, performing jeu de taquin into this now empty cell, and subtracting 1 from the value of each of the cells in the resulting tableau.

Lemma 4.4.5. *[Lam, 1995, Theorem 3.24] Let w be a signed permutation and $a = a_1 \cdots a_p \in R(w)$. Then under Kraśkiewicz insertion*

$$Q'(a_2 \cdots a_p) = \Delta(Q'(a_1 \cdots a_p)). \tag{4.4.2}$$

Lemma 4.4.6. *Let w be a signed permutation, and let $a = a_1 \ldots a_p \in R(w)$. Then $Q(\beta_i(a)) = h_i(Q(a))$ for all integers $1 \leq i \leq p - 3$.*

Proof. Recall that β_i acts trivially on a if and only if $i+1, i+2 \notin peaks(a)$. Similarly, h_i acts trivially on $Q'(a)$ if and only if $i+1, i+2 \notin peaks(Q'(a))$. By Theorem 4.2.8, we then see β_i acts trivially if and only if h_i acts trivially.

Since type B Coxeter-Knuth moves preserve Kraśkiewicz insertion tableaux, we see $Q'(a)|_{[1,i-1]} = Q'(\beta_i(a))|_{[1,i-1]}$, $Q'(a)|_{[i+4,p]} = Q'(\beta_i(a))|_{[i+4,p]}$ and that the shape of $Q'(a)|_{[i,i+3]}$ and $Q'(\beta_i(a))|_{[i,i+3]}$ are the same. In particular, $Q'(a)$ differs from $Q'(\beta_i(a))$ by some rearrangement of the values in $[i, i+3]$. We need to show that this rearrangement is the elementary shifted dual equivalence h_i. The following proof of this fact is presented as a commuting diagram in Figure 4.4.

By omitting any extra values at the end of a, we may assume that $p = i+3$. Now consider the tableaux T and U obtained by adding $i-1$ to each entry in $\Delta^{i-1}(Q'(a))$ and $\Delta^{i-1}(Q'(\beta_i(a))$. Because $Q'(a)|_{[1,i-1]} = Q'(\beta_i(a))|_{[1,i-1]}$, it follows from the definition of Δ

that there is some fixed set of jeu de taquin slides that relates both $Q'(a)|_{[i,i+3]}$ to T and $Q'(\beta(a))|_{[i,i+3]}$ to U. Applying Lemma 4.4.3, we need only show that $T = h_i(U)$ to complete the proof.

By Lemma 4.4.5, we see

$$\Delta^{i-1}(Q'(a)) = Q'(a_i a_{i+1} a_{i+2} a_{i+3})$$

and

$$\Delta^{i-1}(Q'(\beta_i(a))) = Q'(\beta_1(a_i a_{i+1} a_{i+2} a_{i+3})).$$

Since $Q'(\beta_1(a_1a_2a_3a_4))$ and $Q'(a_1a_2a_3a_4)$ are distinct standard tableaux of the same shifted shape with four cells, the shape must be $(3,1)$. Furthermore, there are only two standard tableaux of shifted shape $(3,1)$, so the two tableaux must be related by $Q'(\beta_1(a_1a_2a_3a_4)) = h_1(Q'(a_1a_2a_3a_4))$. Adding $i-1$ to each entry of the two tableaux in this equation changes h_1 to h_i and yields the desired result, $T = h_i(U)$. □

Next, we show that Coxeter-Knuth moves commute with Little bumps.

Lemma 4.4.7. *Let $a = a_1 \dots a_p$ be a reduced word of the signed permutation w, β_k a Coxeter-Knuth move for a and $B^{\delta}_{(i,j)}$ a Little bump for w. Then*

$$B^{\delta}_{(i,j)}(\beta_k(a)) = \beta_k(B^{\delta}_{(i,j)}(a)).$$

Proof. First, observe that a Little bump and the Coxeter-Knuth move β_k will only interact if one of the pushes in the bump is applied to an entry in the window $[k, k+3]$. The inversions introduced by a_i and $\beta_k(a)_i$ are the same when a_i is not acted on by β_k. Therefore, since a and $\beta_k(a)$ are reduced words of the same permutation, we see the inversions introduced by $a_k a_{k+1} a_{k+2} a_{k+3}$ in a and $\beta_k(a_k a_{k+1} a_{k+2} a_{k+3})$ in $\beta_k(a)$ are the same as well. Therefore if an entry with index in $[k, k+3]$ is pushed when a Little bump $B^{\delta}_{(i,j)}$ is applied to a, such an index will also be pushed when $B^{\delta}_{(i,j)}$ is applied to $\beta_k(a)$, though not necessarily the same index. Our argument relies on showing commutation can be reduced to a local check of how β_k interacts with $B^{\delta}_{(i,j)}$. In particular, the result will follow from establishing two properties:

$$\begin{tikzcd}
 & Q'(a)|_{[i,i+3]} \arrow[r, leftrightarrow, dotted, "h_i"] \arrow[dl, "\Delta^{i-1}"'] \arrow[d, leftrightarrow, "j.d.t."] & Q'(\beta_i(a))|_{[i,i+3]} \arrow[d, leftrightarrow, "j.d.t."] \arrow[dr, "\Delta^{i-1}"] & \\
Q'(a_i \cdots a_{i+3}) \arrow[r, "+(i-1)"] \arrow[rrr, leftrightarrow, bend right, "h_1"'] & T \arrow[r, leftrightarrow, "h_i"] & U & Q'(\beta_1(a_i \cdots a_{i+3})) \arrow[l, "+(i-1)"']
\end{tikzcd}$$

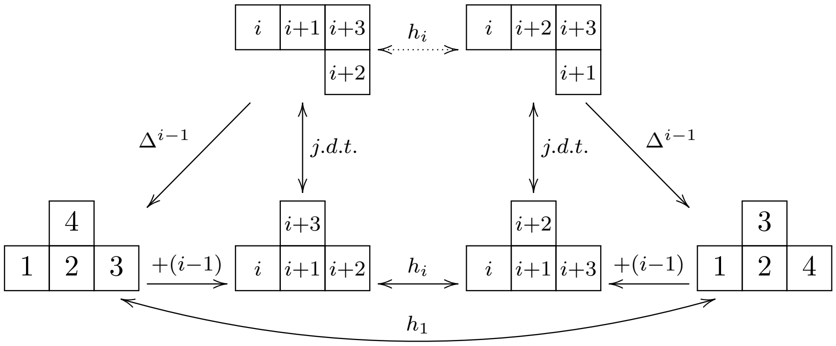

Figure 4.4: The commuting relationships in the proof of Proposition 4.4.6 on top and a generic example on bottom.

1. When the pushes for the Little bump algorithm are applied to swaps acted on by the Coxeter-Knuth move terminate, the resulting subwords differ by a Coxeter-Knuth move.

2. The final push to an swap acted on by the Coxeter-Knuth move has the same effect on w for both a and $\beta_k(a)$, hence would introduce the same defect should the bump continue.

When the entries acted on by a Coxeter-Knuth move differ by two or more, these properties are trivial to confirm. For entries that are closer, the features can be checked for each type of Coxeter-Knuth move either by hand or by computer program. There are as many as four checks for each type of Coxeter-Knuth move, depending on the initial inversion and whether 0 appears in the word. These can be performed by verifying the result for all possible bumps on reduced words in B_5 of length 4. □

Figure 4.5: $B^{-}_{(1,2)}(1010)$ compared to $B^{-}_{(1,2)}(0101)$. The results differ by a Coxeter-Knuth move and the final pushes introduce the same transposition.

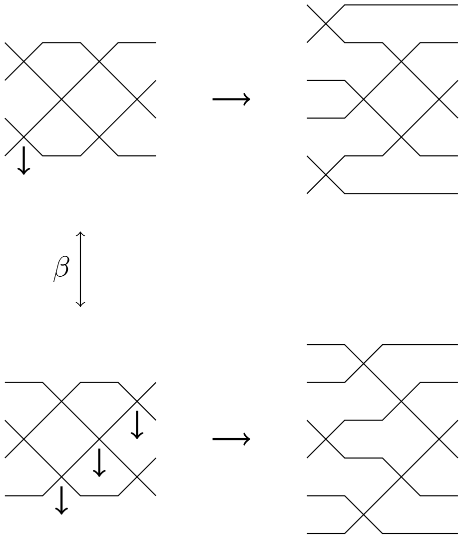

Notice that if we weakly order all shifted standard tableaux of shape λ by their peak sets in lexicographical order, then the unique maximal element U_λ is obtained by placing 1 through λ_1 in the first row, $\lambda_1 + 1$ through $\lambda_1 + \lambda_2$ in the second row, and so on. Further notice that $peaks(U_\lambda) = \{\lambda_1, \lambda_1 + \lambda_2, \lambda_1 + \lambda_2 + \lambda_3, \ldots\}$.

Theorem 4.4.8 (Theorem 4.1.2 (2))**.** *Let w be a signed permutation, $a \in R(w)$, $B^{\delta}_{(i,j)}$ be a Little bump for w and $b = B^{\delta}_{(i,j)}(a)$. Then $Q'(a) = Q'(b)$.*

Proof. We first show that $Q'(a)$ and $Q'(b)$ have the same shape. By Lemma 4.2.8 and Lemma 4.3.6,

$$peaks(Q'(a)) = peaks(a) = peaks(b) = peaks(Q'(b)).$$

Let a' and b' be the reduced words with maximal peak sets in the Coxeter-Knuth class of a and b, respectively. Applying Lemma 4.4.7, we may assume that $a = a'$. The shape of $Q'(a') = U_\lambda$ is determined by its peak set. Hence, the shape of $Q'(b)$ must be at least as large as the shape of $Q'(a)$ in dominance order. By assuming that $b = b'$, we can conclude the converse. Hence, $Q'(a)$ and $Q'(b)$ have the same shape λ. Furthermore, $Q'(a') = U_\lambda = Q'(b')$.

We now proceed to showing that $Q'(a) = Q'(b)$. By Theorem 4.2.10, there exists a sequence of Coxeter-Knuth moves $\beta = \beta_{i_1} \circ \beta_{i_2} \circ \cdots \circ \beta_{i_k}$ such that $\beta(a') = a$. From Lemma 4.4.6, we see

$$Q'(a) = Q'(\beta(a')) = h_{i_1} \dots h_{i_k}(U_\lambda).$$

Applying Lemma 4.4.7, $\beta(b') = b$, so

$$Q'(b) = Q'(\beta(b')) = h_{i_1} \dots h_{i_k}(U_\lambda),$$

from which we conclude that $Q'(a) = Q'(b) = Q'(B^{\delta}_{(i,j)}(a))$. □

As a consequence of Theorem 4.4.8, we prove an analogue of Thomas Lam's conjecture for signed permutations. Two reduced words a and b *communicate* if there exists a sequence of Little bumps $B_1^{\delta_1}, B_2^{\delta_2}, \dots, B_n^{\delta_n}$ such that $b = B_n^{\delta_n}(\dots B_1^{\delta_1}(a))$. Since Little bumps are invertible, this defines an equivalence relation.

Theorem 4.4.9. *Let a and b be reduced words. Then $Q'(a) = Q'(b)$ if and only if they communicate via Little bumps.*

Proof. If a and b communicate, then we see $Q'(a) = Q'(b)$ by Theorem 4.4.8. Therefore we only need to prove the converse.

Let $Q = Q'(a) = Q'(b)$. We show that a and b both communicate with some reduced word c uniquely determined by Q. Since communication is an equivalence relation, this will complete our proof. Recall from Theorem 4.2.15 that by repeated application of the transition equations, we may express any C-Stanley symmetric function as the sum of C-Stanley symmetric functions of increasing permutations. Since canonical Little bumps follow the transition equations by Corollary 4.3.9, repeated applications of canonical Little bumps will transform any reduced word a into some reduced word c of an increasing permutation u. Since a and c communicate, $Q'(a) = Q'(c)$. By Equation (4.2.3) and the fact that u is increasing, $F_u(X) = Q_\mu(X)$ for some μ and the reduced expressions for u are in bijection with the standard tableaux of shifted shape μ under Kraśkiewicz insertion. This implies that c is uniquely determined by $Q'(a)$, and hence for $Q'(b)$ as well. Therefore, every reduced word a with $Q'(a) = Q$ communicates with the same word $c \in R(u)$. □

Corollary 4.4.10. *Every communication class under signed Little bumps has a unique reduced word for an increasing signed permutation.*

For permutations, Theorem 3.32 in [Lam, 1995] shows that Kraśkiewicz insertion coincides with Haiman's shifted mixed insertion. From this, we can conclude the following.

Corollary 4.4.11. *Let* $\{j_1 < j_2 < ... < j_p\}$ *be increasing sequence of* p *distinct non-negative integers. Every communication class containing words of length* p *under signed Little bumps contains a reduced word that is a permutation of* $\{j_1, j_2, ..., j_p\}$*.*

This result can also be proved using Little bumps.

4.5 Axioms for shifted dual equivalence graphs

In this section, we build on the connection between shifted dual equivalence operators h_i and type B Coxeter-Knuth moves β_i as stated in Lemma 4.4.6. In particular, we define and classify the shifted dual equivalence graphs associated to these operators via two local properties. Along the way, we also demonstrate several important properties of these graphs. The approach is analogous to the axiomatization of dual equivalence graphs by Assaf [Assaf, 2013] and refined by Roberts [Roberts, 2014].

Definition 4.5.1. Fix a strict partition $\lambda \vdash n$. By definition, h_i acts as an involution on the standard shifted tableaux of shape λ, denoted $SST(\lambda)$. Given λ, define the *standard shifted dual equivalence graph of degree* n for λ, denoted

$$\mathcal{SG}_\lambda = (V, \sigma, E_1 \cup \ldots \cup E_{n-3})$$

as follows. The vertices $V = SST(\lambda)$ and the labeled edge sets E_i for $1 \leq i \leq n-3$ are given by the nontrivial orbits of h_i on $SST(\lambda)$. To define σ, recall from Section 4.2 that every tableau $T \in SST(\lambda)$ has a peak set, denoted $peaks(T)$. We encode a peak set by a sequence of pluses and minus denoted $\sigma(T) \in \{+,-\}^n$, where $\sigma_i(T) = +$ if and only if i is a peak in T. We refer to $\sigma(T)$ as the *signature* of T. Note that peaks never occur in positions 1 or n and they never occur consecutively. Conversely, any subset of $[n]$ that satisfies these properties is the peak set of some tableau, hence we will call it an *admissible peak set.*

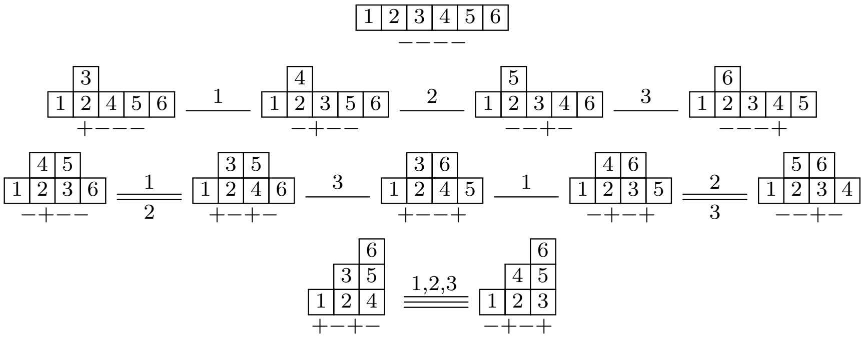

Figure 4.6: The standard shifted dual equivalence graphs of degree 6 omitting σ_1 and σ_6.

In Figure 4.6, all of the standard shifted dual equivalence graphs of degree 6 are drawn and labeled by their signatures omitting σ_1 and σ_6 since 1 and 6 can never be in an admissible peak set. Already from this figure we can see that standard shifted dual equivalence graphs are not always dual equivalence graphs because they can have two vertices connected by 3 edges labeled $i, i+1, i+2$. Also, observe that if vertices v and w are contained in an i-edge, then $\sigma(v)_{[i+1,i+2]} = -\sigma(w)_{[i+1,i+2]}$. Furthermore, notice that the label i of the edge and whether or not it is a double or triple edge can be determined entirely from the peak sets. This fact will be used often in the proofs that follow. From this figure, we also can determine all of the possible standard shifted dual equivalence graphs for $n = 1, 2, 3, 4, 5$ by fixing the higher values.

The standard shifted dual equivalence graphs have several nice properties on par with the dual equivalence graphs and Coxeter-Knuth graphs of type A. By Theorem 4.4.2, each $\mathcal{SG}_\lambda$ is connected. Observe that by Definition 4.2.12, a Schur Q-function Q_λ is the generating function for the sum of peak quasisymmetric functions associated to the labels on the vertices. Recall from Section 4.4 that the lexicographic largest peak set for all standard shifted tableaux of a fixed shape λ is given by the unique tableau U_λ. Thus, the shape λ can be recovered from the multiset of signatures on the vertices.

Each standard shifted dual equivalence graph is an example of the following more general type of graph.

Definition 4.5.2. *Given two finite sets C and S. An S-*signed*, C-*colored graph *consists of the following data:*

1. *a finite vertex set V,*
2. *a signature function $\sigma : V \to \{+,-\}^{|S|}$ associating a subset of S to each vertex,*
3. *a collection E_i of unordered pairs of distinct vertices in V for each $i \in C$.*

We denote a signed colored graph by $\mathcal{G} = (V, \sigma, E)$ where $E = \cup_{i \in C} E_i$. We will say that $\mathcal{G}$ has shifted degree n *if $C = [n-3]$, $S = [n]$, and $\sigma(v)$ is an admissible peak set for an integer sequence of length n for all $v \in V$. We will again encode $\sigma(v)$ by a sequence in $\{+,-\}^n$ where $\sigma(v)_i = +$ if and only if $i \in \sigma(v)$. We use the notation $\sigma_{[i,j]}(v)$ to mean the subset $\sigma(v) \cap [i,j]$ which can be encoded by $+$'s and $-$'s as well. If more than one signed colored graph is defined, we use the notation $V(\mathcal{G}) = V$.*

Definition 4.5.3. Given two S-signed C-colored graphs $\mathcal{G} = (V, \sigma, E)$ and $\mathcal{G}' = (V', \sigma', E')$, a *morphism of signed colored graphs* $\phi \colon G \to G'$ is a map from V to V' that preserves the signature function and induces a map from E_i into E_i' for all $i \in C$. An *isomorphism* is a morphism that is a bijection and whose vertex sets and its inverse is also a morphism.

Definition 4.5.4. *A signed colored graph $\mathcal{G}$ is termed a* shifted dual equivalence graph (SDEG) *if it is isomorphic to a disjoint union of standard shifted dual equivalence graphs.*

The next lemma allows us to classify the isomorphism type of any connected SDEG by a unique standard shifted dual equivalence graph.

Lemma 4.5.5. *Let $\mathcal{SG}_\lambda$ and $\mathcal{SG}_\mu$ be any two standard shifted dual equivalence graphs. If $\phi \colon \mathcal{SG}_\lambda \to \mathcal{SG}_\mu$ is an isomorphism, then $\lambda = \mu$ and ϕ is the identity map.*

Proof. Suppose that $\phi : \mathcal{SG}_\lambda \to \mathcal{SG}_\mu$ is an isomorphism. Then the vertices of $\mathcal{SG}_\lambda$ and $\mathcal{SG}_\mu$ must have the same multisets of associated peak sets. By looking at the unique lexicographically maximal peak set in both, it follows that $U_\lambda = U_\mu$. In particular, $\lambda = \mu$. Thus, ϕ is an automorphism that sends U_λ to itself. Since the h_i defines the i-edges in both $\mathcal{SG}_\lambda$ and $\mathcal{SG}_\mu$ and both graphs are connected, we see that ϕ acts as the identity map. □

The connection between shifted dual equivalence graphs and the type B Coxeter-Knuth graphs stated in Theorem 4.1.3 is now readily apparent. Recall, Theorem 4.1.3 states that every type B Coxeter-Knuth graph $CK_B(w)$ with signature function given by peak sets is a shifted dual equivalence graph, where the isomorphism is given by Q'. It further states that every shifted dual equivalence graph is also isomorphic to some $CK_B(w)$. We give the proof of this theorem now.

Proof of Theorem 4.1.3. We first that show that the map Q' sending vertices in $CK_B(w)$ to their recording tableaux is the desired isomorphism. This follows immediately from the definition of the Kraśkiewicz insertion algorithm, Lemma 4.2.8 and Lemma 4.4.6.

To see the converse statement, observe that $\mathcal{SG}_\mu$ is isomorphic to the Coxeter-Knuth graph $CK_B(w)$ for the increasing signed permutation $w = w(\mu)$ as defined just before Equation 4.2.3. □

Definition 4.5.6. *Given a signed colored graph $\mathcal{G} = (V, \sigma, E)$ of shifted degree n and an interval of nonnegative integers $I = [a, b] \subset [n]$, let*

$$\mathcal{G}^I = (V, \sigma, E_a \cup E_{a+1} \cup \cdots \cup E_{b-3})$$

denote the subgraph of $\mathcal{G}$ using only the i-edges for $a \leq i \leq b-3$. Also define the restriction of $\mathcal{G}$ to I, to be the signed colored graph

$$\mathcal{G}|_I = (V, \sigma', E')$$

1. *$\sigma'(v) = \{s - a + 1 \mid s \in \sigma(v) \cap (a, b)\}$,*

2. *$E'_i = E_{a+i-1}$ when $i \in [|I| - 3]$.*

Notice that the vertex sets of $\mathcal{G}$, $\mathcal{G}^I$ and $\mathcal{G}|_I$ are the same. If $\mathcal{G}$ is a signed colored graph with shifted degree n and $I = [a, b]$ then $\mathcal{G}|_I$ will have shifted degree, but the degree will be at most $|I|$. It could be strictly less than $|I|$ if $n < b$.

Recall the two desirable properties of a signed colored graph $\mathcal{G}$ stated in Theorem 4.1.4. We name them here so we can refer to them easily.

1. **Locally Standard**: If I is an interval of positive integers with $|I| \leq 9$, then each component of $\mathcal{G}|_I$ is isomorphic to a standard shifted dual equivalence graph of degree up to $|I|$.

2. **Commuting**: If $(u, v) \in E_i$ and $(u, w) \in E_j$ then there exists a vertex $y \in V(\mathcal{G})$ such that $(v, y) \in E_j$ and $(w, y) \in E_i$. Thus the components of $(V, E_i \cup E_j)$ for $|i - j| > 4$ are commuting diamonds.

Lemma 4.5.7. *For any standard shifted dual equivalence graph $\mathcal{SG}_\lambda$, both the Locally Standard Property and the Commuting Property hold. In fact, $\mathcal{SG}_\lambda|_{[I]}$ is an SDEG for all intervals I.*

Proof. Consider a standard shifted dual equivalence graph $\mathcal{SG}_\lambda$ for $\lambda \vdash n$. The Commuting Property must hold because h_i acts according to the positions of the values in $[i, i+4]$ only. Hence, h_i and h_j commute provided $|i - j| > 4$.

To demonstrate the Locally Standard Property for a given interval I, observe that we can restrict any $T \in SST(\lambda)$ to the values in I which form a skew shifted tableau and all the data for $\mathcal{SG}_\lambda|_I$ will still be determined. By Lemma 4.4.3, jeu de taquin slides commute with the h_i's. So the isomorphism from $\mathcal{SG}_\lambda|_I$ to a union of standard shifted dual equivalence graph is given by restriction and repeated application of the jeu de taquin operator Δ defined in Definition 4.4.4. $\square$

We note that it is also straightforward to prove Lemma 4.5.7 by appealing to the fact that $\mathcal{SG}_\lambda$ is isomorphic to the Coxeter-Knuth graph $CK_B(w)$ for the increasing signed permutation $w = w(\lambda)$. We know the β_i's satisfy the Commuting Property. Furthermore, restriction on a Coxeter-Knuth graph gives rise to an isomorphism with another Coxeter-Knuth graph since every consecutive subword of a reduced word is again reduced. It is instructive for the reader to consider the alternative proof for the lemmas below using Coxeter-Knuth graphs if that language is more familiar.

Lemma 4.5.8. *Given a strict partition λ of size n, any two distinct components $\mathcal{A}$ and $\mathcal{B}$ of $\mathcal{SG}_\lambda^{[n-1]}$ are connected by an $(n-3)$-edge in $\mathcal{SG}_\lambda$. In particular, any two vertices in*

$\mathcal{SG}_\lambda$ are connected by a path containing at most one $(n-3)$-edge that is not doubled by an $(n-4)$-edge.

Proof. It follows from Theorem 4.4.2 that $\mathcal{A}$ and $\mathcal{B}$ are characterized by the position of n in their respective shifted tableaux. Suppose $\mathcal{A}$ and $\mathcal{B}$ have n in corner cells c and d, respectively with $c < d$. Then there exist tableaux $S \in \mathcal{A}$ and $T \in \mathcal{B}$ that agree everywhere except in the cells containing $n, n-1$ and furthermore $n-2$ lies between n and $n-1$ in the reading word of S and $n-3$ comes before $n-2$. Thus, by definition of h_{n-3}, we have $h_{n-3}(S) = T$ so $\mathcal{A}$ and $\mathcal{B}$ are connected by an $(n-3)$-edge. The final statement of the lemma then follows since any two vertices in a component of $\mathcal{SG}_\lambda|_{[n-1]}$ are connected by a path containing no $(n-3)$-edges by Theorem 4.4.2. $\square$

Lemma 4.5.9. *Let $\mathcal{G} = (V, \sigma, E)$ be a signed colored graph of shifted degree n satisfying the Commuting Property and the local condition that $\mathcal{G}_{[j,j+6]}$ is a shifted dual equivalence graph for all $1 \leq j \leq n-6$. If $v, w \in V$ are connected by an i-edge in $\mathcal{G}$, then $\sigma_k(v) = \sigma_k(w)$ for all $k \notin [i-1, i+5]$.*

Proof. The lemma clearly holds for standard SDEGS by the definition of the shifted dual equivalence moves h_i which determine the i-edges. For $k = i+6$, the lemma holds since $\mathcal{G}|_{[i,i+6]}$ is a shifted dual equivalence graph. Now assume that $i + 6 < k \leq n$. Say $v, w \in V$ are connected by an i-edge in $\mathcal{G}$, and assume $\sigma_j(v) = \sigma_j(w)$ for all $i + 6 \leq j < k$ by induction. By the local condition, the vertex v admits an $(k-2)$-edge if and only if $\sigma_{k-1}(v) = +$ or $\sigma_k(v) = +$, both is not possible since the signature encodes an admissible peak set. Thus, $\sigma_k(v)$ is determined by $\sigma_{k-1}(v)$ and the presence or absence of an adjacent $(k-1)$-edge. Since i-edges and $(k-2)$-edges commute for $k-2-i \geq 5$ by the Commuting Property, we know that v admits a $(k-2)$-edge if and only if w admits a $(k-2)$-edge. Since $\sigma_{k-1}(v) = \sigma_{k-1}(w)$ we obtain $\sigma_k(v) = \sigma_k(w)$ by the same considerations. Therefore, recursively $\sigma_k(v) = \sigma_k(w)$ for all $i + 5 < k \leq n$.

A similar argument works for all $1 \leq k < i-1$. This completes the proof. $\square$

Lemma 4.5.10. *Let λ be a strict partition of n. Let $\mathcal{G} = (V, \sigma, E)$ be a signed colored graph of shifted degree n satisfying the Locally Standard and Commuting Properties. If*

$\phi : \mathcal{G} \longrightarrow \mathcal{SG}_\lambda$ is an injective morphism, then it is an isomorphism.

Proof. Let $v \in V$ and say $\phi(v) = T \in SST(\lambda)$. Since ϕ is signature preserving and $\mathcal{G}$ is Locally Standard, we can apply Lemma 4.5.9 to show that v has an i-neighbor in $\mathcal{G}$ if and only if T has an i-neighbor in $\mathcal{SG}_\lambda$ and a similar statement holds for each of their neighbors. Furthermore, since ϕ is an injective morphism $(v, w) \in E_i \cap E_j$ if and only if $h_i(T) = h_j(T) = \phi(w)$. Thus, ϕ induces a bijection from the neighbors of v to the neighbors of T which preserves the presence or absence of i-neighbors. In particular, every neighbor of T in $\mathcal{SG}_\lambda$ is in the image of ϕ. Since $\mathcal{SG}_\lambda$ is connected, there is a path from T to any other vertex S in $\mathcal{SG}_\lambda$ and by iteration of the argument above we see that ϕ maps some vertex in V to S. Hence, ϕ is both injective and surjective on vertices, and the inverse map is also a morphism of signed colored graphs. Hence, ϕ is an isomorphism. $\square$

With Lemma 4.5.10 in mind, our goal will be to demonstrate the existence of a morphism from any connected signed colored graph satisfying the Locally Standard and Commuting Properties to a standard SDEG and that this morphism is injective. To do this, we will employ an induction on the degree of the signed colored graphs in questions. The next lemma is an important part of that induction.

Lemma 4.5.11. *Let $\mathcal{G} = (V, \sigma, E_1 \cup \ldots \cup E_{n-2})$ be a signed colored graph of shifted degree $n+1$ that satisfies the following hypotheses.*

1. *The Commuting Property holds on all of $\mathcal{G}$.*
2. *Both $\mathcal{G}|_{[n]}$ and $\mathcal{G}|_{[n-6,n+1]}$ are shifted dual equivalence graphs.*

Let $\mathcal{C}$ be a component of $\mathcal{G}^{[n]}$. Then the following properties hold:

1. *There exists a unique strict partition μ of degree $n+1$ and a signed colored graph isomorphism ϕ mapping $\mathcal{C}$ to a component of $\mathcal{SG}_\mu^{[n]}$.*
2. *For every vertex $v \in V(\mathcal{C})$, v has an $(n-2)$-neighbor in $\mathcal{G}$ if and only if $\phi(v)$ has an $(n-2)$-neighbor in $\mathcal{SG}_\mu$.*

We refer to $\mathcal{SG}_\mu$ in this lemma as the *unique extension* of $\mathcal{C}$ in $\mathcal{G}$. The outline of this proof is based on the proof of Theorem 3.14 in [Assaf, 2013], but it uses peak sets in addition to ascent/descent sets for tableaux.

Proof. By hypothesis, $\mathcal{C}|_{[n]}$ is isomorphic to $\mathcal{SG}_\lambda$ for some strict partition $\lambda \vdash n$, so we can bijectively label the vertices of $\mathcal{C}$ by standard shifted tableaux of shape λ in a way that naturally preserves the signature functions σ_i for all $1 \leq i < n$. Since $\mathcal{G}|_{[n-6,n+1]}$ is an SDEG, the lemma is automatically true if $n \leq 7$, so assume $n > 7$.

Partition the vertices of $\mathcal{C}$ or equivalently $SST(\lambda)$ according to the placement of $n-1$ and n. Let D_{ij} be the subgraph of $\mathcal{C}$ with n in row i and $n-1$ in row j with edges in $E_1 \cup \ldots \cup E_{n-5}$, then each D_{ij} is connected since its restriction to $[n-2]$ is also isomorphic to a standard SDEG by hypothesis. Similarly, let D_i be the connected subgraph of $\mathcal{C}$ with vertex set labeled by tableaux with n in row i along with the corresponding edges in $E_1 \cup \ldots \cup E_{n-4}$.

Observe that if $S, T \in V(D_{ij})$, then $(S,T) \notin E_{n-2}$ by the hypotheses and the fact that there are no quadruple edges in SDEGs of shifted degree 7, which can be easily checked. By Lemma 4.5.9, only an $(n-2)$-edge changes σ_n, so σ_n is constant on D_{ij}. The same fact does not hold for each nonempty D_i, but we will show that in order for $\mathcal{G}$ to satisfy the hypotheses of the lemma, the tableaux labeling each D_i can be consistently extended by adding $n+1$ in a row strictly above row i creating a descent from n to $n+1$ or by adding $n+1$ in a row weakly below i creating an ascent.

If $i > j$ and D_{ij} is nonempty, then n appears in a strictly higher row than $n-1$ in all the tableaux labeling vertices of D_{ij}. Furthermore, by the definition of standard shifted tableaux, there is some tableau T labeling a vertex of D_{ij} such that $n-2$ lies in row weakly above row j making position $n-1$ a peak. This implies $\sigma_n(T) = -$ since peaks cannot be adjacent. Since σ_n is constant on D_{ij}, we see that $i > j$ implies $\sigma_n(T) = -$ for all tableaux labeling vertices in D_{ij}.

Fix $i \in [n]$. Say there exists a nonempty D_{ij} such that $i \leq j$ and $\sigma_n(T) = +$ for some $T \in V(D_{ij})$, then $\sigma_n(S) = +$ for all $S \in V(D_{ij})$ because σ_n is constant on D_{ij}. The component D_{ij} is connected to every other $D_{ij'}$ by an $(n-4)$-edge by Lemma 4.5.8. Such an edge $e = (T,U) \in V(D_{ij}) \times V(D_{ij'})$ could be part of a triple edge with an $(n-3), (n-2)$-edge.

In this case, we must have $\sigma_{[n-2,n]}(T) = + - +$ and $\sigma_{[n-2,n]}(U) = - + -$ by the hypothesis on $\mathcal{G}|_{[n-6,n+1]}$ which also implies $\mathcal{G}|_{[n-5,n+1]}$ is an SDEG and Figure 4.6. So, if U is a vertex in $D_{ij'}$, then $i > j'$ since position $n-1$ must be peak of U by Figure 4.6. Therefore, if $n+1$ is added in any row to a tableau $T \in V(D_{ij'})$ it will not create a peak in position n. On the other hand, if D_{ij} is connected to another nonempty $D_{ij'}$ by an $(n-4)$-edge which is not also a $(n-2)$ edge, then again by Figure 4.6 one observes that $\sigma_n(U) = +$ for all $U \in V(D_{ij'})$. Thus, we can consistently extend each vertex in D_i by placing $n+1$ in such a way that it creates a descent from n to $n+1$, so any row strictly above row i will work provided it results in another shifted shape. In this way, an extended tableaux labeling a vertex of D_i will have an $(n-2)$-neighbor using shifted dual equivalence on tableaux if and only if the corresponding vertex of $\mathcal{G}$ has an $(n-2)$-neighbor.

Say there exists a nonempty D_{ij} such that $i \leq j$ and $\sigma_n(T) = -$ for all $T \in V(D_{ij})$. The component D_{ij} is connected to every other $D_{ij'}$ by an $(n-4)$-edge. Assume such an edge $e = (T,U) \in V(D_{ij}) \times V(D_{ij'})$ is part of a triple edge with an $(n-3), (n-2)$-edge. In this case, we must have $\sigma_{[n-2,n]}(T) = - + -$ by the hypothesis on $\mathcal{G}|_{[n-5,n+1]}$ and Figure 4.6. Thus, $n-1$ is a peak of T, but this contradicts the assumption that $i \leq j$. So no $(n-4)$-edge adjacent to a vertex in D_{ij} can be part of a triple edge with an $(n-2)$-edge. We conclude that $\sigma_n = -$ consistently for all nonempty subgraphs $D_{ij'}$ in $\mathcal{C}$. In this case, we can consistently extend each tableau labeling a vertex in D_i by placing $n+1$ in such a way that it creates an ascent from n to $n+1$, so any row weakly lower than i will do. In this way, no extended tableau labeling a vertex of D_i will have an $(n-2)$-neighbor using shifted dual equivalence on tableaux which mirrors the fact that no vertex of D_i has an $(n-2)$-neighbor in $\mathcal{G}$.

To complete the proof, we need an observation about standard shifted dual equivalence graphs. Let T be a shifted standard tableau labeling a vertex in an arbitrary $\mathcal{SG}_\nu$ of degree $n \geq 5$, and assume T admits an $(n-3)$-neighbor. If n was removed from T, one could determine if $n-1$ and n form a descent in T by considering the signed colored graph isomorphism type of the component of $\mathcal{SG}_\lambda|_{[n-4,n]}$. In particular, from Figure 4.6, one can see that either T has a peak in position $n-1$ implying a descent from $n-1$ to n, or T has an ascent from $n-1$ to n.

Let X be the union of all nonempty D_i containing a vertex T with some $\sigma_n(T) = +$, and let Y be the union of all nonempty D_i with no vertex T such that $\sigma_n(T) = +$. Thus, every vertex in X needs a descent from n to $n+1$, and every vertex in Y needs an ascent from n to $n+1$ in order to extend the tableaux labeling of $\mathcal{C}$ in such a way that gives the conclusions of the lemma. Conversely, for each nonempty D_i, we can determine if $D_i \in X$ or $D_i \in Y$ by considering the isomorphism type of a connected component of the subgraph of D_i using only the edges in $E_{n-4} \cup E_{n-3}$ and containing a vertex T which has an $(n-3)$-neighbor. We need to show there exists a unique row where $n+1$ can be placed which satisfies all of the requirements from vertices in X and Y.

If X is empty, then we can extend λ to a strict partition μ by adding one box to the first row of λ. Then $\mathcal{C}|_{[n]}$ is isomorphic to the component of $\mathcal{SG}_\mu|_{[n]}$ with $n+1$ fixed in the first row, and the conclusions of the lemma holds. Thus, we only need to consider the case when X is nonempty.

Say there exists an $i < j$ such that both D_i and D_j are non empty and $D_i \in Y$. Since $\mathcal{C}$ is connected, there exists an edge $e = (S, T)$ connecting $S \in D_i$ and $T \in D_j$. We know D_i and D_j are not connected by edges in $E_1 \cup \ldots \cup E_{n-4}$ since none of these edges change the position of n in $\mathcal{SG}_\lambda$. So, e must be an $(n-3)$-edge which acts as the transposition $t_{n,n-1}$ on S and T by the definition of shifted dual equivalence moves on $\mathcal{SG}_\lambda|_{[n]}$. This implies $S \in D_i$ has $n-1$ in row j and n in row i. In this configuration, $n-1$ cannot be the position of a peak in S, thus S must have a peak in position $n-2$ since it is a vertex of an $(n-3)$-edge and $\mathcal{G}|_{[n-6,n]}$ is an SDEG so we know the possible ways peaks change across edges by the hypotheses and Lemma 4.5.9. If $\sigma_n(S) = +$, then it would contradict the hypothesis that $D_i \in Y$. Thus, $\sigma_{[n-2,n]}(S) = + - -$ which implies $\sigma_{[n-2,n]}(T) = - + -$ by Figure 4.6. Furthermore, from the figure we see that S and T must both have an ascent from n to $n+1$ because there is only one isomorphism type that fits this signature pattern in the last 3 positions. We can conclude that $D_j \in Y$ for all $j > i$. Thus, Y consists of all the D_i for all $i > m$ where m is the maximal value of k for any $D_k \in X$.

Consider a placement of $n+1$ in row $m+1$ extending all the tableaux of shape λ labeling vertices of $\mathcal{C}$. By definition of m, D_m is nonempty and for each tableau in D_m, n must lie in a corner cell of row m, so $\lambda_m > \lambda_{m+1} + 1$. Therefore, $\mu = (\lambda_1, \ldots, \lambda_m, \lambda_{m+1} + 1, \lambda_{m+2}, \ldots)$ is a strict partition of $n+1$. If $n+1$ is placed in μ/λ, then every extended tableau labeling a vertex in X would have a descent between n and $n+1$ as required. Furthermore, every tableau labeling a vertex in Y would have an ascent between n and $n+1$ as required. $\square$

We can also find a *unique lower extension* of a component of $\mathcal{G}|_{[2,n+1]}$ provided similar conditions hold.

Corollary 4.5.12. *Let* $\mathcal{G} = (V, \sigma, E_1 \cup \ldots \cup E_{n-2})$ *be a signed colored graph of shifted degree* $n+1$ *such that the Commuting Property holds on all of* $\mathcal{G}$ *and both* $\mathcal{G}|_{[2,n+1]}$ *and* $\mathcal{G}|_{[1,7]}$ *are shifted dual equivalence graphs. Let* $\mathcal{C}$ *be a component of* $\mathcal{G}^{[2,n+1]}$*. Then, there exists a strict partition* μ *of degree* $n+1$ *and a signed colored graph isomorphism* ϕ *mapping* $\mathcal{C}$ *to a component of* $\mathcal{SG}_{\mu}^{[2,n+1]}$*.*

Proof. This follows from Lemma 4.5.11 provided we reverse all the signatures and the edge labels. Reversing signatures and the edge labels is very natural using the corresponding Coxeter-Knuth graphs since the reverse of a reduced word is reduced. $\square$

Lemma 4.5.13. *Given two shifted standard tableau* T *and* U *of shifted shape* $\lambda \vdash n$*,* T *and* U *are in the same component of* $\mathcal{SG}_{\lambda}|_{[2,n]}$ *if and only if* $\Delta(T)$ *and* $\Delta(U)$ *have the same shape.*

Proof. By definition, T and U are in the same component of $\mathcal{SG}_{\lambda}|_{[2,n]}$ if and only if they are related by a sequence of shifted dual equivalence moves h_i for $2 \leq i \leq n-3$. Since the jeu de taquin operator Δ obeys $\Delta \circ h_i = h_{i-1} \circ \Delta$, by Lemma 4.4.3, T and U are in the same component if and only if $\Delta(T)$ and $\Delta(U)$ are related by a sequence of shifted dual equivalence moves h_i for $1 \leq i \leq n-4$. By Theorem 4.4.2, $\Delta(T)$ and $\Delta(U)$ are related by a sequence of shifted dual equivalence moves h_i for $1 \leq i \leq n-4$ if and only if they have the same shape. $\square$

Lemma 4.5.14. *Let* $\mathcal{G} = (V, \sigma, E_1 \cup \cdots \cup E_{n-3})$ *be a connected SDEG of shifted degree* $n > 9$ *such that* $\mathcal{G}|_{[n-1]}$ *and* $\mathcal{G}|_{[2,n]}$ *are SDEGs, and* $\mathcal{G}$ *satisfies the Commuting Property. Let* $\mathcal{C}$ *be any component of* $\mathcal{G}^{[n-1]}$*, and let* $\mathcal{SG}_{\mu}$ *be the unique extension of* $\mathcal{C}$ *in* $\mathcal{G}$*. Then if two edges connect the image of* $\mathcal{C}$ *to the same component in* $\mathcal{SG}_{\mu}^{[n-1]}$*, then corresponding edges in* $\mathcal{G}$ *must also connect* $\mathcal{C}$ *to the same component in* $\mathcal{G}^{[n-1]}$*.*

Proof. By Lemma 4.5.11, we can label the vertices of $\mathcal{C}$ by standard tableaux of shape μ in a way that preserves signatures and edges in $E_1 \cup \ldots \cup E_{n-4}$ and the set of vertices admitting $(n-3)$-edges. Say S and T are two tableaux in $V(\mathcal{C})$ that admit $(n-3)$-neighbors. Because they both lie in $\mathcal{C}$, n must be in the same cell of both. Assume both are connected to the same component in $\mathcal{SG}_\mu|_{[1,n-1]}$. Then, $n-1$ must also be in the same cell of both, with $n-2$ in some cell between the two in row reading order, and $n-3$ in some cell before that, by the definition of h_{n-3}. Let $U = h_{n-3}(S)$, and let $V = h_{n-3}(T)$. Here U, V are standard tableaux of shape μ, but they do not label vertices in $\mathcal{G}$. Instead, say u, v are vertices in $\mathcal{G}$ such that $(u, S), (v, T) \in E_{n-3}$.

If S and T are connected via edges in $E_1 \cup \ldots \cup E_{n-7}$ then the lemma holds since each of these edges commutes with edges in E_{n-3}. If S and T are connected via edges in $E_2 \cup \ldots \cup E_{n-4}$ then we can assume u and v are also connected by edges in $E_2 \cup \ldots \cup E_{n-4}$ by induction on n.

It remains to show that some S' in the same $(E_2 \cup \ldots \cup E_{n-7})$-component as S and some T' in the same $(E_2 \cup \ldots \cup E_{n-7})$-component as T exists such that S' and T' admit $(n-3)$-edges and are in the same $E_2 \cup \ldots \cup E_{n-4}$-component of $\mathcal{C}$. We divide the possible shapes of μ up into an infinite class and a finite class which needs to be checked to complete the proof.

If $S|_{[n-2]}$ contains $i < n-3$ in a northeast boundary cell c, then we can rearrange the entries of S smaller than i to get S' so that the jeu de taquin process to obtain $\Delta(S')$ removes the cell c and the rest of the jeu de taquin slides are independent of the filling. If $T|_{[n-2]}$ also contains an entry $j < n-3$ in cell c, then rearrange the entries of T to agree with S' in all cells weakly southwest of c to obtain T'. Then $\Delta(T')$ removes the cell c in the jeu de taquin process and by construction, $\Delta(S')$ and $\Delta(T')$ have the same shape since S and T have the same shape. Thus, S' and T' are connected by edges in $E_2 \cup \ldots \cup E_{n-4}$ by Lemma 4.5.13. If the shape of $S|_{[n-2]}$ has 5 or more northeast boundary cells, then such a cell c exists and the lemma holds.

There are only a finite number of strict partitions μ with at most 4 northeast boundary cells after removing 2 corner cells. For example, if μ has 6 or more rows or 9 or more columns then even after removing 2 corner cells there must be at least 5 boundry cells remaining.

Figure 4.7: An example from the proof of Lemma 4.5.14 where $n = 9$ and $\Delta(S)$ does not have the same shape as $\Delta(T)$.

We only need to consider those with at least 10 cells by hypothesis. In each remaining case, one needs to check that no matter how $n, n-1$ are placed in corner cells $\{c, d\}$ of μ then the jeu de taquin argument above may still be applied. We leave the remaining cases to the reader to check to complete the proof. $\square$

In proving the next lemma, we will use Δ^+ to denote the operator that removes the cell containing 1 in a shifted tableau and then performs jeu de taquin into the removed cell. Notice that $\Delta(T)$ is the result of subtracting 1 from each value in $\Delta^+(T)$. This trick allows us to use h_i equivalently on T and $\Delta(T)$ since shifted dual equivalence commutes with jeu de taquin.

Lemma 4.5.15. *Let $\mathcal{G} = (V, \sigma, E_1 \cup \ldots \cup E_{n-3})$ be a connected SDEG of shifted degree $n > 9$ such that $\mathcal{G}$ satisfies the Commuting Property, $\mathcal{G}|_{[n-1]}$ is an SDEG, and $\mathcal{G}|_{[2,n]}$ is an SDEG. Then, there exists a morphism $\phi\colon \mathcal{G} \to \mathcal{SG}_\lambda$ for some strict partition $\lambda \vdash n$.*

Proof. Let $\mathcal{G}^{[n-1]} = (V, \sigma, E_1 \cup \ldots \cup E_{n-4})$ be the subgraph on $\mathcal{G}$ omitting $(n-3)$-edges. If $\mathcal{G}^{[n-1]}$ has only 1 connected component, then $\mathcal{G}^{[n-1]}$ has a unique extension to a standard dual equivalence graph $\mathcal{SG}_\lambda$ of degree n such that $\mathcal{G}^{[n-1]} \cong \mathcal{SG}_\lambda^{[n-1]}$ and neither $\mathcal{G}$ or $\mathcal{SG}_\lambda$ has any $(n-3)$-edges which are not also $(n-2)$-edges. Hence the lemma holds.

Assume $\mathcal{G}^{[n-1]}$ has two distinct components $\mathcal{C}$ and $\mathcal{D}$ which are connected by an $(n-3)$-edge, such a pair exists since $\mathcal{G}$ is connected. By Lemma 4.5.11, both $\mathcal{C}$ and $\mathcal{D}$ have unique extensions with respect to $\mathcal{G}$, say $\mathcal{SG}_\lambda$ and $\mathcal{SG}_\mu$ respectively. Label the vertices of $\mathcal{C}$ and $\mathcal{D}$ by standard shifted tableaux using these extensions. Assume that some $T \in V(\mathcal{C})$ is connected to $U \in V(\mathcal{D})$ by an $(n-3)$-edge. We want to show that $h_{n-3}(T) = U$ proving $\lambda = \mu$.

Applying Lemma 4.5.5, the Commuting Property and the hypothesis that $\mathcal{G}|_{[n-1]}$ is an SDEG, we can assume that $T|_{[n-4]} = U|_{[n-4]}$. We thus only need to show that $T|_{[n-3,n]} =$

$U|_{[n-3,n]}$. To do this, we note that the component of T in $\mathcal{G}|_{[2,n]}$ is isomorphic to the component of the image of T in $\mathcal{SG}_\mu|_{[2,n]}$ by Lemma 4.5.5 and Corollary 4.5.12 and the isomorphism is given by Δ. This follows from the fact that the component of T in $\mathcal{G}^{[2,n-1]}$ is isomorphic to the component of T in $\mathcal{SG}_\mu^{[2,n-1]}$, and this component has a unique lower extension to $\mathcal{G}^{[1,n-1]}$. Thus, $\Delta(T) = h_{n-3}(\Delta(U))$.

By the Commuting Property and Lemma 4.5.14, we can move the values $[1, n-2]$ around in T in any way we like provided the result still admits an $(n-3)$-edge and the follow assumptions will all still hold.

1. $T \in V(\mathcal{C})$ is connected to some $U \in V(\mathcal{D})$ by an $(n-3)$-edge.
2. $T|_{[n-4]} = U|_{[n-4]}$.
3. $\Delta^+(T) = h_{n-3}(\Delta^+(U))$.

We use these conditions to show that U is uniquely determined. If there exists a cell c along the northeast boundary in T containing $i < n-3$, then we can arrange the smaller values so that i slides in the process of applying jeu de taquin in both $\Delta(T)$ and $\Delta(U)$. If c is along the northern boundary, then everything to its right in the same row will shift left by 1 cell and no other entry southeast of c will change. Again, assumptions (2) and (3) above now show that $h_{n-3}(T) = U$ uniquely. Similar reasoning gives the same result when i is along the eastern boundary of T.

If $\lambda = (7, 5, 3, 1)$ and T has $n-3, n-2, n-1, n$ in each of its 4 corners, then both $n-2$ and $n-3$ are between n and $n-1$ in reading order if and only if T admits an $(n-3)$-neighbor. By rearranging $[1, n-2]$ we can find another tableau T' of the same shape as T with n and $n-1$ in the same place and $n-3, n-2$ across the 3rd row. This T' also has an $(n-3)$-neighbor in $\mathcal{D}$ by the shifted dual equivalence relations and Lemma 4.5.14, so it satisfies (1),(2) and (3), and yet it has a corner cell containing $i < n-3$, so we have already considered this case.

Since $n > 8$ by hypothesis, the only remaining case is that $n = 9$ and $\lambda = (5, 3, 1)$. By reversing the roles of T and U, we can also assume $\mu = (5, 3, 1)$. The assumption that T

admits an $(n-3)$-neighbor in $\mathcal{SG}_\lambda$ and $h_{n-3}(T) \neq h_{n-4}(T)$ implies that n and $n-1$ occupy the corners in rows 1 and 3, and $n-3$ and $n-2$ are in the second row by the definition of shifted dual equivalence moves and standard tableaux. The same holds for U. Assumptions (2) and (3) above now show that $h_{n-3}(T) = U$.

To summarize, every component of $\mathcal{G}^{[n-1]}$ has a unique extension to a standard shifted dual equivalence graph for the same shape λ. Thus, we can use $SST(\lambda)$ to label all the vertices of $\mathcal{G}$ in a way that preserves signatures. By hypothesis, if $(S,T) \in E_i$ then as tableaux they are related by h_i for all $1 \le i \le n-4$. We have just shown that if $(S,T) \in E_{n-3}$ then as tableaux they are related by h_{n-3}. Hence there exists a morphism of signed colored graphs from $\mathcal{G}$ to $\mathcal{SG}_\lambda$.

Austin's version of the proof starting just before the itemized portion:

By Lemma 4.5.14, we can consider any T' that results from moving the values $[1, n-2]$ in T such that h_{n-3} acts on T' by switching $n-1$ and n. We may then let U' be the result of applying an $n-3$-edge to T' and show that $h_{n-3}(T') = U'$. The strategy will be consider a several cases to sequentially limit the possible T. In each case, we will identify a cell c of T and form T' such that the jeu de taquin process of applying Δ to T' proceeds through c. At that point, we will leave it to the reader to verify that the following properties completely determine U':

1. $T'|_{[n-4]} = U'|_{[n-4]}$.
2. $\Delta^+(T') = h_{n-3}(\Delta^+(U'))$.
3. $h_{n-3}(T') \notin V(\mathcal{C})$,

First consider the case where there is some T' such that $T'|_{[n-4]}$ has at least two northeast corners c_1 and c_2. We may then form a T_1' and a T_2' such that applying Δ to each proceeds through these two corners, respectively, by rearranging the values in $[n-4]$. Thus, the related tableaux U_1 and U_2 satisfy $U_1|_{[n-3,n]} = U_2|_{[n-3,n]}$. We leave it to the reader to show that $T_1' = h_{n-3}(U_1')$ by considering $\Delta(T_1)$ and $\Delta(T_2)$. We may then assume that $T'|_{[n-4]}$ has exactly one northeast corner, for all choices of T'.

Next let c be the unique northeast corner of $T|_{[n-4]}$. The reader can easily verify the result if $T|_{[n-4]}$ has more than two rows. If $T|_{[n-4]}$ exactly one row, then the fact that $n > 8$ implies that $T|_{[n-4]}$ has at least five columns, and the result follows. We thus need to consider the case where $T|_{[n-4]}$ has exactly two rows.

If $T|_{[n-4]}$ has exactly two rows, then $T|_{[n-4]}$ must have at least four columns, since $n > 9$. If no values in $[n-3, n]$ are in the first row, then the result follows, so assume otherwise. We claim that $n-3$ cannot be in the first cell of the third row of T. If this were the case, we could move $n-3$ or $n-2$ into c to create a T' such that $T|_{[n-4]}$ has two northeast corners. Since $n-3$ is not in the first cell of the third row, then it follows that the third row of T has at most one cell. The result then follows.

Do we care if this proof can be done with $n > 8$? ☐

Remark 4.5.16. For $n = 8$, we may not be able to uniquely determine U in the proof above. See Figure 4.8.

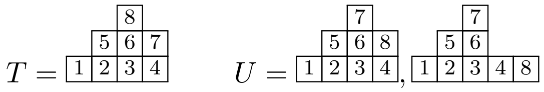

Figure 4.8: An example with $n = 8$ and two distinct possibilities of U from the proof of Lemma 4.5.15.

The following lemma is equivalent to Axiom 6 in Assaf's rules for dual equivalence graphs.

Lemma 4.5.17. *Let $\mathcal{G} = (V, \sigma, E_1 \cup \ldots \cup E_{n-3})$ be a connected SDEG of shifted degree $n > 9$ such that $\mathcal{G}$ satisfies the Commuting Property, $\mathcal{G}|_{[n-1]}$ is an SDEG, and $\mathcal{G}|_{[2,n]}$ is an SDEG. Then, every two distinct components of $\mathcal{G}^{[n-1]}$ are connected by an $(n-3)$-edge.*

This proof is similar to Theorem 3.17 in [Roberts, 2014, p.413] and Lemma 4.5.15 so we only sketch it here.

Proof. The statement in the lemma is equivalent to saying that if a component $\mathcal{A}$ is connected by $(n-3)$-edges to components $\mathcal{B}$ and $\mathcal{C}$, then $\mathcal{B}$ and $\mathcal{C}$ are connected to each other by an

$(n-3)$-edge. Using Lemma 4.5.13, we may apply properties of jeu de taquin to show that this must be the case so long as λ is not a pyramid or λ has more than three northeast corners. The largest example of a shifted shape that violates these two rules is the pyramid $(5,3,1)$ with nine cells. By assumption, $n > 9$, and so the argument is complete. $\square$

Proof of Theorem 4.1.4. The fact that $\mathcal{SG}_\lambda$ satisfies the Commuting Property and the Locally Standard property is proved in Lemma 4.5.7. To prove the converse, assume $\mathcal{G}$ is a signed colored graph with shifted degree n satisfying both of these properties. Proceed by induction on n. For $n \leq 9$, the result is known by the Locally Standard Property. We may then assume $n > 9$.

By Lemma 4.5.15, $\mathcal{G}$ admits a morphism onto $\mathcal{SG}_\lambda$. By Lemma 4.5.10, we need only show that this morphism is injective. The morphism was constructed in such a way that it is the unique extension on any component of $\mathcal{G}^{[n-1]}$ so it is injective on each component automatically. Furthermore, the location of n is constant on each component. Let C and D be two distinct components of $\mathcal{G}^{[n-1]}$, and let $v \in V(C)$ and $w \in V(D)$. By Lemma 4.5.17, there exists an $(n-3)$-edge connecting C to D which necessarily moves n in the tableaux labeling its endpoints under the morphism. Thus, the morphism maps v and w to tableaux with n in two different positions. Hence, the morphism is injective. $\square$

4.6 Open Problems

We conclude with some interesting open problems.

1. What is the geometrical significance of the Little bumps?

2. What interesting symmetric functions expand as a positive sum of Schur Q's? Are there natural expansions of certain symmetric functions first into peak quasisymmetric functions?

3. What is the diameter of the largest connected component of a Coxeter-Knuth graph for permutations or signed permutations of length n?

4.7 Acknowledgments

Many thanks to Mark Haiman, Peter Winkler, and Tao Kai Lam for helpful discussions on this work.

Appendix A

CLASSIFICATION OF DUAL EQUIVALENCE GRAPHS

A.1 Introduction

The question of whether or not a symmetric function can be expressed in the basis of Schur functions using only non-negative coefficients (Schur positivity) is often very challenging to answer; more so if we require the proof to be strictly combinatorial. To this date, there is no known combinatorial proof for the famous Macdonald polynomials (the first non-combinatorial proof can be found in Haiman [2001]). In her upcoming paper Assaf [2013], Sami Assaf develops the theory of dual equivalence graphs as a solution to this problem. In it, she relates Schur functions to connected components of graphs, which she associates to various polynomials. We seek to recreate some of the crucial results of this paper, ameliorating several errors in the 2011 draft along the way.

After some preliminaries, we describe the dual equivalence graph of a partition, which is intimately tied to standard tableaux and Schur functions. After this, we state Assaf's axioms for dual equivalence graphs. The main theme of this paper will be the classification of all such abstract dual equivalence graphs in terms of partitions.

While some familiarity with symmetric functions is assumed, the proofs themselves require no prior knowledge. We hope that this paper will lay out the technical groundwork of dual equivalence graphs, so that they may be used more freely in future research. As a new combinatorial object, they have the inherent promise of providing new enumerative approaches to old problems, especially in the fields of symmetric functions and algebraic combinatorics.

(A note: Some of the shorter definitions, theorem statements, and some diagrams were taken verbatim from Assaf [2013]. The paper itself, though it contains original elements, thoroughly leans on the content and structure of Assaf [2013].)

A.2 Prelimaries

A.2.1 Tableaux

This section quickly introduces the key notation and definitions that underlie the rest of the paper. A *partition* λ is a weekly decreasing finite sequence of non-negative integers $\lambda_1 \geq \ldots \geq \lambda_k \geq 0$. If $\sum \lambda_i = n$, we say that λ is a partition of n and write $\lambda \vdash n$. Partitions are often expressed in terms of diagrams where λ_i is the number of boxes, or *cells*, in the i^{th} row, from bottom to top, as in the left diagram of Figure A.1.

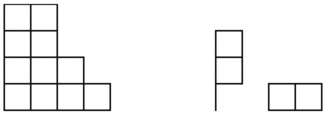

Figure A.1: The diagrams for (4,3,2,2) and (4,1,1)/(2)

If the diagram of ρ is contained in the diagram of λ, equivalently $\rho_i \leq \lambda_i$ for all i, then we may consider the *skew diagram* λ/ρ defined by omitting the boxes of ρ from λ, as in the right skew diagram of Figure A.1.

A *filling* assigns a positive integer to each cell of a partition, usually written inside of the cell. A partition together with a filling is called a *tableau*. We will only be concerned with *standard* fillings, which use each number in $[n]$ exactly once and are increasing up columns and across rows from left to right, as in Cartesian coordinates (see Figure A.2). Standard fillings are often referred to as *standard Young tableau*, and the set of all standard Young tableau with partition λ is denoted SYT(λ). The notion of a standard Young tableau extends to skew shapes, as seen in the second tableau of Figure A.2.

The *content* of a cell is $j - i$, where j is the column and i is the row. For example, in a standard Young tableau, the 1-cell is in the first row and column, and so has content 0. In other words, each diagonal going southwest to northeast has the same content, with the uppermost having the largest content. We then define the *content reading word* of a tableau by reading off each entry from highest content to lowest, moving northeast along each diagonal, as in Figure A.2. The content reading word of a standard Young tableaux is necessarily a permutation.

Figure A.2: Tableaux with content words 438162957 and 3214.

The *signature* of a permuation is a string of 1's and -1's, or +'s and -'s for short, where there is a + in the i^{th} position if and only if i comes before $i+1$ in the word. Notice that the word is one entry longer than the signature. We may then define the signature of a tableau T, denoted $\sigma(T)$, as the signature of the content reading word of T. For example, the signatures of the tableau in Figure A.2 are $+--+-+-+$ and $--+$, respectively.

A.2.2 Symmetric Functions

This section motivates the importance of the main results but is not necessary for an understanding of the proofs in the rest of the paper. Tableaux are of a crucial importance in giving combinatorial descriptions of symmetric functions, as is well laid out in Macdonald [1995] and Fulton [1997]. The ring of symmetric functions, Λ, has several well known basis with ties to tableaux. The one of most concern to us - and many others as well - is the set of Schur functions $\{s_\lambda\}$. We will take the unorthodox approach of defining these functions using a result of Gessel. While less immediately intuitive than standard definitions, this definition contains the only properties that we need. We will first need a preliminary definition.

Definition A.2.1. Given any signature $\sigma \in \{\pm 1\}^{n-1}$, we define the *fundamental quasi-symmetric function* $F_\sigma(X) \in \mathbb{Z}[x_1, x_2, \ldots]$ by

$$F_\sigma(X) := \sum_{\substack{i_1 \leq \ldots \leq i_n \\ i_j = i_{j+1} \Rightarrow \sigma_j = +1}} x_{i_1} \cdots x_{i_n}.$$

The set of fundamental quasi-symmetric functions of degree n forms a homogeneous basis for the *ring of quasi-symmetric functions* of degree n. The extent that we need this ring to motivate our results is limited to a few facts. The first is the promised definition of Schur functions.

Definition A.2.2 (Gessel [1984]). Given any partition λ, we may define a *Schur function* by

$$s_\lambda(X) := \sum_{T \in SYT(\lambda)} F_{\sigma(T)}(X).$$

While it is not obvious from this definition that Schur functions are symmetric or that they form a basis for Λ, what we have gained from this definition is a clear connection to the signatures of fillings of tableaux.

The second fact is that the important Lascoux-Lecrec-Thibon (LLT) polynomials and Macdonald polynomials (as introduced in Lascoux et al. [1997] and Macdonald [1988], respectively) may also be expressed using the sum of fundamental quasi-symmetric polynomials over fillings of a fixed partition (see Haglund et al. [2005a] for details). It is then desirable to be able gather the fillings that generate such a symmetric function, say a Macdonald polynomial, in a such a way that they correspond to Schur functions. That is, recollect terms by collecting all of the standard fillings of various partitions. The goal of the next section will be to develop dual equivalence graphs for the expressed purpose of allowing such a recollection. More specifically, we want a tool to express LLT and Macdonald polynomials as the sum of Schur functions.

A.3 Dual Equivalence Graphs

There is a well known equivalence class on words from the alphabet of intergers, which can be defined using the so called elementary equivalence operations or *jeu de taquin*. A fuller description can be found in Fulton [1997] or Stanley [2012b]. This equivalence assigns a unique class to every standard Young tableau (actually, every semi-standard tableaux). We begin this section by defining a similar equivalence that relates any two standard tableau of the same shape.

Definition A.3.1. Given a permutation of length n consisting of all numbers in $[n] = \{1, \ldots, n\}$, define an *elementary dual equivalence* as an involution d_i that interchanges $i-1, i$, and $i+1$ as

$$d_i(\ldots i \ldots i \pm 1 \ldots i \mp 1 \ldots) := (\ldots i \mp 1 \ldots i \pm 1 \ldots i \ldots)$$

when i is not between $i-1$ and $i+1$, and d_i is the identity otherwise. Two words are *dual equivalent* if one may be transformed into the other by successive elementary dual equivalences.

As an example, 21345 is dual equivalent to 51234 because $d_4(d_3(d_2(21345))) = d_4(d_3(31245)) = d_4(41235) = 51234$. We may also let the d_i act on the entries of a standard Young tableau by applying them to the content reading word. It is not hard to check that the result is again a standard Young tableau. This fact, as well as the following theorem are proved in Haiman [1992].

Theorem A.3.2 (Haiman [1992])**.** *Two standard tableaux on partition shapes are dual equivalent if and only if they have the same shape.*

By using each d_i as an edge colored by i, we may then treat each set of all standard young tableaux with n boxes as colored graph with each $\mathrm{SYT}(\lambda)$ as a connected component. We may further label each vertex with its signature to create a *standard dual equivalence graph* that we will denote $\mathcal{G}_\lambda$ (see Figure A.3).

Here, $\mathcal{G}_\lambda$ is an example of the following broader class of graphs.

Definition A.3.3. A *signed, colored graph of type (n,N)*, also called an *(n,N)-signed, colored graph*, consists of the following data:

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

We denote such a graph by $\mathcal{G} = (V, \sigma, E_2 \cup \cdots \cup E_{n-1})$ or simply (V, σ, E).

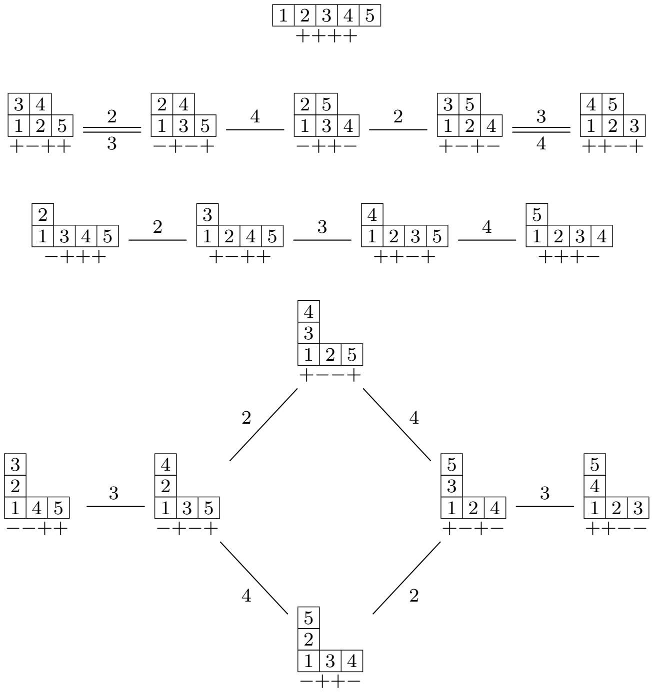

Figure A.3: Some dual equivalence graphs with standard Young tableaux as vertices.

We may now specialize this definition to the main object of this paper - an abstract generalization of the structure inherent in any $\mathcal{G}_\lambda$. We go through the pains of this abstract definition because it provides a specific set of properties to check when working with less understood graphs.

Definition A.3.4. A signed, colored graph $\mathcal{G} = (V, \sigma, E)$ of type (n, N) is a *dual equivalence graph of type* (n, N) if $n \leq N$ if 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 an $x \in V$ such that $\{w, x\} \in E_i$. The choice of x is also required to be unique.

(ax2): If $\{w, x\} \in E_i$, then $\sigma(w)_j = -\sigma(x)_j$ for $j = i-1, i$, and $\sigma(w)_h = \sigma(x)_h$ for $h < i-2$ or $h > i+1$.

(ax3): For $\{w, x\} \in E_i$, if $\sigma(w)_{i-2} = -\sigma(x)_{i-2}$, then $\sigma(w)_{i-2} = -\sigma(w)_{i-1}$, and if $\sigma(w)_{i+1} = -\sigma(x)_{i+1}$, then $\sigma(w)_{i+1} = -\sigma(w)_i$.

(ax4): Every connected component of $(V, \sigma, E_{i-1} \cup E_i)$ appears in Figure A.4 , and every connected component of $(V, \sigma, E_{i-2} \cup E_{i-1} \cup E_i)$ appears in Figure A.5.

(ax5): If $\{w, x\} \in E_i$ and $\{x, y\} \in E_j$ for $|i-j| \geq 3$, then $\{w, v\} \in E_j$ and $\{v, y\} \in E_i$ for some $v \in V$.

(ax6): Any two vertices of a connected component of $(V, \sigma, E_2 \cup \cdots \cup E_i)$ may be connected by some path crossing at most one E_i edge.

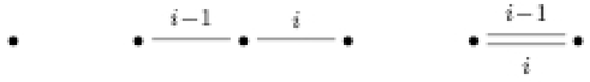

Figure A.4: Allowable $E_{i-1} \cup E_i$ components of axiom 4

Before justifying this definition, we point out a few properties about the structure of dual equivalence graphs. Graphs of different types may often be related by ignoring some of the information. If $M \leq N$, then the (m, M)-*restriction* of a graph, is the result of excluding E_i for $i \geq m$ and projecting each signature onto its first $M-1$ coordinates.

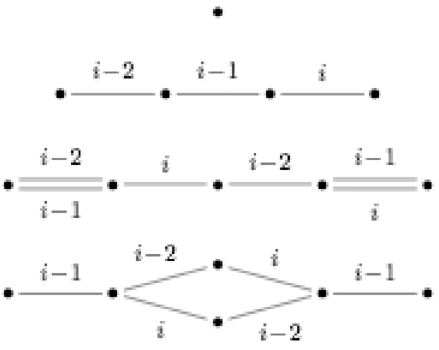

Figure A.5: Allowable $E_{i-2} \cup E_{i-1} \cup E_i$ components of axiom 4

Remark A.3.5. *The following are immediate consequences of the definition:*

1. *A connected component of an* (n, N)*-dual equivalence graph is also an* (n, N)*-dual equivalence graph.*
2. *For any* $m \leq n$ *and* $M \leq N$*, the* (m, M) *restriction of a* (n, N)*-dual equivalence graph is an* (m, M)*-dual equivalence graph.*
3. *The restrictions given in axiom 4 are precisely those found in some* G_λ*, where* $\lambda \vdash 4$ *or* $\lambda \vdash 5$*.*

The goal is to connect the tableau definition of dual equivalence graphs to the abstract definition. As an initial justification, we prove the following theorem.

Theorem A.3.6 (Assaf [2013], Theorem 3.5)**.** *For any* $\lambda \vdash n$*,* $\mathcal{G}_\lambda$ *is an* (n, n)*-dual equivalence graph.*

Proof : We show that $\mathcal{G}_\lambda$ satisfies each of the axioms in order. For the most part, this will just be a matter of referring to the definition of d_i.

Fix some tableau $T \in SYT(\lambda)$. The entry i is not between the $i-1$ and $i+1$ in the content reading word of T precisely when $\sigma(T)_i = -\sigma(T)_{i-1}$. Furthermore, d_i is an involution, so T is connected to at most 1 other tableau by an i edge. These two properties imply that axiom 1 is satisfied.

The action of d_i, when not trivial, always reverses σ_i and σ_{i-1} while leaving σ_h unchanged for $h \geq i+2$ and $h \leq i-3$. Thus axiom 2 is satisfied.

For d_i to change σ_{i-2}, it must move $i-1$ to the opposite side of $i-2$ in the content reading word of T, but for d_i to move $i-1$ at all, $i+1$ must be between $i-1$ and i. Putting these together, we see that we must have $\sigma(T)_{i-2} = -\sigma(T)_{i-1}$, satisfying the first half of axiom 3. The second half is identical.

Because d_i can be thought of as acting on permutations and ignoring most of the numbers, we may think of axiom 4 as a description on the possible graphs produced by d_2, d_3 and d_4 starting from any permutation of {1, 2, 3, 4, 5}. The fact that each such permutation can be actualized in a standard dual equivalence graph is easily shown but unnecessary. Direct inspection shows that axiom 4 holds. Notice that each permutation can be realized as a standard ribbon tableau. This inspection is made much simpler by noticing that reversing the order of a permutation simply reverses the order on the edges. Also, the graph of a transpose of a graph will also have the same shape with edge numbers inverted. These facts and the information provided in Figure A.3 reduced the inspection to a more manageable number of cases.

Axiom 5 follows since $|i-j| \geq 3$, implies d_i and d_j commute. This is simply because {i-1, i, i+1} and {j-1, j, j+1} are disjoint.

If T may be transformed into some chosen U using only d_j, with $j \leq i$, then the standard tableaux formed by removing the $i+2, \ldots, n$ from T and U must be the same shape. If $i+1$ occurs in the same cell in U and T, then Theorem A.3.2 guarantees that U and T may be connected with only edges in $E_2 \cup \ldots \cup E_{i-1}$. Now suppose that $i+1$ is in a different cell of U and T. By Theorem A.3.2, we may then move the i-entry in T to the position of $i+1$ in U and move $i-1$ so that it occurs between i and $i+1$ in the content word by using only d_j with $j < i$. By standardness, such a position for $i-1$ must exist, and by the restriction on j, the cells of $i+2, \ldots, n$ have not been changed. We are then free to use d_i to switch the places of i and $i+1$, and then employ Theorem A.3.2 again to recover the rest of U. This demonstrates the desired connectedness condition of axiom 6. $\square$

In order to fully classify dual equivalence graphs, we will first need to define isomorphisms.

Definition A.3.7. A map ϕ between two signed, colored graphs of type (n, N) is called a *morphism* if it preserves signatures and i-edges. That is, $\sigma(\phi(v)) = \sigma(v)$ and $\{v, w\} \in E_i$ implies $\{\phi(v), \phi(w)\} \in E_i$. A morphism is an *isomorphism* if it is a bijection on the vertex set and admits an inverse morphism.

We are now able to state the main theorems of this paper.

Theorem A.3.8 (Assaf [2013], Theorem 3.10)**.** *Every connected component of an* (n, n)*-dual equivalence graph is isomorphic to a unique* G_λ.

To prove this, we will build up a sequence of lemmas and propositions. We are, however, immediately able to give a proof of uniqueness.

Proposition A.3.9 (Assaf [2013] Proposition 3.9)**.** *If* $\phi : \mathcal{G}_\lambda \to \mathcal{G}_\mu$ *is an isomorphism, then* $\lambda = \mu$ *and* $\phi = id$.

Proof: If we fill the diagram of λ in numerical order, always avoiding minus signs in the signature when possible, we are given a unique filling T_λ with 1 through λ_1 in the first row and continuing up, row by row. We may then read off the shape from the signature as $\sigma(T_\lambda) = +^{\lambda_1 - 1}, -, +^{\lambda_2 - 1}, -, \ldots$. We will refer to any signature that can be written with these weakly decreasing exponents as a *partition type signature.* Notice that in lexographic order, the largest partition type signature of any filling of λ is λ itself, as was seen by its construction. Isomorphisms create a bijection on vertices, so both $\mathcal{G}_\lambda$ and $\mathcal{G}_\mu$ have the same largest partition shape signature. Thus $\lambda = \mu$. Since T_λ is the unique filling with this signature in $\mathcal{G}_\lambda$, it is necessarily a fixed point of ϕ. By the uniqueness statement in axiom 1, all points connected to T_λ are also fixed points. Inductively, all of $\mathcal{G}_\lambda$ is fixed, and so $\phi = id$. $\square$

As we move forward, we will often assume as little as possible in the statement of theorems. The following lemmas demonstrate the relationship between axioms and morphisms and will be useful in proving later theorems.

Lemma A.3.10. *Assume that $\mathcal{G}$ and $\mathcal{H}$ are (n,N)-signed, colored graphs that obey axiom 1. Also, suppose that there exists a morphism $\phi : \mathcal{G} \to \mathcal{H}$.*

1. *If $\mathcal{H}$ is connected, then ϕ is surjective.*
2. *If ϕ is surjective and either $\mathcal{G}$ or $\mathcal{H}$ obey axiom 2 or 3, they both do.*
3. *If ϕ is surjective and $\mathcal{G}$ obeys axiom 5 or 6, then $\mathcal{H}$ does as well.*
4. *If ϕ is bijective, then it is an isomorphism.*
5. *If $\mathcal{G}$ is connected and obeys axioms 5 and 6, and $\mathcal{H}$ is a standard dual equivalence graph, then ϕ is an isomorphism.*

Proof: We begin with Part 1 and continue in order. Choose any vertex $v \in \mathcal{G}$ and assume that vertex $w \in \mathcal{H}$ is connected to $\phi(v)$ by an i-edge. Because v and $\phi(v)$ have the same signature, axiom 1 guarantees that v is connected to a unique w' by an i-edge. This edge must be preserved by ϕ, being sent to a unique edge (by axiom 1) of w. Thus $\phi(w')$ must be w. Because $\mathcal{H}$ is connected, we may then induct to any vertex, proving Part 1.

Notice that we have just shown a bijective correspondence between edges and vertices connected to v and the those of $\phi(v)$. That is, ϕ restricts to an isomorphism between the set of vertices sharing an edge with v and the set of vertices sharing an edge with $\phi(v)$, where these colored edges are also preserved by ϕ. Because axioms 2 and 3 are strictly concerned with signature changes over edges, we have shown that if either graph obeys one, they both do, proving Part 2.

For axioms 5 and 6 in Part 3, it is simply a matter of checking that a violation in $\mathcal{H}$ would create a violation in the pull back. As we will not specifically need these properties, we leave the details to the reader.

Part 4 follows from the fact that axiom 1 gaurantees that ϕ restricts to local isomorphisms, as discussed above. This and bijectivity guarantees an inverse morphism.

For Part 5, it suffices to prove the ϕ is injective, since then Parts 1 and 5 guarantee that it is an isomorphism. Suppose that $\phi(x) = \phi(y)$ and chose the smallest i such that x and y are in the same component of the (i+1, i+1) restriction. If $x \neq y$, we may employ axiom 6 to find a path connecting x and y with only one i-edge, call it $(u, v) \in E_i$. In the image of this path, $(\phi(u), \phi(v))$ is the only edge that can change the position of the $i+1$-cell, which would lead to $\phi(x) \neq \phi(y)$. Thus, we may assume that the $i+1$-cell is not changed.

Because $d_i(\phi(u))$ does not change the position of the $i+1$-cell, $i+1$ must be between i and $i-1$ in the content words of $\phi(u)$ and $\phi(v)$. By Theorem A.3.2, the $i-2$-cell in $\phi(u)$ and $\phi(v)$ can be moved between i and $i-1$ using only j-edges with $j \leq i-3$. By their construction, these new tableaux are connected by an $i-1$-edge. Axioms 1 and 5, guarantee that this sequence of elementary dual equivalences pulls back, edge by edge, to a new path connecting u and v contained in the (i, i) restriction. This contradicts our choice of i, concluding our proof. $\square$

In light of this lemma, we would like to find a morphism for any (n, n)-dual equivalence graph to some $\mathcal{G}_\lambda$. We will build up such a morphism inductively. It will be helpful to embed dual equivalence graphs into larger structures using the following definition.

Definition A.3.11. Fix any partitions $\lambda \subset \rho$ with $|\lambda| = n$ and $|\rho| = N$ and standard skew tableau A of shape ρ/λ with entries $n+1, \ldots, N$. We then define the set of Young tableaux *augmented by* A, denoted $\mathrm{ASYT}(\lambda, A)$, as the set of $T \in \mathrm{SYT}(\rho)$ such that the restriction of T to ρ/λ is A. Further, we define its signed colored graph given by the action of d_i, for $i < n$, as $\mathcal{G}_{\lambda,A}$

Remark A.3.12. $\mathcal{G}_{\lambda,A}$ is just a connected component of the (n, N)-restriction of $\mathcal{G}_\rho$, and is therefore a dual equivalence graph with (n, n)-restriction isomorphic to $\mathcal{G}_\lambda$.

We can now use the notion of augmentation to prove that we may extend some morphisms from an (n, n)-restriction to the entire dual equivalence graph.

Proposition A.3.13 (Assaf [2013] Proposition 3.13)**.** *Let $\mathcal{G} = (V, \sigma, E)$ be a connected (n, N)-signed colored graph obeying axiom 1 and 6, and let ϕ be a morphism from the (n, n)-restriction of $\mathcal{G}$ to $\mathcal{G}_\lambda$ for some partition λ of n. Then ϕ extends to a morphism $\tilde{\phi} : \mathcal{G} \to \mathcal{G}_{\lambda,A}$, where $|\lambda| + |A| = N$. Furthermore, the position of the $(n+1)$-cell of A is unique.*

Proof : We may think of this proof as finding a suitable location and filling for the cells of A. First notice that Lemma A.3.10 gives us that $\mathcal{G}$ satisfies axiom 2 and 3. By axiom 2 and connectedness, σ_h is constant for $h \geq n+1$, so we are free to place the $n+1, \ldots, N$-cells anywhere that satisfies this fixed choice of signs. The difficulty is in showing that there is an acceptable choice for the $n+1$ cell. Throughout, we will use $\mathcal{G}$ for itself and for its (n, n) restriction. The distinction should be clear from context.

$\mathcal{G}$ and $\mathcal{G}_\lambda$ obey axiom 1, and so ϕ is a surjection by Lemma A.3.10. Thus $\mathcal{G}$ has vertices that map onto tableaux with an n cell in each of the northeast corners of λ, by Theorem A.3.2. Notice that σ_n dictates on which side (content-wise) of each of these corners an $n+1$ box may be placed. That is, if a tableau T has $\sigma_n(T) = +1$, then the only allowed augmentations to this tableau will have an $n+1$ cell to the southeast of the n-cell of T. We have the opposite result if $\sigma_n(T) = -1$. When a collection of tableaux with n in all of the northeast corners are considered, there can be at most one possible augmentation consistent with the signatures at each of these tableaux. A consistent placement for an $n+1$-cell, if one exists, must then be unique.

If σ_n is constant on $\mathcal{G}$, then simply place the $n+1$ cell in the first column if it is -1 on all vetices or the first row if it is +1. Both cases are consistent with the signatures of every vertex. Now suppose $x, y \in \mathcal{G}$ such that $\sigma(x)_n = -\sigma(y)_n = +1$. To demonstrate an allowable placement of an $n+1$-cell, it suffices to show that the n-cell of $\phi(x)$ must be northwest of the n-cell of $\phi(y)$. We could then place the $(n+1)$-cell southeast of all such x's and northwest of all such y's. By axiom 6, x can be connected to y by a path that has only one $n-1$ edge. This is the only edge that can change the placement of the n-cell, and by axiom 2, the only edge that can change σ_n. The problem thus reduces to the action of this $(n-1)$-edge, so we may assume that it actually connects x and y, relabeling if necessary. By axioms 2 and 3, $\sigma(x)_{n-2,n-1,n} = -\sigma(y)_{n-2,n-1,n} = + - +$. These signatures, and the fact that the $(n-1)$ edge acts by switching n and $n-1$, force the position of the n cell in $\phi(x)$ to be northwest of the position of the n-cell in $\phi(y)$, concluding our proof. □

This proposition will be used in an induction argument in the proof of Theorem A.3.8, but first we need more language to describe the restrictions of graphs where there may not be an edge despite the signature permitting one.

Definition A.3.14. Let $G = (V, \sigma, E)$ be a signed, colored graph of type (n, N) satisfying axiom 1. We say that $w \in V$ *admits an i-neighbor* if $\sigma(w)_{i-1} = -\sigma(w)_i$.

We may now proceed to the proof of Theorem A.3.8.

Proof of Theorem A.3.8 : Let $\mathcal{G}$ be a connected dual equivalence graph of type $(n+1, n+1)$. It suffices to create a morphism $\phi : \mathcal{G} \to \mathcal{G}_\lambda$ for some $\lambda \vdash n+1$. Then by Lemma A.3.10, ϕ is an isomorphism, and proposition A.3.9 gives uniqueness.

Up to $n+1 = 4$, the result is largely dictated by axiom 4. There are only 8 possible signatures, so the number of cases is easy to handle. We then proceed with $n \geq 4$ by strong induction, assuming the result for (m, m)-dual equivalence graphs with $m \leq n$.

Let $\mathcal{C}$ and $\mathcal{D}$ be connected components of the $(n, n+1)$ restriction of $\mathcal{G}$ that are connected by an edge, $(u, v) \in E_n$ with $u \in \mathcal{C}, v \in \mathcal{D}$. By induction, there are unique partitions μ and ν permitting isomorphims $\mathcal{C} \cong \mathcal{G}_\mu, \mathcal{D} \cong \mathcal{G}_\nu$ on the (n, n)-restrictions. By Proposition A.3.13, we may uniquely extend these to $\mathcal{G}_{\mu,A}$ and $\mathcal{G}_{\nu,B}$, where $|A| = |B| = 1$. Let λ be the shape of (μ, A). We wish to show that (ν, B) also has shape λ, allowing us to patch together the desired morphism.

We will simultaneously think of $\mathcal{C}$ as an $(n, n+1)$-component of $\mathcal{G}$ and $\mathcal{G}_\lambda$. In $\mathcal{G}_\lambda$, the vertex u must still admit an n neighbor, v', in a distinct connected component $\mathcal{D}'$. It then suffices to show that $\mathcal{D}$ and $\mathcal{D}'$ may be embedded into the same standard dual equivalence graph, namely $\mathcal{G}_\lambda$. For convenience, we will denote the connected (m, M)-components of v and v' as $\mathcal{D}|_{(m,M)}$ and $\mathcal{D}'|_{(m,M)}$, respectively.

By axiom 4, the $E_n \cup E_{n-1} \cup E_{n-2}$ component of u dictates the $E_n \cup E_{n-1} \cup E_{n-2}$ component of both v and v'. Axiom 5 says that v and v' admit an i-neighbor exactly when u does, for any $i \leq n-3$. Thus v admits an i-neighbor if and only if v' does. Since $\sigma(v)_n = \sigma(v')_n = -\sigma(u)_n$, axioms 1 and 2 fix the entire signature as $\sigma(v) = \sigma(v')$. Axiom 5 then determines the edge structure of $\mathcal{D}|_{(n-2,n+1)}$ and $\mathcal{D}'|_{(n-2,n+1)}$. Also by axiom 5, any vertex in either of these components is still connected to $\mathcal{C}$ by an n-edge. Repeating the argument above, the signatures of all of these vertices must be fixed. Thus, $\mathcal{D}|_{(n-2,n+1)} \cong \mathcal{D}'|_{(n-2,n+1)}$. Proposition A.3.13 embeds these components into augmented standard dual equivalence graphs whose shapes can only differ in the placement of the n

and $n+1$ boxes. It then suffices to show that the position of these boxes is unique.

To see that the location of the n-box is fully determined, recall the proof of Proposition A.3.13. We need only a fixed choice of σ_n on a set of vertices that correspond to tableaux with the $n-1$-cell in each northeast corner of the $(n-1, n-1)$-restristions. Such a set must exist by Theorem A.3.2. Axiom 5 may be used to rearrange the first $n-2$ entries while staying connected to $\mathcal{C}$ by an n-edge. Then axiom 6 allows $n-1$ to be moved to any corner by a single $n-2$-edge, thus staying in the $E_n \cup E_{n-1} \cup E_{n-2}$ components of a vertex in $\mathcal{C}$. Axiom 4 then fixes σ_{n-1}, in turn fixing the position of the n-box.

To see that the location of the $n+1$-box is fully determined, we first consider the case were n is between $n-1$ and $n+1$ in the tableau corresponding to v. Axiom 5 allows $n-2$ to be moved between $n-1$ and $n+1$ while still connecting to $\mathcal{C}$ by an n-edge. But now this n-edge is doubled by an $n-1$ edge, contradicting the fact that $\mathcal{C}$ is in a different $(n, n+1)$ component than $\mathcal{D}$. The same argument holds for v'.

We may then assume that $n-1$ is between n and $n+1$ in both v and v'. Suppose that $n+1$ is in cell x of v and cell y of v'. If $x \neq y$, $n-2$ may be positioned between x and y in the content reading word in such a way that it is not between $n-1$ and n. $n-3$ may then be placed between $n-1$ and $n-2$, as seen in Figure A.6. Such a location is guaranteed by standardness. As before, this can be done while maintaining an n-edge with $\mathcal{C}$. Applying an $n-2$-edge moves $n-1$ between x and y. Changing the location of $n+1$ between x and y will change whether or not this vertex admits a doubled $n, n-1$-edge. In Figure A.6, there will be a double edge if $n+1$ is in the y-cell but not if it is in the x-cell. The vertex is still in the $E_n \cup E_{n-1} \cup E_{n-2}$ components of a vertex of $\mathcal{C}$, and so the existence of a double edge is dictated by axiom 4. Thus, the position of $n+1$ is fully determined. With all boxes locations fixed, (μ, A) and (ν, B) must both have shape λ.

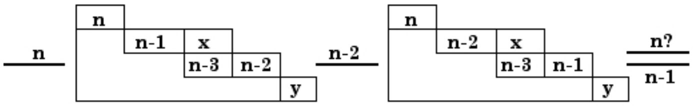

Figure A.6: Two possibilities for $n+1$. Choosing y provides the double edge.

Therefore, each $(n, n+1)$ component of $\mathcal{G}$ may be embedded into $\mathcal{G}_\lambda$, and because this embedding has already been shown to respect n edges, it lifts to the desired morphism. □

We also have the immediate corollaries.

Corollary A.3.15. *Let* $\phi : \mathcal{G} \to \mathcal{H}$ *be any morphism between connected signed, colored graphs satisfying axiom 1. If each of axioms 2, 3, and 4 is obeyed by at least one of* $\mathcal{G}$ *or* $\mathcal{H}$*, and* $\mathcal{G}$ *obeys axioms 5 and 6, then* ϕ *is an isomorphism between standard dual equivalence graphs.*

Proof: It may be helpful to recall that morphisms only exist between signed, colored graphs of the same type. By Lemma A.3.10, both signed, colored graphs of the morphism satisfy all of the dual equivalence axioms. By Theorem A.3.8, the image is isomorphic to some $\mathcal{G}_\lambda$. By Lemma A.3.10, ϕ is an isomorphism. □

Corollary A.3.16. *Any morphism from a connected dual equivalence graph to a connected signed colored graph satisfying axiom 1 is an isomorphism.*

Proof: This is just a special case of Corollary A.3.15. □

Corollary A.3.17. *Two* (n, N)*-dual equivalence graphs are isomporphic if and only if they have same maximal partition type signature.*

Proof: This follows from Theorem A.3.8 and the fact that the isomorphism type of a standard dual equivalence graph can be read off of its maximal partition type signature, as demonstrated in the proof of Proposition A.3.9. □

A.4 Conclusions

At this point, the power of the main result may not be entirely obvious, and it is worth a few words to hint at its future uses. As mentioned before, both LLT polynomials and Macdonald polynomials may be expressed in terms of quasi-symmetric functions using signatures of tableaux. In the upcoming paper, Assaf [2013], a method is given to turn these tableaux into a graph with connected dual equivalent graphs as connected components. By counting these components we may then calculate an expression for LLT and Macdonald polynomials as a sum of Schur functions. A combinatorial proof that this, Schur positivity, is even possible has never been published, though an algebraic proof can be found in Haiman [2001].

It also appears likely that dual equivalence graphs can be used to prove the Schur positivity of other functions. As a combinatorial object, they offer many opportunities for enumeration, generalization, and specialization. In this regard, the purpose of this paper is to free up future research by providing a firm foundation for dual equivalence graphs.

BIBLIOGRAPHY

S. H. Assaf. *Dual equivalence graphs, ribbon tableaux and Macdonald Polynomials.* PhD thesis, University of California at Berkeley, 2007.

S. H. Assaf. A combinatorial realization of Schur-Weyl duality via crystal graphs and dual equivalence graphs. In *20th Annual International Conference on Formal Power Series and Algebraic Combinatorics (FPSAC 2008)*, Discrete Math. Theor. Comput. Sci. Proc., AJ, pages 141–152. Assoc. Discrete Math. Theor. Comput. Sci., Nancy, 2008.

S. H. Assaf. A generalized major index statistic. *Sém. Lothar. Combin.*, 60, 2008/09.

S. H. Assaf. Dual equivalence graphs I: A combinatorial proof of LLT and Macdonald positivity. *arXiv preprint arXiv:1005.3759*, 2013.

S. H. Assaf. Shifted dual equivalence and Schur P-positivity. *arXiv:1402.2570*, 2014.

S. H. Assaf and S. Billey. Affine dual equivalence and k-Schur functions. *J. Combin.*, 3(3): 343–399, 2012.

S. H. Assaf, N. Bergeron, and F. Sottile. Multiplying Schubert polynomials by Schur polynomials. In *The 5th Mathematical Society of Japan Seasonal Institute*. The 5th Mathematical Society of Japan Seasonal Institute, 2012. (manuscript).

J. Bandlow. *Combinatorics of Macdonald polynomials and extensions.* PhD thesis, University of California, San Diego, 2007.

L. J. Billera, S. K. Hsiao, and S. van Willigenburg. Peak quasisymmetric functions and Eulerian enumeration. *Adv. Math.*, 176(2):248–276, 2003.

S. Billey. Transition equations for isotropic flag manifolds. *Discrete mathematics*, 193(1): 69–84, 1998.

S. Billey and M. Haiman. Schubert polynomials for the classical groups. *J. Amer. Math. Soc.*, 8:443–482, 1995.

S. Billey, Z. Hamaker, A. Roberts, and B. Young. Coxeter-Knuth graphs and a signed Little map for type B reduced words. Manuscript.

A. Björner and F. Brenti. *Combinatorics of Coxeter groups*, volume 231 of *Graduate Texts in Mathematics*. Springer, New York, 2005.

A. Borodin and I. Corwin. Macdonald processes. *Probability Theory and Related Fields*, pages 1–176, 2011.

A. Borodin and V. Gorin. Lectures on integrable probability. *arXiv preprint arXiv:1212.3351*, 2012.

J. O. Carbonara. A combinatorial interpretation of the inverse t-Kostka matrix. *Discrete Math.*, 193(1-3):117–145, 1998.

C. Carré and B. Leclerc. Splitting the square of a Schur function into its symmetric and antisymmetric parts. *J. Alg. Combin.*, 4(3), 1995.

C. Chevalley. Sur les Décompositions Cellulaires des Espaces G/B. *Proceedings of Symposia in Pure Mathematics*, 56(1), 1994.

M. Chmutov. Type A Molecules are Kazhdan-Lusztig. *ArXiv e-prints*, 2013.

A. J. Dalal and J. Morse. The ABC's of affine Grassmannians and Hall-Littlewood polynomials. *DMTCS Proceedings*, (01):935–946, 2012.

P. Edelman and C. Greene. Balanced tableaux. *Advances in Mathematics*, 63(1):42–99, 1987.

S. Fishel. Statistics for special q,t-Kostka polynomials. *Proc. Amer. Math. Soc*, 123(10): 2961–2969, 1995.

S. Fomin and A. N. Kirillov. Combinatorial B_n analogues of Schubert polynomials. *Trans. of AMS*, 348:3591–3620, 1996.

W. Fulton. *Young Tableaux, with Applications to Representation Theory and Geometry.* London Mathematical Society Student Texts 35, Cambridge University Press, Cambridge, 1997.

A. Gadde, L. Rastelli, S. S. Razamat, and W. Yan. Gauge theories and Macdonald polynomials. *Comm. Math. Phys.*, 319(1):147–193, 2013.

A. Garsia. *The Saga of Reduced Factorizations of Elements of the Symmetric Group.* Laboratoire de combinatoire et d'informatique mathématique, 2002.

A. M. Garsia and M. Haiman. Factorizations of Pieri rules for Macdonald polynomials. *Discrete Math.*, 139(1-3):219–256, 1995.

I. M. Gessel. Multipartite P-partitions and inner products of skew Schur functions. In *Combinatorics and algebra (Boulder, Colo., 1983)*, volume 34 of *Contemp. Math.*, pages 289–317. Amer. Math. Soc., Providence, RI, 1984.

C. Godsil and G. Royle. *Algebraic Graph Theory*, volume 207 of *Graduate Texts in Mathematics.* Springer-Verlag, New York, 2004.

I. Grojnowski and M. Haiman. Affine Hecke algebras and positivity of LLT and Macdonald polynomials. *Preprint*, 2007.

J. Haglund. A combinatorial model for the Macdonald polynomials. *Proc. Natl. Acad. Sci. USA*, 101(46):16127–16131, 2004.

J. Haglund. *The q,t-Catalan Numbers and the Space of Diagonal Harmonics*, volume 41 of *University Lecture Series.* American Mathematical Society, Providence, RI, 2008.

J. Haglund, M. Haiman, and N. Loehr. A combinatorial formula for Macdonald polynomials. *J. Amer. Math. Soc*, 18(3), 2005a.

J. Haglund, M. Haiman, N. Loehr, J. B. Remmel, and A. Ulyanov. A combinatorial formula for the character of the diagonal coinvariants. *Duke Math. J*, 126(2):195–232, 2005b.

M. Haiman. Hilbert schemes, polygraphs and the Macdonald positivity conjecture. *Journal of the American Mathematical Society*, 14(4):941–1006, 2001.

M. D. Haiman. On mixed insertion, symmetry, and shifted young tableaux. *Journal of Combinatorial Theory, Series A*, 50(2):196–225, 1989.

M. D. Haiman. Dual equivalence with applications, including a conjecture of Proctor. *Discrete Mathematics*, 99(1):79–113, 1992.

P. Hall. The algebra of partitions. In *Proc. 4th Canadian Math. Congress*, pages 147–159, 1957.

Z. Hamaker and B. Young. Relating Edelman-Greene insertion to the Little map. *Journal of Algebraic Combinatorics*, 2013. accepted pending minor revisions.

D. E. Knuth. Permutations, matrices, and generalized Young tableaux. *Pacific J. Math.*, 34: 709–727, 1970.

A. Knutson, T. Lam, and D. E. Speyer. Positroid varieties: juggling and geometry. *Compos. Math.*, 149(10):1710–1752, 2013.

W. Kraskiewcz. Reduced decompositions in hyperoctahedral groups. *Comptes rendus de l'Académie des sciences. Série 1, Mathématique*, 309(16):903–907, 1989.

W. Kraśkiewicz. Reduced decompositions in hyperoctahedral groups. *CR Acad. Sci. Paris Sér. I Math*, 309:903–904, 1989.

W. Kraśkiewicz. Reduced decompositions in Weyl groups. *European Journal of Combinatorics*, 16:293–313, 1995.

T. Lam. Stanley symmetric functions and Peterson algebras. *Arxiv preprint arXiv:1007.2871*, 2010.

T. Lam and M. Shimozono. A Little bijection for affine Stanley symmetric functions. *Sém. Lothar. Combin.*, 54A, 2005/07.

T. K. Lam. *B and D analogues of stable Schubert polynomials and related insertion algorithms*. PhD thesis, Massachusetts Institute of Technology, 1995.

L. Lapointe and J. Morse. Tableaux statistics for two part Macdonald polynomials. *Alg. Comb. and Quantum Groups*, pages 61–84, 2003.

A. Lascoux and M.-P. Schützenberger. Sur une conjecture de H. O. Foulkes. *C. R. Acad. Sci. Paris Sér. A-B*, 286(7):A323–A324, 1978.

A. Lascoux and M.-P. Schützenberger. Polynômes de Schubert. *Comptes Rendus des Séances de l'Académie des Sciences. Série I. Mathématique*, 294:447–450, 1982.

A. Lascoux, B. Leclerc, and J.-Y. Thibon. Ribbon tableaux, Hall-Littlewood functions, quantum affine algebras, and unipotent varieties. *J. Math. Phys*, 38(2):1041–1068, 1997.

C. Lenart and M. Shimozono. Equivariant K-Chevalley Rules for Kac-Moody Flag Manifolds. *ArXiv e-prints*, Mar. 2012.

D. Little. Combinatorial aspects of the Lascoux-Schützenberger tree. *Advances in Mathematics*, 174(2):236–253, 2003.

D. E. Littlewood. On certain symmetric functions. *Proc. London Math. Soc.*, 3(1):485–498, 1961.

N. A. Loehr and G. S. Warrington. Quasisymmetric expansion of Schur function plethysms. *Proc. Amer. Math. Soc*, 140(4):1159–1171, 2012.

N. A. Loehr, L. G. Serrano, and G. S. Warrington. Transition matrices for symmetric and quasisymmetric Hall-Littlewood polynomials. *J. Comb. Theory Ser. A*, 120(8):1996–2019, 2013.

I. G. Macdonald. A new class of symmetric functions. *Actes du* 20^e *Sém. Lothar.*, 372: 131–171, 1988.

I. G. Macdonald. *Symmetric functions and Hall polynomials.* Oxford Mathematical Monographs. The Clarendon Press Oxford University Press, New York, second edition, 1995.

A. Nakayashiki and Y. Yamada. Kostka polynomials and energy functions in solvable lattice models. *Selecta Mathematica*, 3(4):547–599, 1997.

J. B. Remmel and R. Whitney. Multiplying Schur functions. *J. Algorithms*, 5(4):471–487, 1984.

A. Roberts. Computer proofs: LLT polynomials and dual equivalence graphs. <http://www.math.washington.edu/∼austinis/ Proof_LLTandDEG.sws>, 2012.

A. Roberts. On the Schur expansion of Hall-Littlewood and related polynomials via Yamanouchi words. *arXiv:1404.1036*, 2013a.

A. Roberts. Computer proof: dual equivalence graphs by strict patterns. <http://www.math.washington.edu/∼austinis/Proof_DEG_by_strict_patterns.pdf>, 2013b.

A. Roberts. Dual equivalence graphs revisited and the explicit Schur expansion of a family of LLT polynomials. *Journal of Algebraic Combinatorics*, 39(2):389–428, 2014.

B. E. Sagan. Shifted tableaux, Schur q-functions, and a conjecture of R. Stanley. *Journal of Combinatorial Theory, Series A*, 45(1):62–103, 1987.

B. E. Sagan. *The Symmetric Group.* Wadsworth, Inc., Belmont, CA, 1991.

M. Schocker. The peak algebra of the symmetric group revisited. *Adv. Math.*, 192(2):259–309, 2005.

M.-P. Schützenberger. La correspondance de Robinson. In *Combinatoire et représentation du groupe symétrique (Actes Table Ronde CNRS, Univ. Louis-Pasteur Strasbourg, Strasbourg, 1976)*, volume 579, pages 59–113. Lecture Notes in Math., Vol. 579. Springer, 1977.

R. P. Stanley. On the number of reduced decompositions of elements of Coxeter groups. *Eur. J. Comb.*, 5:359–372, 1984.

R. P. Stanley. *Enumerative Combinatorics*, volume 2 of *Cambridge Studies in Advanced Mathematics.* Cambridge University Press, Cambridge, 2001.

R. P. Stanley. *Catalan Addendum.* <www-math.mit.edu/ rstan/ec/catadd.pdf>, 2012a.

R. P. Stanley. *Enumerative Combinatorics*, volume 1 of *Cambridge Studies in Advanced Mathematics.* Cambridge University Press, Cambridge, second edition, 2012b.

E. Steinitz. Zur theorie der Abel'schen gruppen. *Jahresbericht der Deutschen Mathematiker-Vereinigung*, 9(1):80–85, 1901.

J. R. Stembridge. Enriched P-partitions. *Trans. Amer. Math. Soc.*, 349(2):763–788, 1997.

M. A. van Leeuwen. Some bijective correspondences involving domino tableaux. *J. Combin.*, 7(1):R35–R35, 2001.

D. R. Worley. *A theory of shifted Young tableaux.* PhD thesis, Massachusetts Institute of Technology, 1984.

B. G. Wybourne. Symmetry principles in atomic spectroscopy. *Le Journal de Physique Colloques*, 31(C4):C4–33, 1970.

M. Zabrocki. A Macdonald vertex operator and standard tableaux statistics for the two-column (q, t)-Kostka coefficients. *J. Combin.*, 5:661–706, 1998.

M. Zabrocki. Positivity for special cases of (q, t)-Kostka coefficients and standard tableaux statistics. *Electron. J. of Combin.*, 6(R41):2, 1999.

Zeidler, Eberhard. *Quantum field theory. III. Gauge theory: A bridge between mathematicians and physicists.* Springer, Heidelberg, 2011.